CDF/DOC/TRACKING/CDFR/5701
September 4, 2001

SiliMon User's Manual

Ronan McNulty [1], Tara Shears [1],
Stephane Tourneur [2]

---

[1] University of Liverpool
[2] University of California at Davis

Contents

1 Introduction

SiliMon (**Sili**con **Mon**itoring) is a silicon monitoring program [1] which produces histograms of raw, cluster-based and track-based information. The histograms are designed to allow fast and accurate evaluation of the silicon performance. SiliMon is run as a consumer monitor online, so that any operational problems can be quickly observed. A small set of summary histograms is continuously checked by the CO shift during data-taking. SiliMon is also run as part of the Production executable offline so that high statistics are available to understand the silicon response. Experts access these offline files and more detailed histograms to check the silicon behaviour. All histograms are described in this manual.

In order to provide fast and logical navigation through the thousands of available histograms a browser has been developed. The browser presents histogrammed information in a series of screens of increasing detail. Buttons are used to choose the area of the detector to study. Reference histograms can be overlaid at any stage so that deviations from normal behaviour become apparent. In this way problems can be quickly and efficiently isolated. This manual includes a user's guide to operating the browser.

The manual is organised as follows. Section 2 contains a description of all histograms produced by SiliMon. CO histograms are listed first, with guidelines for interpretation, expert (more detailed) histograms follow. The SiliMon browser itself is described in section 3. Finally, section 4 lists where to find the online and offline histograms.

2 SiliMon Histograms

SiliMon produces histograms which allow the performance of the silicon system to be monitored and evaluated. These are organised by information type: raw information histograms; cluster information histograms; track information histograms. Each type of information consists of a "tree" of histograms of differing levels of detail, from summary histograms which are produced for non-expert (CO shift) consumption to more detailed histograms (eg. per layer, per barrel, per ladder and per chip histograms) which are available for experts.

SiliMon is run online during data-taking by the CO shift crew. Nine canvases of raw, cluster and track-based summary histograms are available to them to monitor the silicon system performance. In section 2.1 these canvasses are documented and simple guidelines for interpretation given. In order to view these histograms online, the Consumer display can be used. Instructions for running the display are given in section 4.1.

SiliMon is run offline after data-taking by the Production data reprocessing executable. All histograms are available with high statistics in Production output files. Section 2.2 documents all the histograms which are produced by SiliMon, with examples. These files cannot be viewed by the Consumer display. In order to browse these files efficiently and intuitively, a browser has been developed to navigate through the large amount of available information. Instructions for using this browser are given in section 3.

2.1 Canvasses for the CO shift

This section describes the summary canvasses (screens of histograms) which the CO shift monitors. These histograms are produced online, and can be browsed with the Consumer Display. An example screen shot of the Consumer Display is shown in figure 1.

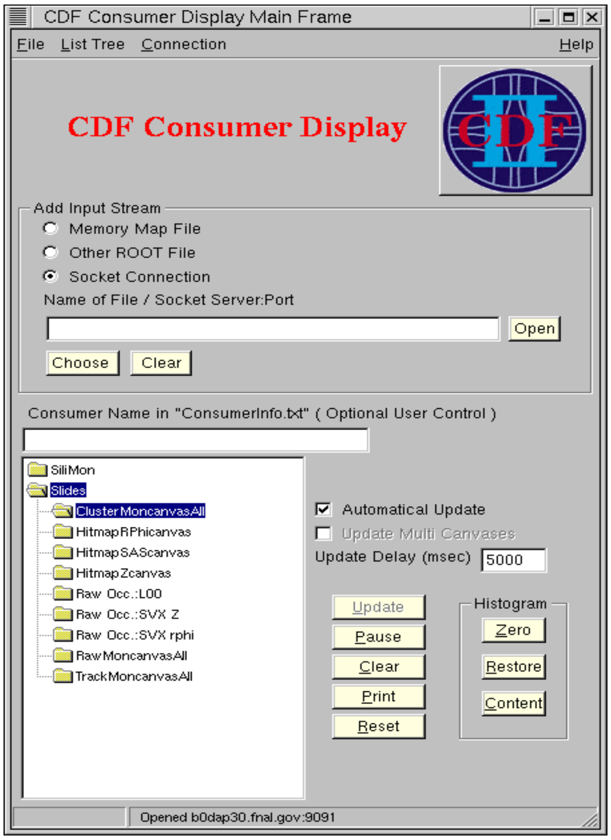

○ Memory Map File
○ Other ROOT File
● Socket Connection
Name of File / Socket Server:Port

Consumer Name in "ConsumerInfo.txt" (Optional User Control)

☑ Automatical Update
☐ Update Multi Canvases
Update Delay (msec) 5000

Opened b0dap30.fnal.gov:9091

Figure 1: Consumer Display screen showing SiliMon CO canvases.

Note that there are nine canvasses listed under the SiliMon “Slides” folder:

- ClusterMoncanvasAll - general cluster information
- HitmapRPhicanvas - axial cluster hitmap
- HitmapSAScanvas - stereo cluster hitmap
- HitmapZcanvas - 90° cluster hitmap
- RawMoncanvasAll - raw ADC count information
- Raw Occ.:L00 - raw hitmap for layer 00
- Raw Occ.:SVX Z - raw hitmap for SVX and ISL stereo and 90° wafers

- Raw Occ.:SVX Z - raw hitmap for SVX and ISL axial wafers
- TrackMoncanvasAll - general track information

These canvasses are monitored by the CO shift. They will be described, together with some simple guidelines for interpretation, in the following subsections in the order in which they appear in the SiliMon Slides list. The guidelines are taken from the SiliMon CO documentation on the web [1]. The web documentation should be checked for the most recent guidelines, documentation and reference canvasses. In this manual, the data shown in the example canvasses come from run 119636.

2.1.1 ClusterMoncanvasAll - General cluster information

The following histograms are displayed on this CO canvas:

- Number of clusters / event for all clusters
- Fraction of clusters / event associated to tracks
- Number of strips / cluster for all clusters
- Number of strips / cluster for clusters on tracks
- Mean number of ADC counts / cluster for all clusters
- Mean number of ADC counts / cluster for clusters on tracks

An example canvas is shown in figure 2. The histograms contained within the canvas summarise general cluster properties for both all clusters in the event, and for clusters which have been associated to tracks. The average number of clusters per layer and side, the average number of strips in a cluster per layer and side, and the average number of ADC counts contained in a cluster per layer and side are shown. The x-axis of each histogram is layer number, with axial wafers occupying the left hand side of the plot, and stereo and 90° wafers occupying the right hand side. Note that as layer 00 only contains one (axial) side, layer 0 on the right hand side of the histograms will always be empty.

The standard CDF clustering code is run to reconstruct silicon clusters. Note that this requires appropriate input calibration files for both pedestals and dead/hot channels in order to give sensible output. If an inappropriate calibration file is used then the resultant clusters may be meaningless. This level of histogrammed information therefore not only tests the silicon response, but the quality of the offline reconstruction.

Guidelines

• The number of clusters / event is typically 80 - 100. This is a function of the amount of biased, operational silicon, and the beam conditions. Greater luminosity will increase the number of clusters, as there will be a greater probability of physics interactions. In this case, the number of clusters on all layers should increase. Bad beam backgrounds (eg. beam halo) may also increase the number of clusters, however here one would expect the innermost layers of silicon to be most affected.

• The number of strips / cluster is typically about 3 (all clusters)/ about 2 (clusters on tracks) for r-ϕ information, and 3-4 (all clusters)/ 2-3 (clusters on tracks) for z information. Anything much larger than this implies noise or a bad calibration.

• The mean charge / cluster is typically 50-70 ADC counts (all clusters) / 30-60 ADC counts (clusters on tracks). Anything much larger than this may be noise or a bad calibration. If so, this should also be noticeable in the number of strips / cluster histograms.

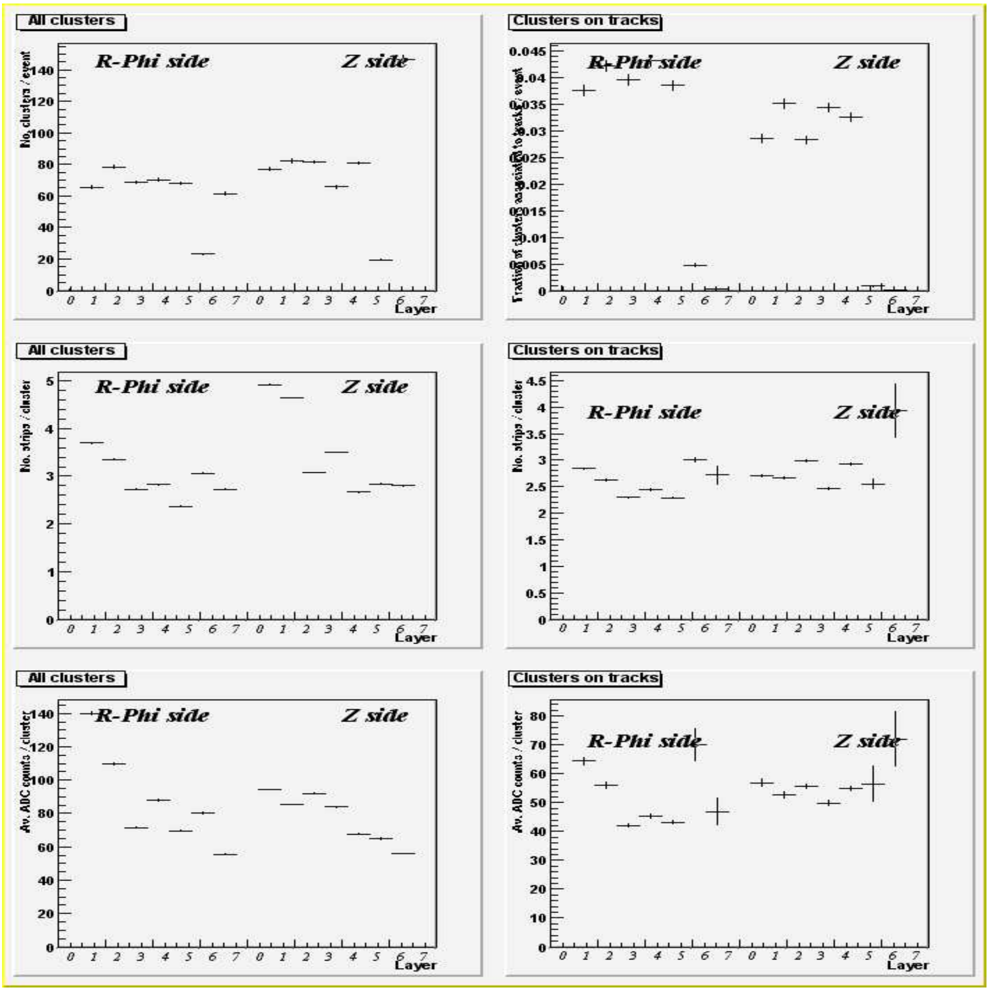

Figure 2: General cluster properties: (top) number of clusters / event; (centre) number of strips / cluster; (bottom) number of ADC counts / cluster. In all cases histograms are shown for all clusters in the event (left), and for clusters associated to tracks (right).

2.1.2 HitmapRPhicanvas - Cluster hitmap for axial clusters

This canvas (figure 3), contains nine histograms which show axial cluster hitmaps. The weighted mean cluster position is plotted in x-y, for each barrel separately. Histograms are shown for all clusters on the top row, and for clusters associated to tracks on the centre row. Once a track has been associated to a hit it is possible to calculate the r-z position of the cluster. This view is shown on the bottom row of the canvas.

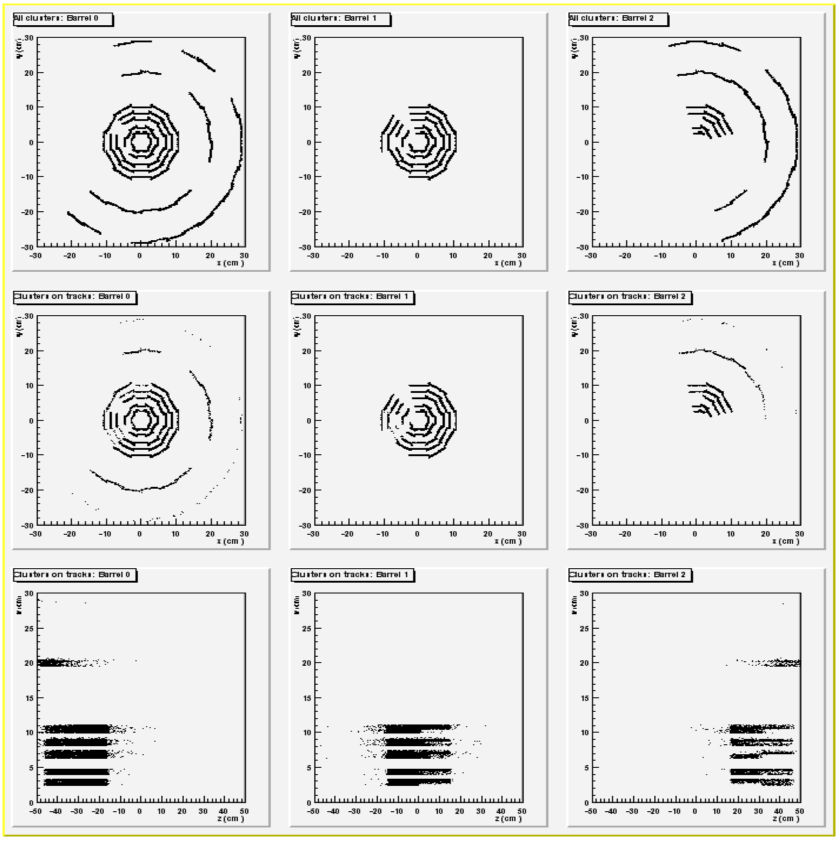

Figure 3: X-Y hitmap of axial clusters, for all clusters (top row), and clusters associated to tracks (centre row). R-Z hitmap for clusters associated to tracks is shown on the bottom row. In all cases the left column shows clusters in barrel 0 of the detector, the centre column clusters in barrel 1, and the right column clusters in barrel 2.

Guidelines

- If no tracks are found in the event, only the top three histograms will be filled.
- Any gaps in these distributions denote dead/unbiased areas in the detector.

2.1.3 HitmapSAScanvas: Cluster hitmap for stereo clusters

This canvas (figure 4), contains nine histograms which show stereo cluster hitmaps. There are four stereo layers in the silicon system; two in SVX (layers 3 and 5), and two in ISL. The weighted mean cluster position is plotted in x-y, for each barrel separately. Histograms are shown for all clusters on the top row, and for clusters associated to tracks on the centre row. Once a track has been associated to a hit it is possible to calculate the r-z position of the cluster. This view is shown on the bottom row of the canvas.

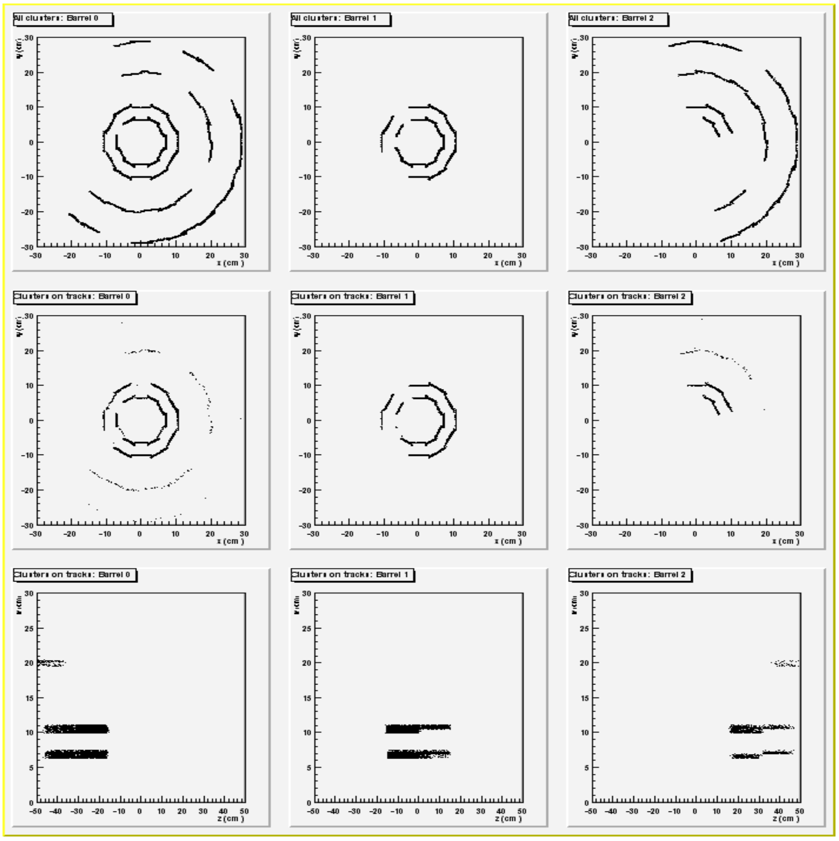

Figure 4: X-Y hitmap of stereo clusters, for all clusters (top row), and clusters associated to tracks (centre row). R-Z hitmap for clusters associated to tracks is shown on the bottom row. In all cases the left column shows clusters in barrel 0 of the detector, the centre column clusters in barrel 1, and the right column clusters in barrel 2.

Guidelines

- If no tracks are found in the event, only the top three histograms will be filled.
- Any gaps in these distributions denote dead/unbiased areas in the detector.

2.1.4 HitmapZcanvas: Cluster hitmap for 90° clusters

This canvas (figure 5), contains nine histograms which show 90° cluster hitmaps. There are three 90° layers in SVX; layer 1, layer 2 and layer 4. The weighted mean cluster position is plotted in r-z, for each barrel separately. Histograms are shown for all clusters on the top row, and for clusters associated to tracks on the centre row. Once a track has been associated to a hit it is possible to calculate the x-y position of the cluster. This view is shown on the bottom row of the canvas.

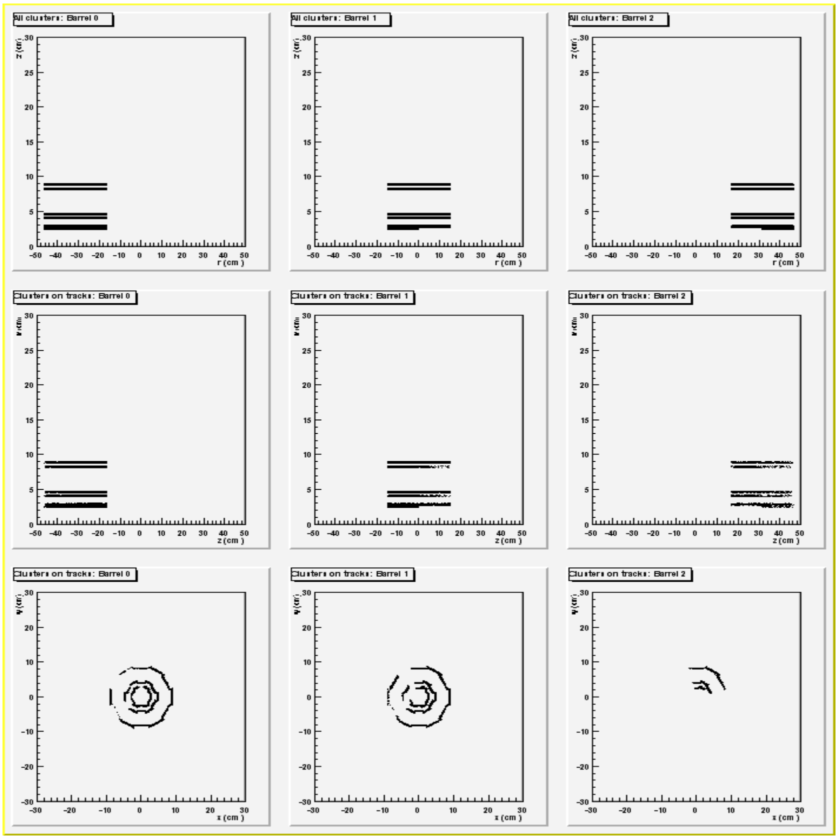

Figure 5: R-Z hitmap of 90° clusters, for all clusters (top row), and clusters associated to tracks (centre row). X-Y hitmap for clusters associated to tracks is shown on the bottom row. In all cases the left column shows clusters in barrel 0 of the detector, the centre column clusters in barrel 1, and the right column clusters in barrel 2.

Guidelines

- If no tracks are found in the event, only the top three histograms will be filled.
- Any gaps in these distributions denote dead/unbiased areas in the detector.

2.1.5 RawMoncanvasAll - Raw ADC count monitoring

The following histograms are displayed on the CO canvas:

- Mean ADC count / event, for layer 00
- Mean ADC count / event, for layer 1
- Mean ADC count / event, for layer 2
- Mean ADC count / event, for layer 3
- Mean ADC count / event, for layer 4
- Mean ADC count / event, for layer 5
- Mean ADC count / event, for layer 6
- Mean ADC count / event, for layer 7

An example canvas is shown in figure 6. These histograms summarise the ADC counts read out from the detector in each layer. Each histogram displays the mean ADC count per two or three chips (depending on the number of ladder segments). In all cases the x-axis is arranged in order of increasing side number, within each side by barrel, within each side by phi wedge, and within each phi wedge by chip number.

Guidelines

• DPS (Dynamic Pedestal Subtraction) is switched on for normal silicon operation. During this mode of running, pedestals are automatically subtracted at the silicon front end. The subtracted pedestals will be read out. In this case the typical mean ADC count / chip should be close to zero. If the mean ADC count / chip is significantly different to this, this could mean that the pedestal download has failed, or that an inappropriate file was used.

• During DPS off running, the full pedestal is read out. This is typically 30 - 100 ADC counts / chip. If the mean ADC count / chip is significantly different to this, this could mean that a chip is dead, unbiased, or very noisy.

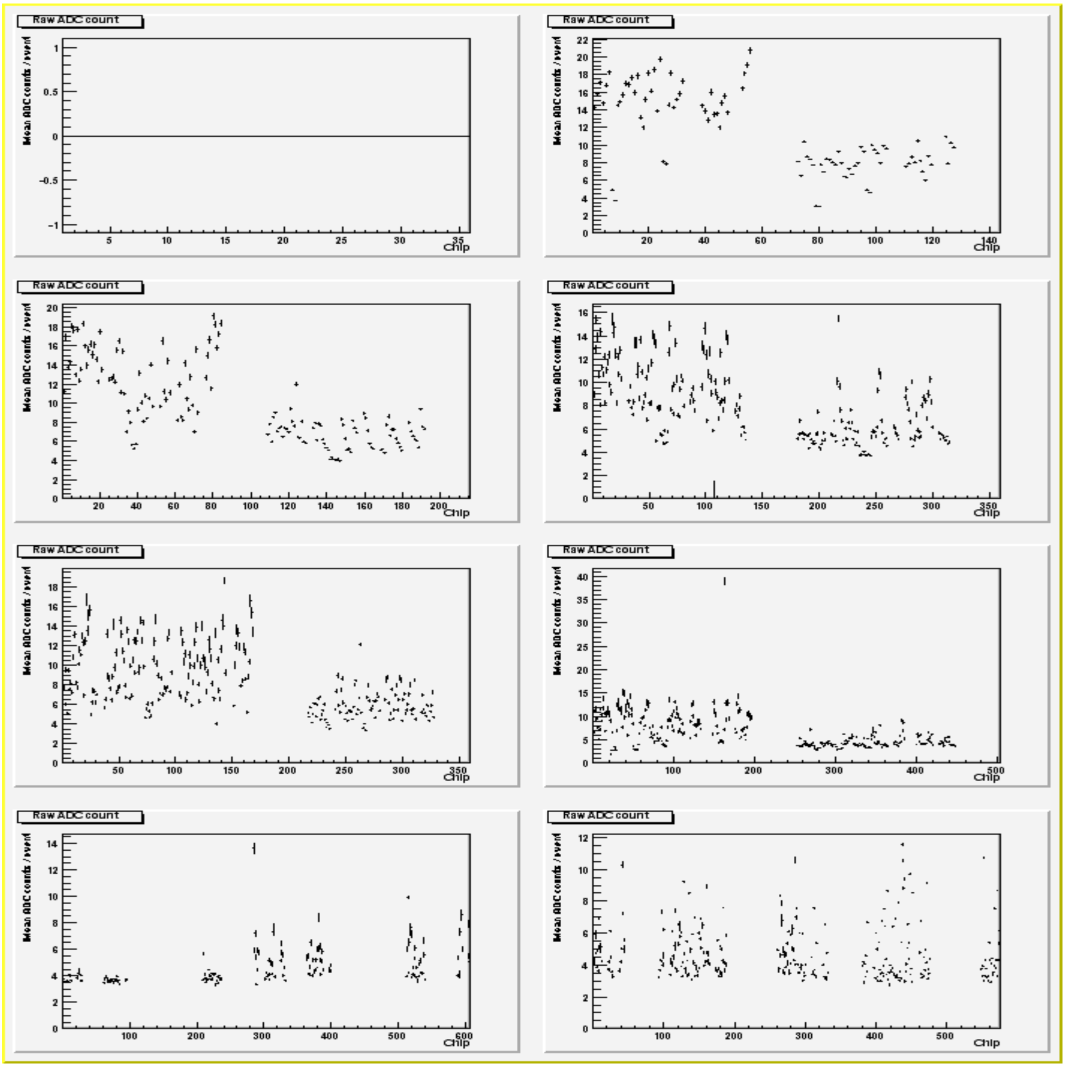

Figure 6: Mean ADC counts / event, for (from left to right) layer 00, SVX layer 1, SVX layer 2, SVX layer 3, SVX layer 4, SVX layer 5, innermost ISL layer (layer 6), outermost ISL layer (layer 7).

2.1.6 Raw Occ.: L00 - Raw hitmap in layer 00

This canvas (figure 7), contains one histogram which summarises the raw hit occupancy in layer 00. A "rolled-out" view of the detector is shown; by increasing phiwedge on the y-axis, and by increasing chip number (ie. left to right) on the x-axis. In layer 00 phi wedges are wide (2 chips per halfladder) and narrow (1 chip per halfladder) alternately. Non-physical areas of the detector are marked in blue and need not be checked. Each bin is filled with the average ADC count on the corresponding chip. As this is a two-dimensional plot, it is the relative areas of each filled bin that are important.

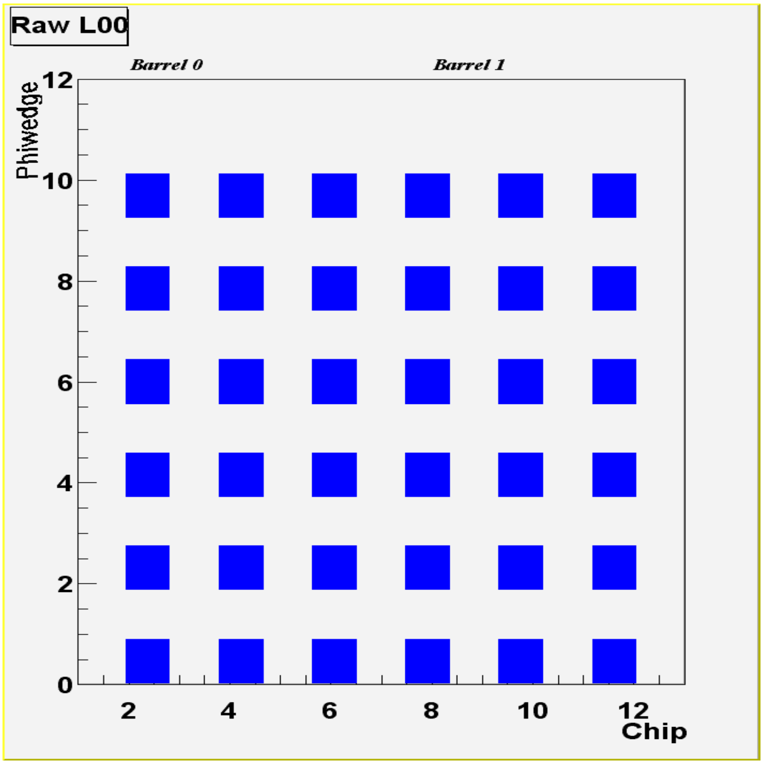

Figure 7: Raw hit occupancy for layer 00. The y axis denotes phi wedge, and the x-axis chip number. Chips in barrel 0 occupy the left hand side of the plot, and chips in barrel 1 the right hand side. The chips are arranged in order of increasing half ladder. There are three half ladders in each barrel. The blue areas denote areas of the detector where there are no chips.

Guidelines

- If a bin is empty, the corresponding chip is dead, disconnected, or unbiased.
- The area of the box is proportional to the number of counts recorded. If one box is much larger than its neighbours, the chip it corresponds to is noisy.

2.1.7 Raw Occ.: SVX Z - Raw hitmap in stereo and 90° SVX and ISL wafers

This canvas (figure 8), contains one histogram which summarises the raw hit occupancy in SVX and ISL stereo and 90° wafers. A "rolled-out" view of the detector is shown; by increasing phiwedge on the y-axis, and by increasing chip number (ie. left to right) on the x-axis. There are between 2 and 7 chips per halfladder in these detectors, depending on their spatial extent. Non-physical areas of the detector are marked in blue and need not be checked. Each bin is filled with the average ADC count on the corresponding chip. As this is a two-dimensional plot, it is the relative areas of each filled bin that are important.

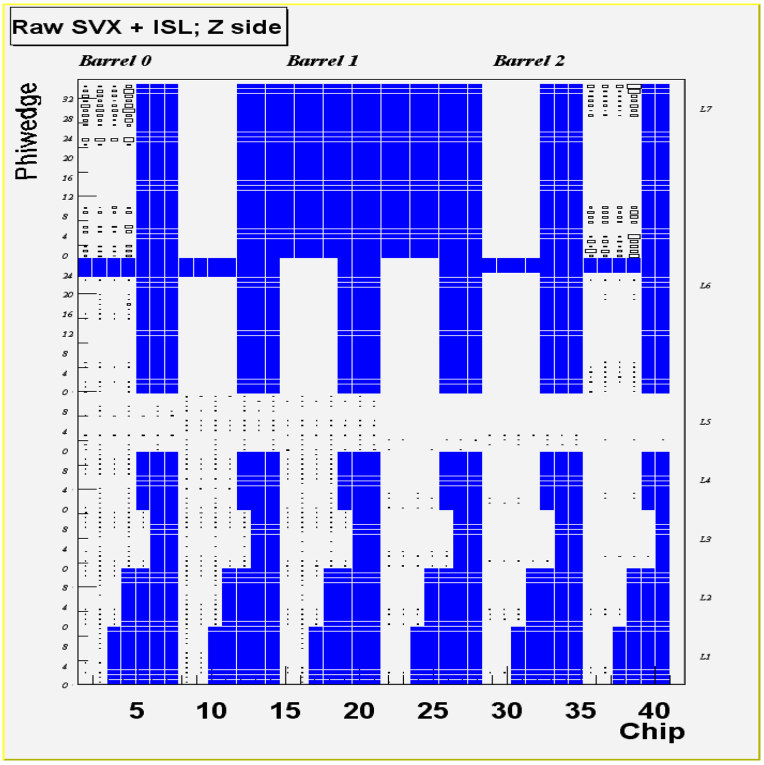

Figure 8: Raw hit occupancy in SVX and ISL for z side silicon. The y axis denotes phi wedge, in order of increasing layer number, and the x-axis chip number. Chips in barrel 0 occupy the left hand side of the plot, chips in barrel 1 the middle, and chips in barrel 2 the right hand side. The chips are arranged in order of increasing half ladder. There are two half ladders in each barrel. The blue areas denote areas of the detector where there are no chips.

Guidelines

- If a bin is empty, the corresponding chip is dead, disconnected, or unbiased.
- The area of the box is proportional to the number of counts recorded. If one box is much larger than its neighbours, the chip it corresponds to is noisy.

2.1.8 Raw Occ.: SVX rphi - Raw hitmap in axial SVX and ISL wafers

This canvas (figure 9), contains one histogram which summarises the raw hit occupancy in SVX and ISL axial wafers. A "rolled-out" view of the detector is shown; by increasing phiwedge on the y-axis, and by increasing chip number (ie. left to right) on the x-axis. There are between 2 and 7 chips per halfladder in these detectors, depending on their spatial extent. Non-physical areas of the detector are marked in blue and need not be checked. Each bin is filled with the average ADC count on the corresponding chip. As this is a two-dimensional plot, it is the relative areas of each filled bin that are important.

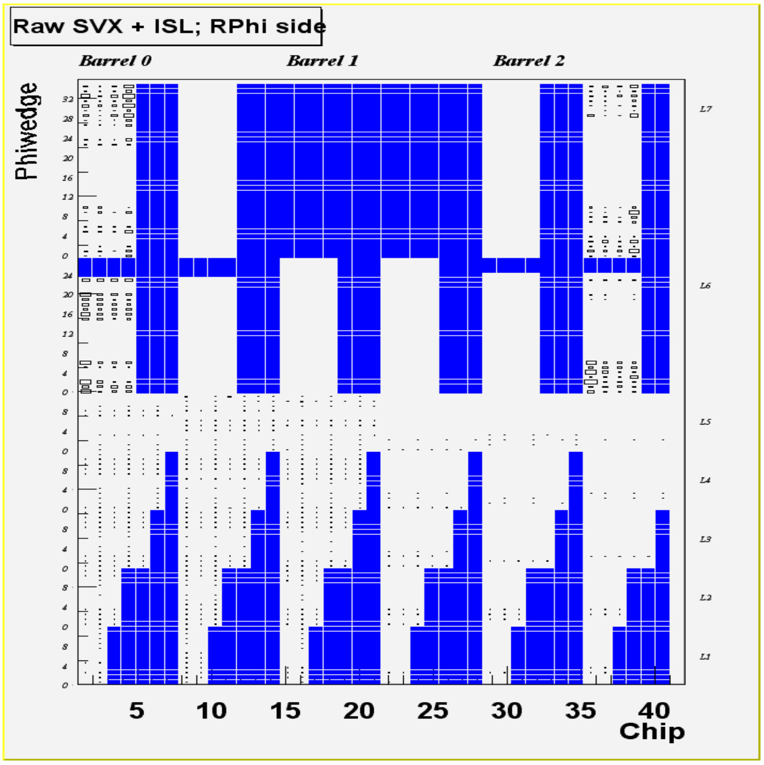

Figure 9: Raw hit occupancy in SVX and ISL for axial silicon wafers. The y axis denotes phi wedge, in order of increasing layer number, and the x-axis chip number. Chips in barrel 0 occupy the left hand side of the plot, chips in barrel 1 the middle, and chips in barrel 2 the right hand side. The chips are arranged in order of increasing half ladder. There are two half ladders in each barrel. The blue areas denote areas of the detector where there are no chips.

Guidelines

- If a bin is empty, the corresponding chip is dead, disconnected, or unbiased.
- The area of the box is proportional to the number of counts recorded. If one box is much larger than its neighbours, the chip it corresponds to is noisy.

2.1.9 TrackMoncanvasAll - General track information

The following histograms are displayed on this CO canvas:

- No. tracks / event, for all tracks. Note the highest x bin counts the overflows
- No. tracks with >=2 r-ϕ silicon hits / events, for all tracks. Note the highest x bin counts the overflows
- No. silicon hits on track, for all tracks
- Track Pt, for all tracks (solid), and tracks with silicon hits (points)
- Track charge, for all tracks (solid), and tracks with silicon hits (points)
- Track hitmap eta vs. ϕ for all tracks
- ϕ distribution of tracks, for all tracks (solid), and tracks with silicon hits (points)
- η distribution of tracks, for all tracks (solid), and tracks with silicon hits (points)
- Track hitmap eta vs. ϕ for tracks with silicon hits.

An example canvas is shown in figure 10. The histograms contained within the canvas summarise general track-based information in the event. The numbers of tracks found in the event, the hitmap in $\eta - \phi$ space where the tracks have their point of closest approach, the separate η and ϕ distributions, the transverse momentum (Pt) and charge distributions are shown. Histograms are shown for all the tracks in the event, and for those tracks with silicon hits associated. A track is defined as having silicon hits associated if at least three axial hits are included in the fit.

The standard CDF track reconstruction code is run to reconstruct tracks. Note that if an inappropriate alignment file or other calibration file are used, the output may not be sensible. This level of histogrammed information therefore not only tests the silicon response, but the quality of the offline reconstruction and calibration.

Guidelines

• If no tracks with silicon hits have been found, but the silicon is biased and operational and clusters look sensible, there could be a problem with alignment between the COT and the silicon system. Alternatively, there could be a problem with the reconstruction geometry which affects the tracking performance.

• If the track Pt is softer for tracks with silicon hits, and there are more positively charged tracks with silicon hits than there are for all tracks, this could be a sign of beam halo.

• If the track hitmap distribution shows any depleted regions, but the silicon clusters look sensible, this may point to a problem with the COT. Inefficiencies here will lead to inefficiencies in tracking, as outside-in algorithms are usually run by default.

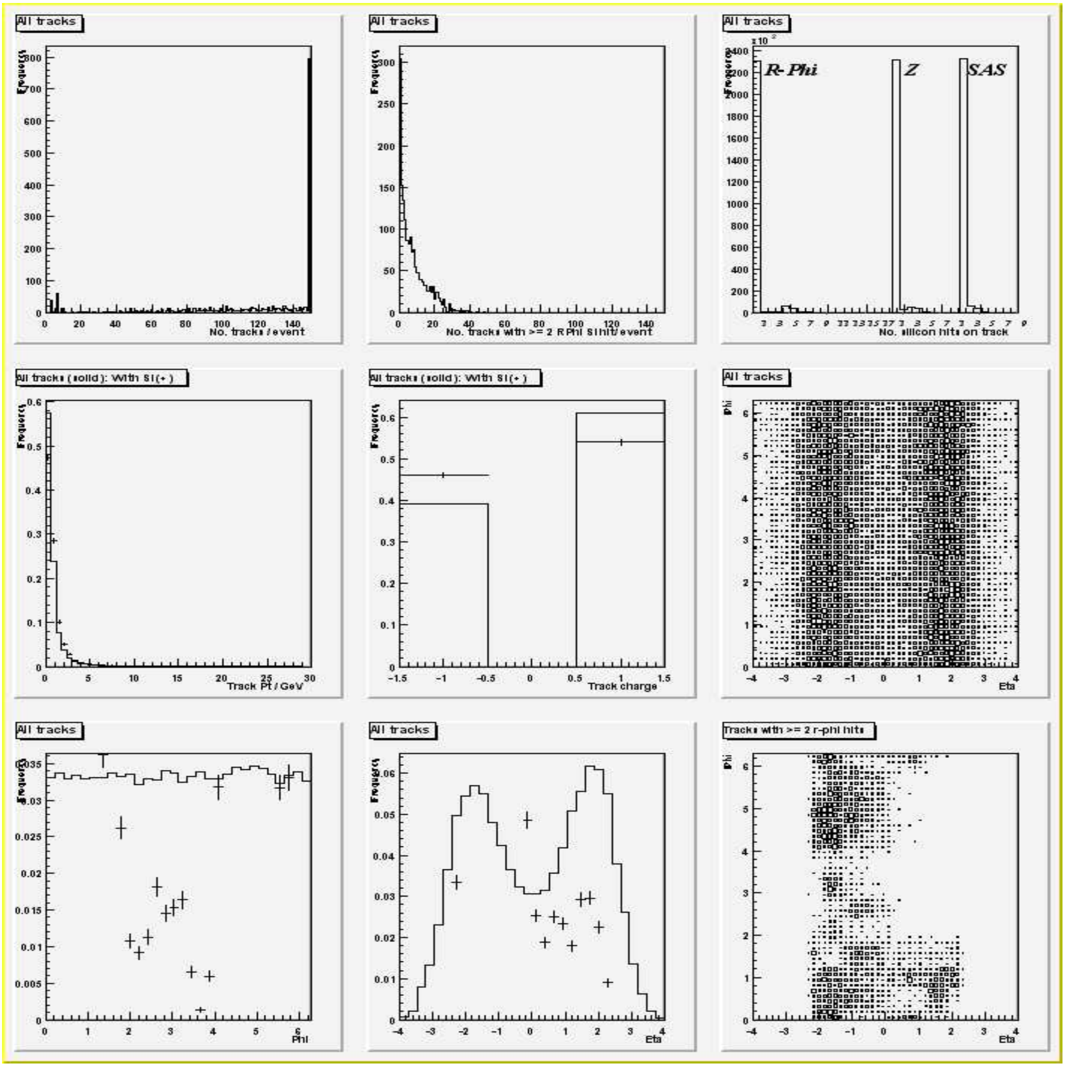

Figure 10: Histograms for general track monitoring (from left to right): track multiplicity and silicon hit content (top row), track Pt, charge for all tracks and tracks with silicon hits, and hitmap for all tracks(centre row), track ϕ, η for all tracks and tracks with silicon hits, and hitmap for tracks with silicon hits (bottom row).

2.2 Histograms for experts

SiliMon produces some 3175 histograms. These can be sorted into ten categories, which will be described in the following subsections:

- Raw ADC count histograms in section 2.2.1
- Cluster occupancy histograms in section 2.2.2
- Cluster extent histograms in section 2.2.3
- Cluster ADC count histograms in section 2.2.4
- Track occupancy histograms in section 2.2.5
- Track Pt histograms in section 2.2.6
- Track charge histograms in section 2.2.7
- Track η and ϕ histograms in section 2.2.8
- Track residual histograms in section 2.2.9
- Track impact parameter and global alignment histograms in section 2.2.10

Each histogram can be updated in one of three modes. We define the following: a snapshot histogram as a histogram whose contents are reset every 50 (2000) events online (offline); an accumulation histogram as a histogram whose contents accumulate throughout a run; a time history histogram as a histogram which is filled with the mean of the corresponding snapshot histogram each time the snapshot histogram is reset. Up to 200 such updates are displayed, with the most recent rightmost. Within each of the following subsections the available histogram types and update modes will be described.

2.2.1 Raw ADC count histograms

Raw ADC counts are read out directly from the silicon front end. No offline processing aside from bank unpacking is performed. Available raw ADC count histograms are listed in table 1. The first three sets of histograms in table 1 have already been described in section 2.1. Histograms showing the raw ADC count per layer are described in section 2.2.1.1, per layer barrel and side in section 2.2.1.2, and per strip in section 2.2.1.3.

Level of detail	Snapshot	Time history	Accumulation
Summary: L00 hitmap	No	No	Yes
Summary: SVX/ISL rϕ hitmap	No	No	Yes
Summary: SVX/ISL rz hitmap	No	No	Yes
ADC count / 2 chips: per layer	Yes	Yes	Yes
ADC count / chip: per layer, barrel and side	Yes	Yes	Yes
ADC count / strip: per chip	No	No	Yes

Table 1: Histograms available for raw ADC count monitoring.

2.2.1.1 Per layer histograms

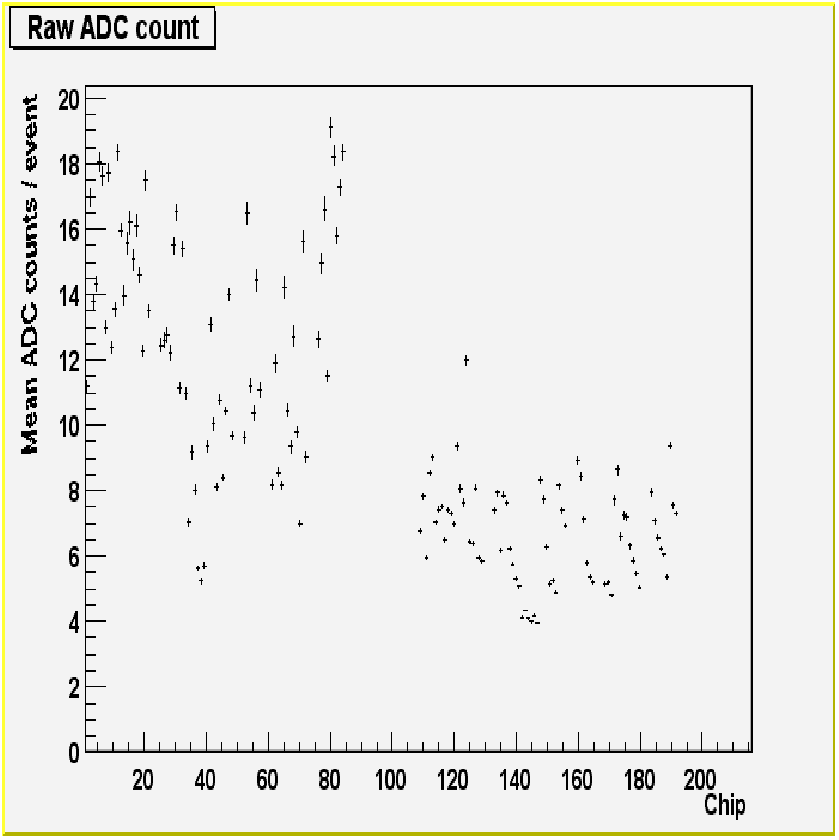

Figure 11: Mean number of ADC counts / chip for layer 2 of the silicon system. Here the x-axis represents chip number (see text for an explanation), and the y-axis gives the mean ADC count recorded on that chip.

There are eight histograms of raw ADC counts per layer, each corresponding to one layer of the silicon. The x-axis is organised in order of increasing chip number. The bins are arranged by side (so the left half of the plot corresponds to axial chips), within each side by barrel, then phiwedge, then chip number. In order that there not be too many bins, each bin contains two chips. The y-axis represents the mean number of ADC counts recorded per event. An example of this type of histogram is shown in figure 11.

2.2.1.2 Per layer barrel and side histograms

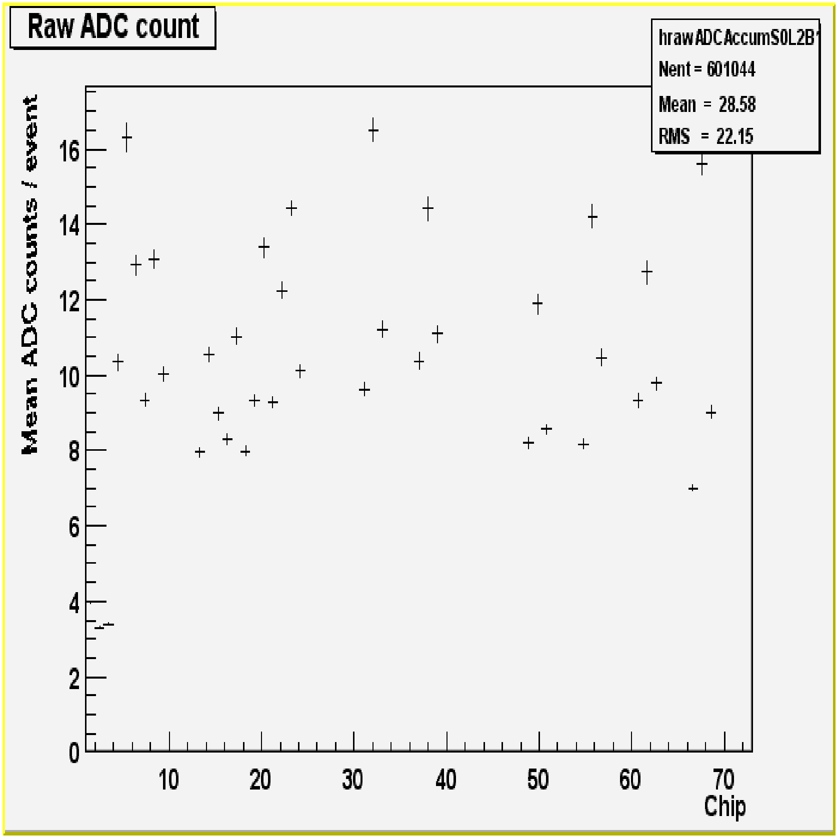

Figure 12: Mean number of ADC counts / chip for a given layer, barrel and side inside the silicon system. Here the x-axis represents chip number (see text for an explanation), and the y-axis gives the mean ADC count recorded on that chip. In this example, data from layer 2, barrel 1, axial wafers is shown.

There are forty-four histograms of raw ADC counts per layer, barrel and side. Each histogram contains information from a given barrel, layer, and side. The x-axis is organised in order of increasing chip number. The chips are arranged in order of increasing phi wedge, within each phiwedge by increasing half ladder number and within each halfladder by chip. Each chip occupies one bin. The y-axis represents the mean ADC count found within a chip. An example histogram is shown in figure 12.

2.2.1.3 Per strip histogram

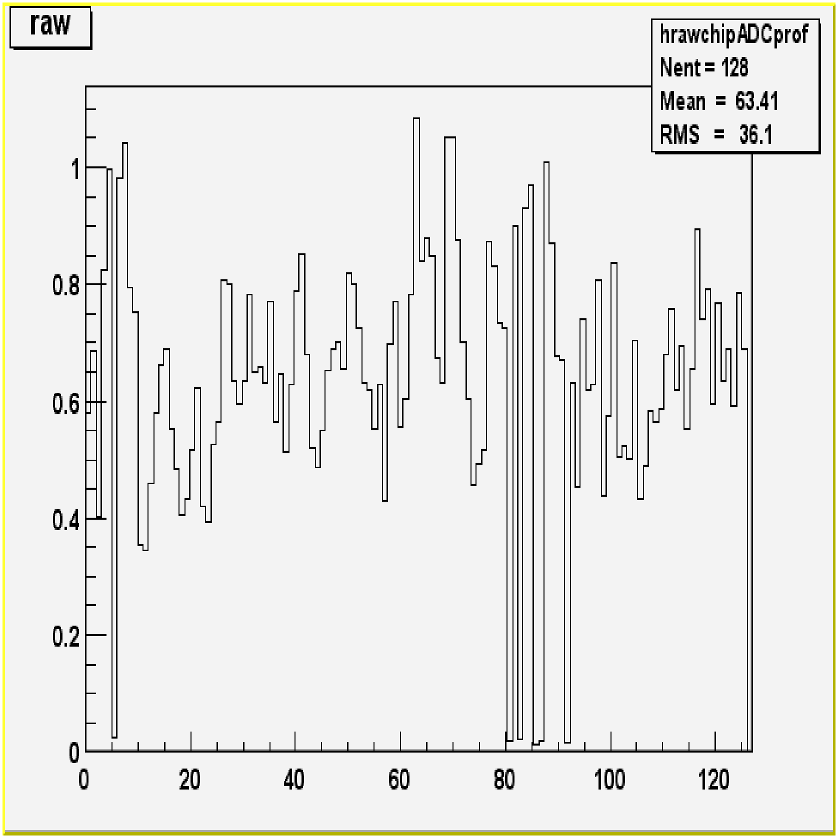

Figure 13: Projected mean ADC count chip histogram from per strip two dimensional histogram. The x-axis represents chip strip number, and the y-axis the mean ADC count recorded.

Raw ADC count information from all 5644 chips in the silicon system is stored in a two dimensional histogram. The x-axis represents strip number on a chip. The y-axis is organised in order of increasing chip number. The chips are in order of increasing layer, within each layer by side, within each side by barrel, within each barrel by phiwedge, within each phiwedge by halfladder and finally within each halfladder in order of increasing strip number. Each chip occupies one bin on the y axis. The z axis represents the average number of ADC counts recorded per event.

Information for individual chips can be obtained by making a projection along the y-axis in ROOT, or by using the SiliMon browser (section 3) where this is done automatically. An example of a projection is shown in figure 13.

2.2.2 Cluster occupancy histograms

Cluster occupancy is measured by observing which silicon strips are associated to a cluster. In order for a cluster to be formed the raw information must be processed by the offline silicon clustering code in conjunction with appropriate calibrations for noisy and dead strips. Available cluster occupancy histograms are listed in table 2. The first four sets of histograms in table 2 have already been described in section 2.1. Histograms showing the cluster occupancy per layer are described in section 2.2.2.1, per layer barrel and side in section 2.2.2.2, and per strip and BE (back end) state in section 2.2.2.3. All histograms with the exception of occupancy / strip per chip and BE (back end) state (all clusters only) are produced both for all clusters in the event, and for clusters associated to tracks. It should be noted that in all cases, all strips in a cluster enter into the histograms.

Level of detail	Snapshot	Time history	Accumulation
Summary: r-ϕ cluster hitmap per barrel	No/Yes	No	Yes/No (*)
Summary: sas cluster hitmap per barrel	No/Yes	No	Yes/No (*)
Summary: r-z cluster hitmap per barrel	No/Yes	No	Yes/No (*)
Summary: No. clusters / event per side and layer	Yes	Yes	Yes
No. clusters / 2 chips: per layer	Yes	Yes	Yes
No. clusters / chip: per layer barrel and side	Yes	Yes	Yes
Occupancy / strip: per chip	No	No	Yes
Occupancy / strip: per chip and BE state	No	No	Yes

Table 2: Histograms available for cluster occupancy monitoring. (*)The summary hitmaps by default are set to fill accumulation histograms; however, this can be changed to snapshot histograms in the SiliMon directives file.

2.2.2.1 Per layer histograms

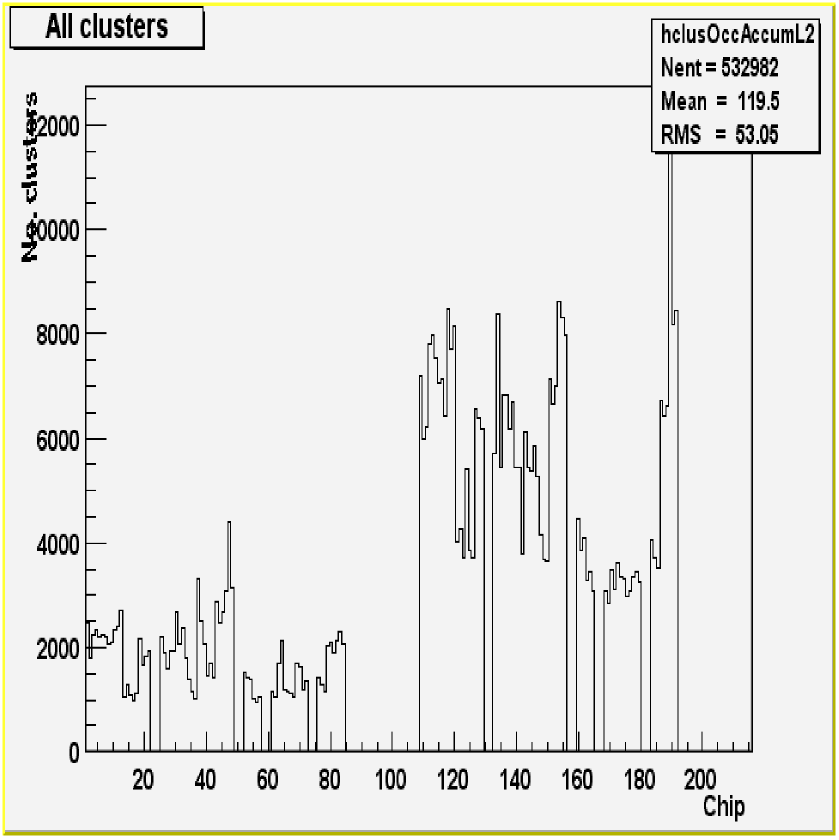

Figure 14: Cluster occupancy for layer 2 of the silicon. The x-axis represents chip number (as explained in the text), and the y-axis represents the number of clusters found in a given bin.

There are eight histograms of cluster occupancy per layer, each corresponding to one layer of the silicon. The x-axis is organised in order of increasing chip number. The bins are arranged by side (so the left half of the plot corresponds to axial chips), within each side by barrel, then phiwedge, then chip number. In order that there not be too many bins, each bin contains two chips. The y-axis represents the number of clusters found on a given chip in this layer. An example of this type of histogram is shown in figure 14. In this case occupancy for all clusters is shown, although the histogram exists for clusters associated to tracks as well.

2.2.2.2 Per layer barrel and side histograms