A High-performance Track Fitter for Use in Ultra-fast Electronics

E. Clement[a], M. De Mattia[b], S. Dutta[c], R. Eusebi[b,*], K. Hahn[d], Z. Hu[e], S. Jindariani[e], J. Konigsberg[f], T. Liu[e], J. Low[f], R. Patel[g], D. Rathjens[b], L. Ristori[e], L. Skinnari[h], M. Trovato[d], K. A. Ulmer[g], S. Viret[i]

[a]University of Bristol, Bristol, United Kingdom
[b]Texas A&M University, College Station, TX 77843, U.S.A.
[c]Saha Institute of Nuclear Physics, HBNI, Kolkata, India
[d]Northwestern University, Evanston, IL 60208, U.S.A.
[e]Fermi National Accelerator Lab, Batavia, IL 60510, U.S.A.
[f]University of Florida, Gainesville, FL 32611, U.S.A.
[g]University of Colorado, Boulder, CO 80309, U.S.A.
[h]Cornell University, Ithaca, USA
[i]Université de Lyon, Université Claude Bernard Lyon 1, CNRS-IN2P3, Institut de Physique Nucléaire de Lyon, Villeurbanne, France

Abstract

This article describes a new charged-particle track fitting algorithm designed for use in high-speed electronics applications such as hardware-based triggers in high-energy physics experiments. Following a novel technique designed for fast electronics, the positions of the hits on the detector are transformed before being passed to a linearized track parameter fit. This transformation results in fitted track parameters with a very linear dependence on the hit positions. The approach is demonstrated in a representative detector geometry based on the CMS detector at the Large Hadron Collider. The fit is implemented in FPGA[☆] chips and optimized for track fitting throughput and obtains excellent track parameter performance. Such an algorithm is potentially useful in any high-speed track-fitting application.

Keywords: Linearized Fitter, Tracker, Algorithm, Principal Component Analysis

1. Introduction

Determining the parameters that describe the trajectory of charged particles traversing a magnetic field is a common problem in high-energy physics applications. There are several solutions that obtain accurate track parameter resolution with methods such as least squares fits utilized in Kalman filters, which employ iterative algorithms to converge toward a desired precision [1, 2, 3]. In the context of offline processing, where finding accurate parameters

---

takes priority over the computing time and resources spent on it, these powerful algorithms are routinely used in experiments such as CMS and ATLAS.

In the online environment of these high-energy physics experiments, where high-speed electronics are used to make ultra-fast decisions about which particle collision events to store and which to reject, such sophisticated techniques can easily exceed the limits of available processing time. For example, the CMS Phase-2 upgraded level-1 trigger [4, 5, 6] aims to find and fit hundreds of tracks within about five microseconds after each proton-proton bunch crossing, which occur every 25 ns. Under such stringent conditions, multi-step iterative track fitting algorithms can quickly surpass the allowed time.

An alternative approach that trades accuracy for reduced computation time is to use a linearized fit. For a track parametrized with p parameters, this type of fit assumes that the true track parameters, $p_i^{\text{true}}, i = 1, ..., p$, can be estimated from a linear transformation of the track hit positions, x_j, where the index j runs over the components of the position of each of the N hits that are part of the fit.

This is expressed as

$$p_i = \Sigma_{j=1} A_{ij} \delta x_j + \overline{p}_i, \tag{1}$$

where A_{ji} is a pre-determined set of constants, δx_j is defined as $\delta x_j = x_j - \overline{x}_j$ and $\overline{p}_i$ and $\overline{x}_j$ are the track parameter and hit positions about which the linear expansion is performed.

The pre-determined elements of the matrix A are obtained from a sample of tracks representing the true correspondence between track hits and true track parameters. These matrix elements are obtained by minimizing the distance $\min(\langle (p_i^{\text{true}} - p_i)^2 \rangle)$ between the true track parameters p_i^{true} and the estimates of Eq. 1 for all track parameters.

In addition, following the principal component analysis described in [7] and [8] it is possible to determine a fit quality parameter that in the limit of validity of the linear approximation is distributed as a χ^2. This parameter is constructed as the quadrature sum of $N - p$ components, $\chi^2 = \sum_{i=1}^{N-p} \chi_i^2$, where each of the components can be expressed as a linear combination of the hit coordinates

$$\chi_i = \Sigma_{j=1}^N B_{ij} \delta x_j, \tag{2}$$

where the index i runs up from 1 to $N - p$, and the matrix B_{ij} is a predetermined set of constants obtained following [8] (page 111-112).

The use of such linear transformations is an especially appealing approach in fast electronics applications since it is based on a small number of addition and multiplication operations that can be computed very quickly and is easily encoded in modern FPGAs. A linearized track fit has been previously utilized in the CDF experiment [9] where a hardware track trigger was developed [7, 10, 11, 12]. The approach is also being utilized in the ATLAS experiment [8].

However, the assumption of linearity implies that the matrix A is independent of the track parameters, and any such dependence, for example on the track momentum, will introduce nonlinearities. In general, the problem of track fitting in most detector geometries involves nonlinearities due to the limited size, the shape, and the orientation of detector sensors, and thus the performance of any linearized fit degrades with the degree of these nonlinearities.

In this paper, we describe a new technique that corrects for these nonlinearities and thus improves the performance and overall resource utilization of linearized track fits in hardware.

The paper is organized as follows. Section 2 describes the origin of nonlinearities and the limitations of a linearized track fit using the CMS detector as an example. Section 3 describes a new approach to transform track hit positions, while Section 4 shows the performance of this new algorithm. Section 5 counts the memory requirements of the technique and Section 6 details the implementation and resource utilization of the algorithm in hardware. Section 7 presents the conclusions.

2. Origin of Nonlinearities

The trajectory of a charged particle in a uniform magnetic field can be described by a helix (when neglecting energy loss and other stochastic effects) and is described by five parameters. These are taken to be the charge over the transverse momentum of the particle (q/p_T), the azimuthal angle of the momentum vector at the point of closest approach to the origin (ϕ_0), the transverse impact parameter, defined as the minimum distance between the trajectory and the origin in the transverse plane (d_0), the cotangent of the angle between the momentum vector and the z axis at the point where d_0 is evaluated ($\cot\theta$), and the z coordinate of the trajectory at the point where d_0 is evaluated (z_0).

In this note we focus on the use of a linearized fitter with four parameters, where the d_0 parameter is assumed to be zero. This assumption is made for simplicity and does not affect the results or the conclusions of this paper. In the following sub-sections we discuss the sources that give rise to a nonlinear relationship between hit positions and track parameters.

2.1. Nonlinearities from Geometry

We will use the layout of the CMS Phase-2 upgrade detector described in the Technical Proposal [13] to illustrate the significant effects played by nonlinearities arising from irregular detector geometry. It should be noted that while the proposed CMS tracker geometry has changed since the writing of this article, the source of nonlinearities in the new geometry are identical to the ones discussed here. In the context of a full level-1 trigger, the track fit technique described in this paper would follow a pattern recognition step, which is not discussed here.

The CMS Phase-2 upgrade detector exemplified in this article is a multipurpose particle detector design with a "barrel" consisting of six cylindrical and concentric layers of silicon surrounding the beampipe and two "endcaps" consisting of five parallel disks of silicon on the forward and backward ends of the detector, as shown in Fig. 1. The layers and disks are composed of silicon modules that overlap slightly with neighboring modules to obtain full coverage for charged particles coming from the central luminous region of the detector.

The modules themselves are composed of a pair of silicon sensors separated by about $1-4$ mm to be able to have a local coarse measurement of the track's p_T for a given combination of hits in both sensors. This information is read out by the front-end electronics together with the position of the hit in the inner of the double sensors. In this note we do

not use or reference that coarse p_T estimate, and utilize only the position (R, ϕ, z) of the hit in the inner of the double sensors.

The inner modules (drawn in blue in Fig. 1) consist of a macro-pixel sensor and a strip sensor on top of each other. The macropixels and strips are $1.5\,\mathrm{mm}$ and $2.4\,\mathrm{cm}$ long in the z direction, respectively. The pitch is $100\,\mu\mathrm{m}$ for both sensor types. The outer modules (drawn in red) have parallel silicon strips of $90\,\mu\mathrm{m}$ pitch measuring $50\,\mathrm{mm}$ in length along the z direction.

We use the official right-handed coordinate system of CMS in which the origin lies at the geometrical center of the detector, the z axis lies longitudinally along the detector and along the direction of the magnetic field, and the x and y axes form the transverse plane. Given the cylindrical symmetry of the detector and its internal magnetic field we also use cylindrical coordinates (R, ϕ, z) in this note, where R and ϕ are constructed from the x and y coordinates of the official CMS coordinate system.

For a perfectly cylindrical detector, the radial coordinates do not provide information useful for a principal component analysis of the hit correlations since they are at fixed values for each layer. In this perfect detector it is sufficient to consider the hit position information (ϕ, z) from different layers as separate inputs. This approximation is almost always used because it reduces the number of constants in the matrix and the resulting number of operations needed to obtain the result, making the matrix multiplication faster and less memory intensive.

In a real detector, however, the hit coordinates for a given layer are not at the same radius because of two effects: the modules are flat, and they are staggered to allow for overlaps and complete coverage with active material. This is shown in Fig. 2 (left) where the barrel layers are depicted and pairs of modules at the same angle ϕ correspond to different R positions and have significant variation in radius.

The staggering of the modules results in nonlinearities in the transformation from tracker hits to track parameters. As an example, Fig. 2 (center) illustrates the size of these nonlinearities in the outermost layer of the detector for a $p_T = 2\,\mathrm{GeV}$ track. The assumption that the modules in that layer are at a constant radius introduces a p_T-dependent misestimation in the ϕ coordinate of the hit of about four mrad in either direction.

Likewise, an assumption of a constant radius introduces a change in the estimated p_T track parameter. To estimate the change in p_T that would compensate for the shift in radius, we consider a track assumed to come from the center of the detector in the transverse plane, that is with $d_0 = 0$. In that case we can write the coordinates ϕ and z of a point on the trajectory as a function of its radius R and the track parameters q/p_T, ϕ_0, $cot(\theta)$, and z_0 as

$$\phi = \phi_0 - \arcsin\left(\frac{R}{2\rho}\right) , \text{ and} \tag{3}$$

$$z = z_0 + 2\rho \arcsin\left(\frac{R}{2\rho}\right) \cot\theta , \tag{4}$$

where $\rho = 0.33\,\mathrm{m} \cdot (p_T/\mathrm{GeV})/(B/\mathrm{T})$ is the radius of curvature of the trajectory in the
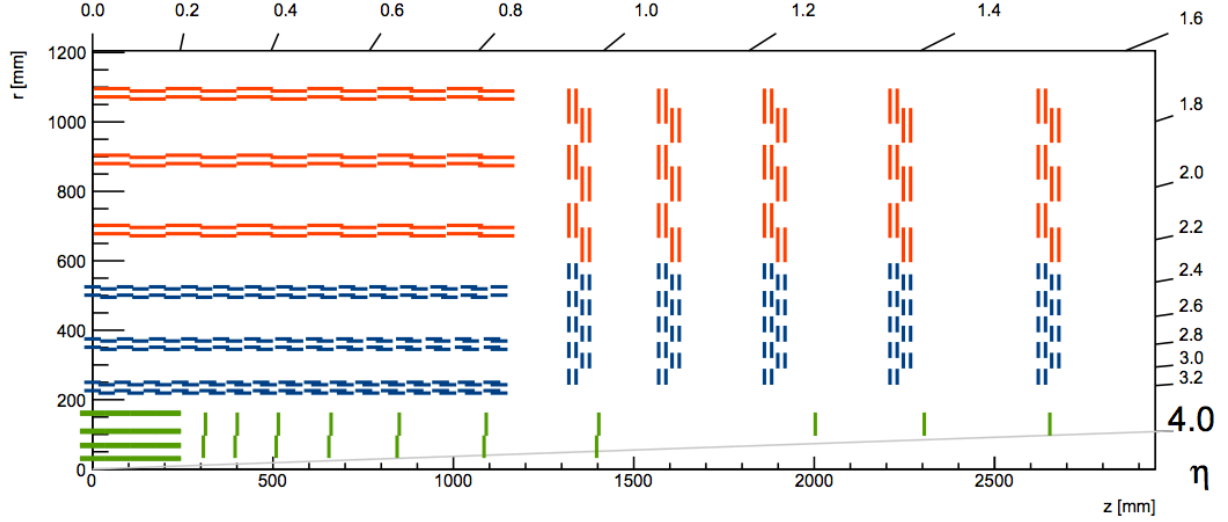
Figure 1: Sketch of one quarter of the CMS Phase-2 upgrade silicon tracker as described in the Technical Proposal [13]. The red and blue lines represent silicon strip modules that are used in the track trigger, while the green lines represent silicon pixel modules, which are not considered in the track fit proposed here.
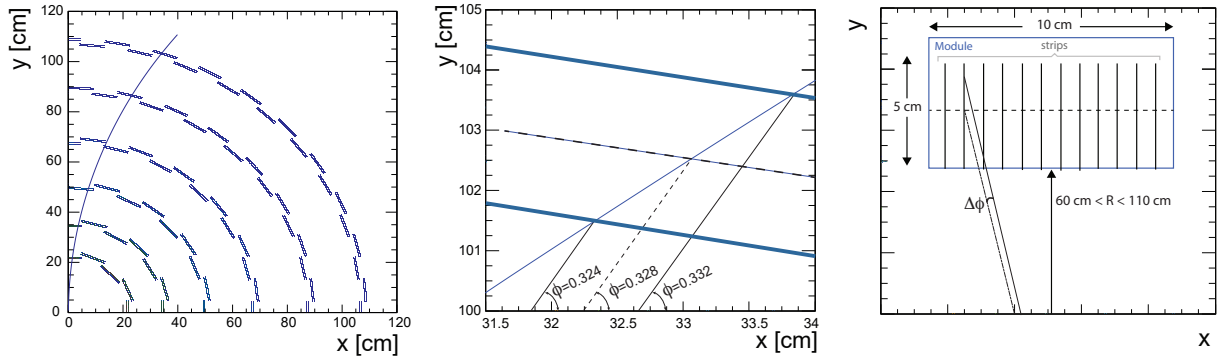

Figure 2: Geometrical origin of nonlinearities. **Left:** staggering and flatness of the modules in the barrel as indicated by the positions of reconstructed hits from simulated single muon events. The line depicts a 2 GeV track. **Center:** zoomed view of the left figure in the area of the tracks crossing the outermost layer. The thick lines represent the average radius of the innermost and outermost modules in the layer, while the dotted line shows the average radius of all modules in the layer. The angles are the ϕ coordinates of the hits that would be registered for this track depending where the hit is actually produced. **Right:** Parallel-strip modules in the disk. Sketch of a 2D module in the disks of the endcap region and some of the strips within. The figure shows the difference in angle of the ϕ coordinate due to the parallel arrangement of the strips.

transverse plane and is directly proportional to the p_T of the track and inversely proportional to the magnetic field B.

From Eq. 3 we can compute the p_T of a track that would produce, on the outermost staggered module, a hit with the same ϕ of the hit produced by the 2 GeV track in the

5

innermost module. This relative variation in p_T is approximately independent of p_T and is approximately inversely proportional to the variation of the radius of the hit. The biggest variation in p_T is seen when considering the innermost layer where it can be as big as about 11%, while in the outermost layer it is the smallest and it is about 2.4%. These effects significantly degrade the online performance and need to be addressed to reach the typical offline track p_T resolution of around 0.5% (2%) for 2 (100) GeV tracks [14].

A second effect that causes biases and reduces the precision of the measured parameters is the planar shape of the modules, which causes a variation of the radius of the hit within a single module. This effect is most important for the innermost layer, where 16 modules cover the full ϕ angle with each module covering approximately 0.4 radians. Assuming the center of the module to be its closest point to the origin and having a radius of approximately $23\,\mathrm{cm}$, the radius at the edge is approximately $23.5\,\mathrm{cm}$, a variation of about 2%. For a $2\,\mathrm{GeV}$ track this translates into a similar effect on the p_T of about 2%. For the second innermost layer this effect is reduced to less than 1% by the increase in the number of modules (24) and it is reduced further as the number of modules per layer increases at even larger radii.

Furthermore, Eq. 4 shows that the effect of the variation of the radius on $\cot\theta$ is similar to the effect on p_T in the case of the ϕ coordinates. The effect is maximal in the outermost layer where it is about 11% and minimum in the innermost layer where it is about 2.5%. Again, this is a significant effect in the performance of a linearized fit.

A third source of nonlinearities from detector geometry stems from the parallel orientation of the strips on the silicon module. This is problematic in the endcap disks, where the outer modules are arranged so that their $\sim 5\,\mathrm{cm}$ strips roughly point towards the z axis. However, since the strips are parallel to each other, if one strip points straight to the z axis, all others will not. Figure 2 (right) shows the geometry of an outer module in the disk as seen in the transverse plane. Since the outer modules of the disk do not provide information on the position of the charge deposit along the strip, each hit on each strip in the disks is assigned the R coordinate of the center of the strip drawn as a dashed line in the figure. The figure shows how the parallel strips arrangement causes the correct radius R' of the hit to deviate from the assigned radius R if the particle did not pass through the center of the strip. As a consequence, the ϕ coordinate of the hit is also biased, as shown in the figure. This source of nonlinearity does not exist in the barrel, where the strips are directed along the z axis. It would also disappear from the disks if either the strips were arranged in a radial fashion all pointing toward the z axis or if a more accurate position of the hit along the strip was provided, as happens in the inner modules.

2.2. Nonlinearities from Functional Form

Another source of nonlinearities comes from the track equations themselves. Equation 3 shows that the equations that describe the trajectory of tracks are nonlinear and that simply using linearized equations will not give the optimal performance. It should be noted that other effects, such as a realistic magnetic field that is not fully homogeneous over the tracking volume, will increase the complexity of the functional form of Eq. 3.

While this source of nonlinearity seems irreducible in the usage of a linearized track algorithm, we will show in the following sections how it can be circumvented.

6

2.3. Standard Mitigating Mechanisms

To mitigate these effects the most common approach taken by the high-energy physics field is to generate many different sets of constants, each to be applied in a smaller region of the detector. By reducing the region of application in the detector the effects of the nonlinearities in that region are effectively reduced. The cost of such an approach is the need to generate, store, and then quickly retrieve a large number of sets of constants and apply each to the correct region of the detector. The constants must also be stored in the hardware and be accessible to the fit, therefore limiting the amount of memory available for the fitter engine itself. A tradeoff must be made between the resolution of the fitted parameters, which depends on the size of the region of application, and the total number of constants needed. The typical number of regions used in high-energy physics experiments ranges from a few hundred thousand to several million.

Furthermore, it should be noted that the nonlinearities, which are the root-cause of the lack of performance of any linearized track fitter, are simply not addressed. In this approach, the underlying nonlinearities remain in place reducing the performance even within the small regions; it is just the interval of application of each set of constants that is reduced.

3. Hit-position Transformations

Instead of reducing the region of application of each set of constants while increasing the total number of regions, we describe here an approach that directly addresses the nonlinearities of the system, allowing for a considerably more performant track fitter for a given region of application, or alternatively allowing the region of application to be extended significantly for a given desired performance.

Our approach starts by restoring the cylindrical symmetry at the input stage of the linearized fitter by transforming the locations of the hits on a physical layer to new positions that lie on an ideal cylindrical layer of fixed radius, R'. In addition, the transformation is made such that the coordinates of the new hit position (R', ϕ', z') have a linear functional dependence on the track parameters:

$$\phi' = \phi_0 - \frac{R'}{2\rho}, \tag{5}$$

$$z' = z_0 + R' \cot \theta \,. \tag{6}$$

These linear parameterizations are chosen because they are the first order approximations of Eqs. 3 and 4. Using Eqs. 3 and 4, and Eqs. 5 and 6, we write the hit coordinates on the ideal layer (R', ϕ', z') in relationship to the coordinates on the physical layer (R, ϕ, z) and the track parameters as:

$$\phi' = \phi - \frac{R'}{2\rho} + \arcsin\left(\frac{R}{2\rho}\right), \tag{7}$$

$$z' = z + 2\rho \cot(\theta) \left[\frac{R'}{2\rho} - \arcsin\left(\frac{R}{2\rho}\right)\right]. \tag{8}$$

7

Let's consider the first order approximation of Eq. 7 around $R/(2\rho) = 0$:

$$\phi' = \phi + \frac{R - R'}{2\rho}. \tag{9}$$

The size of the shift in the ϕ hit coordinate required to extrapolate to the ideal layer is $s = \phi' - \phi = (R - R')/2\rho$. However, we note that the relative accuracy of the shift, $\delta s/s$, is directly proportional to the relative accuracy of the radius of curvature, which by construction is also that of the transverse momentum:

$$\frac{\delta s}{s} = \frac{\delta \rho}{\rho} = \frac{\delta p_{\mathrm{T}}}{p_{\mathrm{T}}}. \tag{10}$$

This shows that an a-priori value, or "pre-estimate" of the transverse momentum with a large relative resolution can be used to evaluate the shift in the ϕ coordinates to that same relative precision. For example, a p_{T} pre-estimate with a 3% relative resolution can produce a first-order correction that is only 3% off from ideal. Our technique develops directly from this fact.

We start by performing a fast and coarse pre-estimate of the track quantities $q/(2\rho)$ and $\cot \theta$. These pre-estimates are also performed as a linear fit with a single set of constants determined from averaging over the entire region of application. As described in Section 2.1 these pre-estimates have sub-optimal resolution on the track parameters providing a transverse momentum relative resolution of ~ 2–3%. While this coarse measurement is worse than the ultimately desired precision on the track parameters, it will allow the transformation of the projected ϕ position to the same relative precision.

We then transform the hit positions to the ideal layer. To minimize processing time in real hardware applications instead of directly using Eqs. 7 and 8 we utilize the following polynomial approximations:

$$\phi' = \phi + \frac{R - R'}{2\rho} + \frac{1}{6}\left(\frac{R}{2\rho}\right)^3, \tag{11}$$

$$z' = z - \cot(\theta)(R - R') - \frac{\cot(\theta)R^3}{6(2\rho)^2}. \tag{12}$$

The result of these transformations is illustrated in Fig. 3 where the top panel shows the original hit locations and the bottom panel shows those hits transformed onto common cylindrical surfaces. The (positive z) half of the CMS detector has been divided into 14 regions. The plot shows the nine major geometrical regions depending on the combination of modules in the barrel layers and disks, which are shown in different colors in the figure. Five of these regions are further separated in two but are left out of this figure for visualization purposes. Notice how the whole barrel layer six constitutes a single region and how the hits in the disks were also extrapolated to ideal cylindrical layers in each individual region.

We can also use the pre-estimated quantities to reduce the nonlinearities coming from having parallel-strip arrangements in the outer modules of the disks as discussed at the end
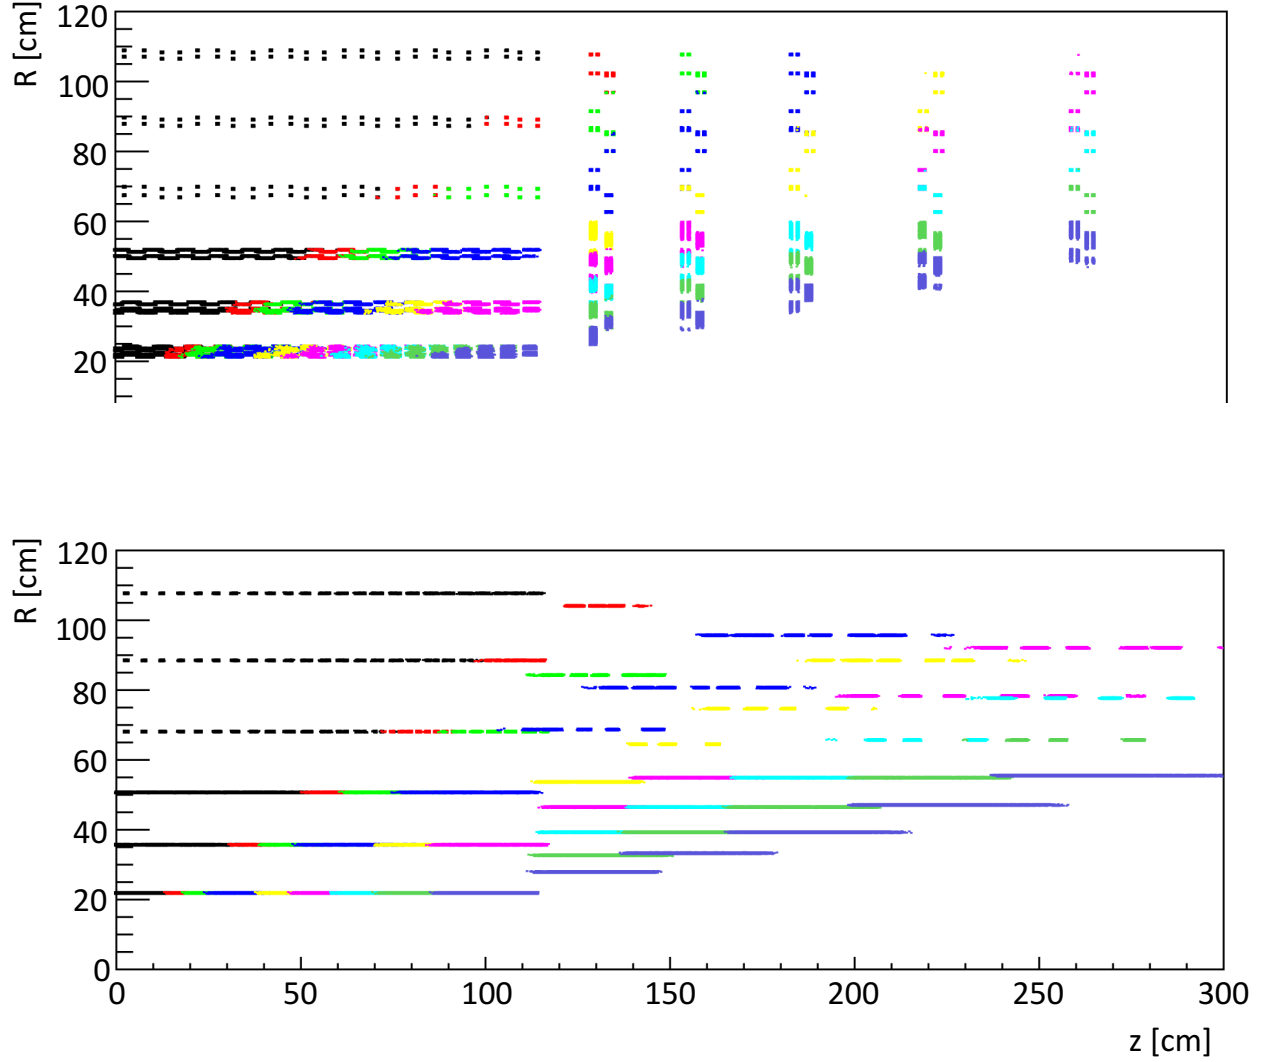
Figure 3: Hit positions in the (r, z) plane. The colors represent hits assigned to the different regions. Top: original positions. Bottom: transformed hit coordinates.

of Section 2.1. For those modules located at a nominal radius R we estimate the correct radial value of the hit coordinate, R_{ex}, as a linear extrapolation from the hit z coordinate in the outermost barrel layer or the outermost inner disk as:

$$R_{ex} = (z - z_0) \tan \theta. \qquad (13)$$

We then correct further the ϕ coordinate of hits in these modules as:

$$\Delta\phi = p \cdot (h_{sn} - m_{sn}) \cdot \frac{R_{ex} - R}{R^2}, \qquad (14)$$

9

where h_{sn} is the strip number of the actual hit, m_{sn} is the strip number of the center strip in the module, and p is the distance separating any two strips in the module. Note that the computation of R_{ex} requires a pre-estimate value of $\tan\theta$, rather than $\cot\theta$. To avoid performing a costly division in the hardware, we simply compute $\tan\theta$ as another pre-estimate quantity together with $q/(2\rho)$ and $\cot\theta$.

It is important to notice that this technique transforms the individual position of a hit to a cylindrical geometry using only a track pre-estimate and the individual hit position. The technique does not use any information from other hits and is therefore blind to the arrangement of hits that stems from a given detector geometry. That is, since it will perform an intermediate transformation that changes that geometry into an ideally cylindrical barrel, this technique can be deployed in any detector geometry.

A second linearized track fit is then performed on the transformed hit coordinates, in each region, which is used to obtain the final track parameters. This technique directly addresses the geometrical nonlinearities described in Section 2.1. In addition, Eqs. 7 and 8 directly target the nonlinearities arising from the functional forms described in Section 2.2. This technique mitigates both sources of nonlinearities. The net result is a least-squares fit to the track hits where the transformed coordinates have been used to significantly reduce the nonlinearities based on the pre-estimate fit.

4. Performance of a Track Fitter Based on Transformed Hits

The approach of projecting hit positions onto idealized layers before a linearized track fit has been explored using the CMS Phase-2 upgrade tracker geometry. Details of the implementation of this proof of principle and the track fitting results are presented in this section.

A sample of simulated single muons, generated evenly in $1/p_{\mathrm{T}}$ between $0.5/\mathrm{GeV}$ and $0.005/\mathrm{GeV}$, is used to determine the optimal coefficients needed in each region to implement this technique. In general, however, any set of good tracks can be used, and a sample of tracks reconstructed directly in data would allow for detector misalignment effects to be taken into account directly. The sample of muons was simulated using the official CMS Monte Carlo simulation described in [14] using the official CMS software package CSSMW [15].

In the case of the specific example of the CMS detector discussed here the tracker is divided into a total of 14 different regions which correspond to different combinations of barrel and endcap detector modules. Tracks that pass entirely through the barrel are considered in one region, tracks that pass through five barrel layers and one endcap disk are considered in another region, etc. Additionally, regions containing at least one endcap disk are further split in two to consider the different characteristics of modules in the inner and outer regions of the detector. For each region, an individual set of constants is determined. Within each region, the entire detector utilizes the same set of constants for all ϕ coordinates. To further suppress nonlinearities at low p_{T}, different sets of constants are trained for tracks above and below $p_{\mathrm{T}} = 10\,\mathrm{GeV}$ and the correct set of constants is determined on the fly based on the p_{T} pre-estimate.

In the CMS geometry layers have overlapping modules allowing a single track to produce hits in several modules of a single layer. Since our technique uses a single hit per layer (as discussed in Section 2.1) we mask off fixed regions in the outer modules where hits are likely to be coming from tracks that also leave hits on an inner, partly overlapping, module. This results in having almost exclusively one hit per layer at a negligible efficiency loss.

Tracks in this geometry produce hits in up to seven different layers or disks, with the typical case producing hits in six layers. In this study we consider tracks with one hit in five or six different layers. Tracks with hits in seven different layers are reduced to six by removing the hit in the outermost layer.

An independent sample of muons generated evenly across η, ϕ, and $1/p_T$ was used to determine the performance of the fit. Figure 4 shows the fitted resolutions of the track p_T, ϕ_0, z_0, and η as a function of true track η and for three ranges of true track p_T. In all cases, the obtained resolutions compare very well with the performance of offline simulations and are thus excellent for a track trigger application. The worsening of resolutions as a function of η also follows closely the offline simulation [14].

The relative p_T resolution is about 0.5% for a $2 < p_T < 5\,\text{GeV}$ track in the region of $|\eta| < 0.2$. It may be interesting to notice that if Eq. 11 is truncated to the first order the relative p_T resolution will jump to about 1.6%. The improvement of the relative resolution with a second order correction is the net effect of addressing the nonlinearities stemming from the functional form of the track equations.

No significant bias is observed in the fits with the mean central value within the resolution of the true central value for all parameters. The distribution of the reduced χ^2 values is shown in Fig. 5 and is observed to be nicely peaked at values near one, and thus appropriate for a selection requirement of high quality track fits.

5. Memory Requirements

The number of constants needed to estimate the parameters of a track in a given region of the detector is an important consideration in the hardware deployment of linearized track fitters.

As described in Sect. 3, our technique begins by computing an initial pre-estimate to obtain coarse track p_T, $\cot\theta$, and $\tan\theta$ parameters using Eq. 1. The p_T is computed using the ϕ coordinates of the hits in the six layers and requires only thirteen coefficients (six A_j, six $\overline{\phi}_j$, and one $\overline{p}_T$). Following Eq. 1 the over-barred quantities $\overline{\phi}_j$ and $\overline{p}_T$ represent the average ϕ hit positions and track p_T parameters about which the linear expansion is performed.

To achieve sufficient precision for the $\cot\theta$ and $\tan\theta$ pre-estimates we utilize both the z and R hit-coordinates using a total of 25 coefficients per pre-estimate (twelve A_j, six $\overline{z}_j$, six $\overline{R}_j$, and $\overline{\cot\theta}$). The total number of constants to estimate the pre-estimates for the 14 regions of the detector is then 882.

This is then followed by the computation of the ϕ and z positions of the hits at their projected locations on the idealized layers. This step requires no pre-stored constants except
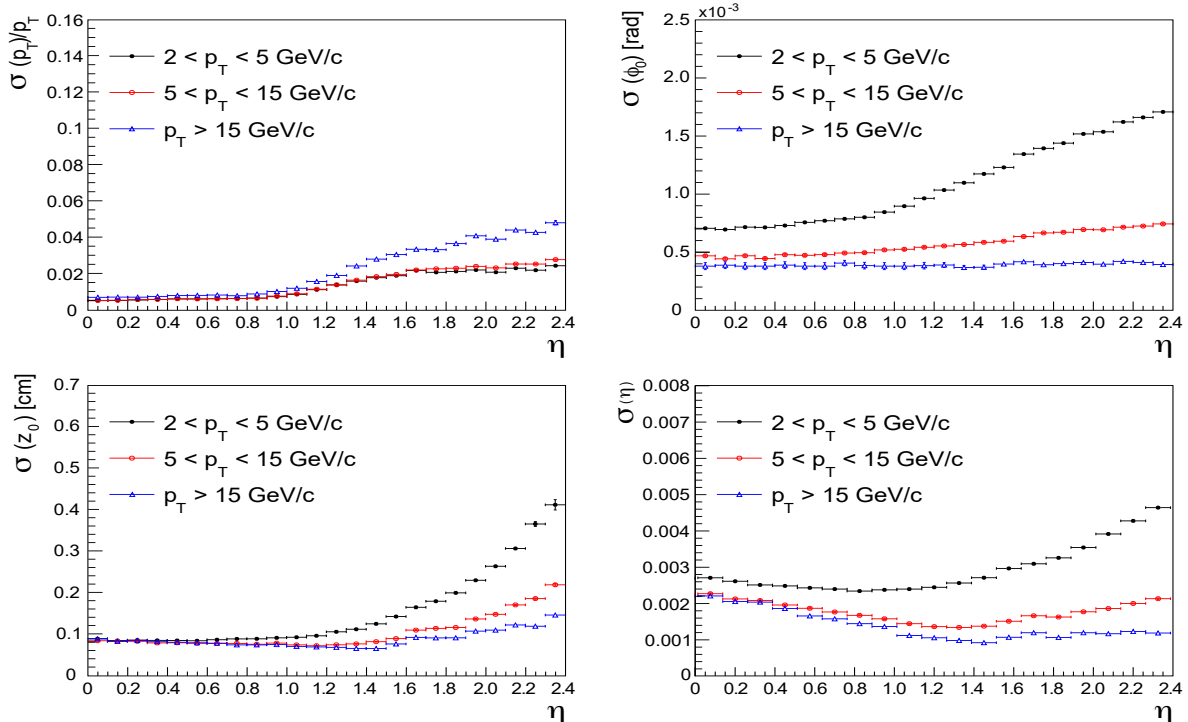
Figure 4: Track parameter resolutions for single muon events.

for the radius of the idealized layers. This requires at most six ideal layers per region for a total of 84 constants.

Because the pre-estimate returns values of q/p_{T} and $\cot(\theta)$ the transformations of the hit position in Eqs. 11 and 12 require only additions and multiplications and are thus particularly suited for fast applications in FPGAs. The only exception is the $1/R^2$ term in Eq. 14 that is computed via a lookup table of 16 constants for the whole detector.

The final fit is then performed on the transformed hit positions. Since all remaining nonlinearities are negligible and the R coordinates are fixed, all track parameters and χ^2 components are obtained with either the six ϕ or six z transformed coordinates. In the transverse plane the computation of p_{T}, ϕ_0, and the transverse χ^2 coefficients could be performed with a single set of a 6×6 matrix of coefficients, six $\overline{\phi}_{\mathrm{j}}$ and a value for $\overline{p_{\mathrm{T}}}$ and $\overline{\phi}_0$ for a total of 44 constants. Central values for the χ^2 terms are zero by construction. However, to improve performance at low transverse momentum, we use two sets of constants, for p_{T} below and above $10\,\mathrm{GeV}$, and dynamically choose one based on the $1/p_{\mathrm{T}}$ pre-estimate. This step then requires 88 constants per region for a total of 1232 constants.

In the longitudinal plane the computations of the z_0, $\cot\theta$, and transverse χ^2 coefficients use another 6×6 matrix plus six $\overline{z}_{\mathrm{j}}$, $\overline{z_0}$ and $\overline{\cot\theta}$ for a total of 44 constants per region and a total of 616 constants for all regions.

Notice that because the transformation restores cylindrical symmetry to the detector we can use a single set of coefficients to estimate the final track parameters and χ^2 components
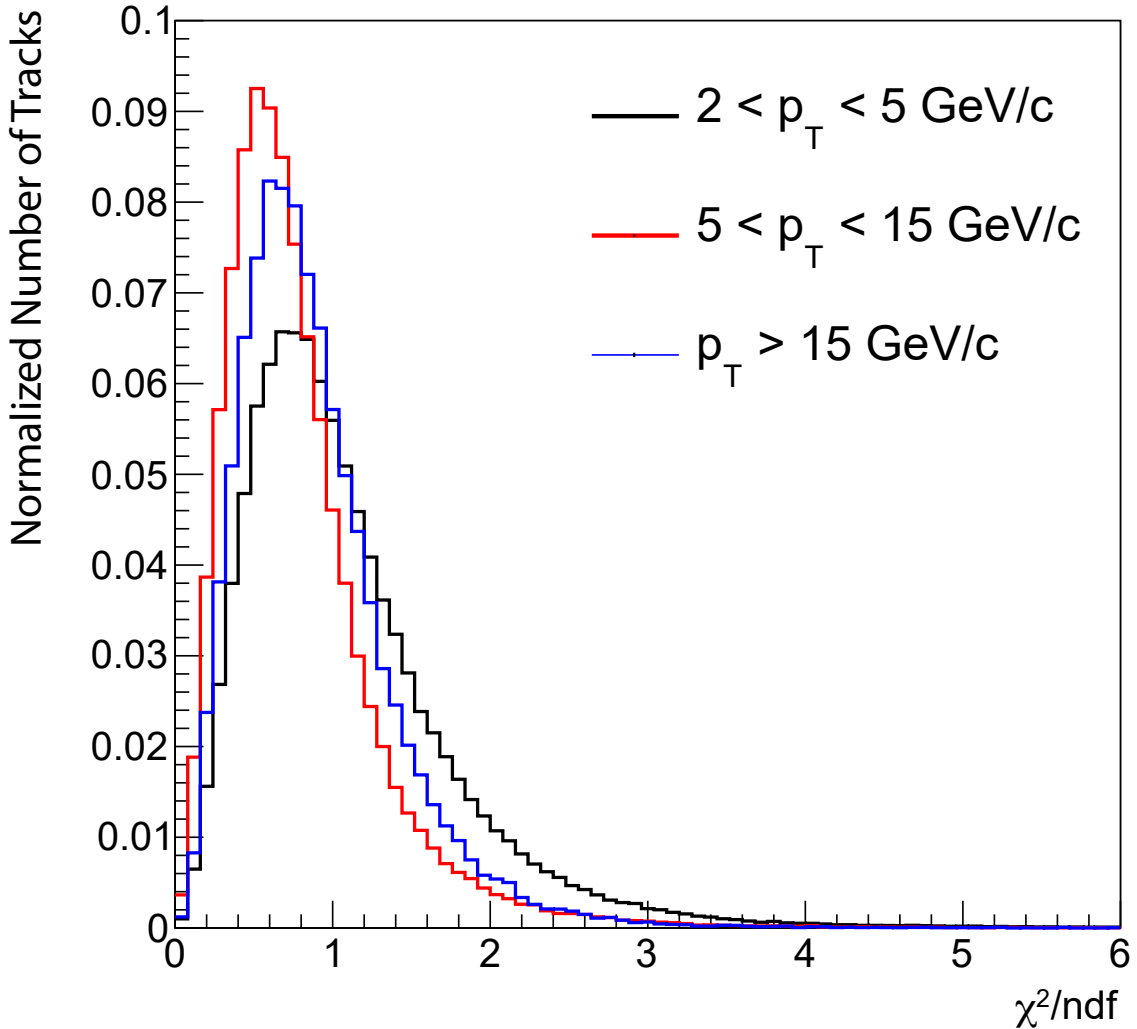
12

Figure 5: Fitted values for track χ^2 for single muon events.

for the full 2π azimuthal range in each region. This is a distinct advantage over other approaches that consider each possible combination of modules as a separate entity with a dedicated set of coefficients.

The total number of constants used for the full detector is 2830. It is expected that some tracks will leave only five hits instead of six. Depending on the particular hit that is missing it is possible to have six other sets of combinations, which brings the grand total number of constants to 19,810.

In comparison, the number of constants used in the track-trigger system of the ATLAS

detector is close to a million [8], and the number used in the CDF experiment for their track trigger system was about eight million [10]. As described in the next section, the number of constants to be stored has implications on the firmware implementation of the linearized track fitter.

6. Hardware Implementation and Timing

The firmware implementation of this technique for the transverse components is shown schematically in Fig. 6. The algorithm starts by inputting the raw hit positions (R_j, ϕ_j, z_j) at the left, which are then used to compute the $q/(2\rho)$ and $\tan\theta$ pre-estimates. These in turn are used to transform the hit positions which are then passed to the linearized track fit to determine the track parameters q/p_T and ϕ_0, and the goodness of fit components χ_i. A similar schematic is used for the longitudinal components. This firmware is implemented in the Kintex Ultrascale KU060 FPGA. A set of 10,000 tracks was passed through both the firmware and emulator obtaining 100% bit-to-bit matches.

Within the Kintex Ultrascale class FPGA, the digital signal processing (DSP) structure allows for multiplication of inputs and addition to other inputs, which is a structure that accommodates well the matrix multiplication needed for the linear track fits. The DSPs are arranged in columns and the scalar products are implemented as a set of chained MACC (multiply-accumulate) DSPs, taking advantage of dedicated DSP interconnections within a single column. Because the DSPs are chained they cannot all operate in parallel on the same coordinates and the initial latency is higher than more straightforward configurations. However, the system is fully pipelined and the configuration allows higher clock speeds thus maximizing the throughput and minimizing the overall latency over a large number of consecutive fits.

In addition to track hit positions, each operation also needs access to the relevant constants for the conversion from hit positions to track parameters as in Eq. 1. The number of needed constants described in Section 5 is small enough so that they can be stored in the fabric of the KU060 instead of utilizing the larger block RAM. Less than 5% of the total LUT, LUTRAM, and flip-flop resources of the KU060 are needed. This accomplishes three things: it allows the block RAM to be used for other more memory-intensive portions of track finding, it reduces routing delay by placing the constants close to the DSP operators, and it allows a higher clock speed. The design described here was run at a speed of 500 MHz and utilizes 166 DSPs, or 6% of the DSPs available on the KU060. The design is fully pipelined and is able to accept a new set of track hits each clock cycle. After the initial fixed latency of 39 clock cycles (78 ns), a new track fit is returned each clock cycle, thus allowing a total of 461 track fits per μs per single instance of the track fitter. Eight instances of the fitter were implemented on a single KU060 while still running at 500 MHz.

7. Conclusion

We describe here a novel approach to linearized track fitting based on transformations of the track hit positions to reduce inherent nonlinearities stemming from a combination
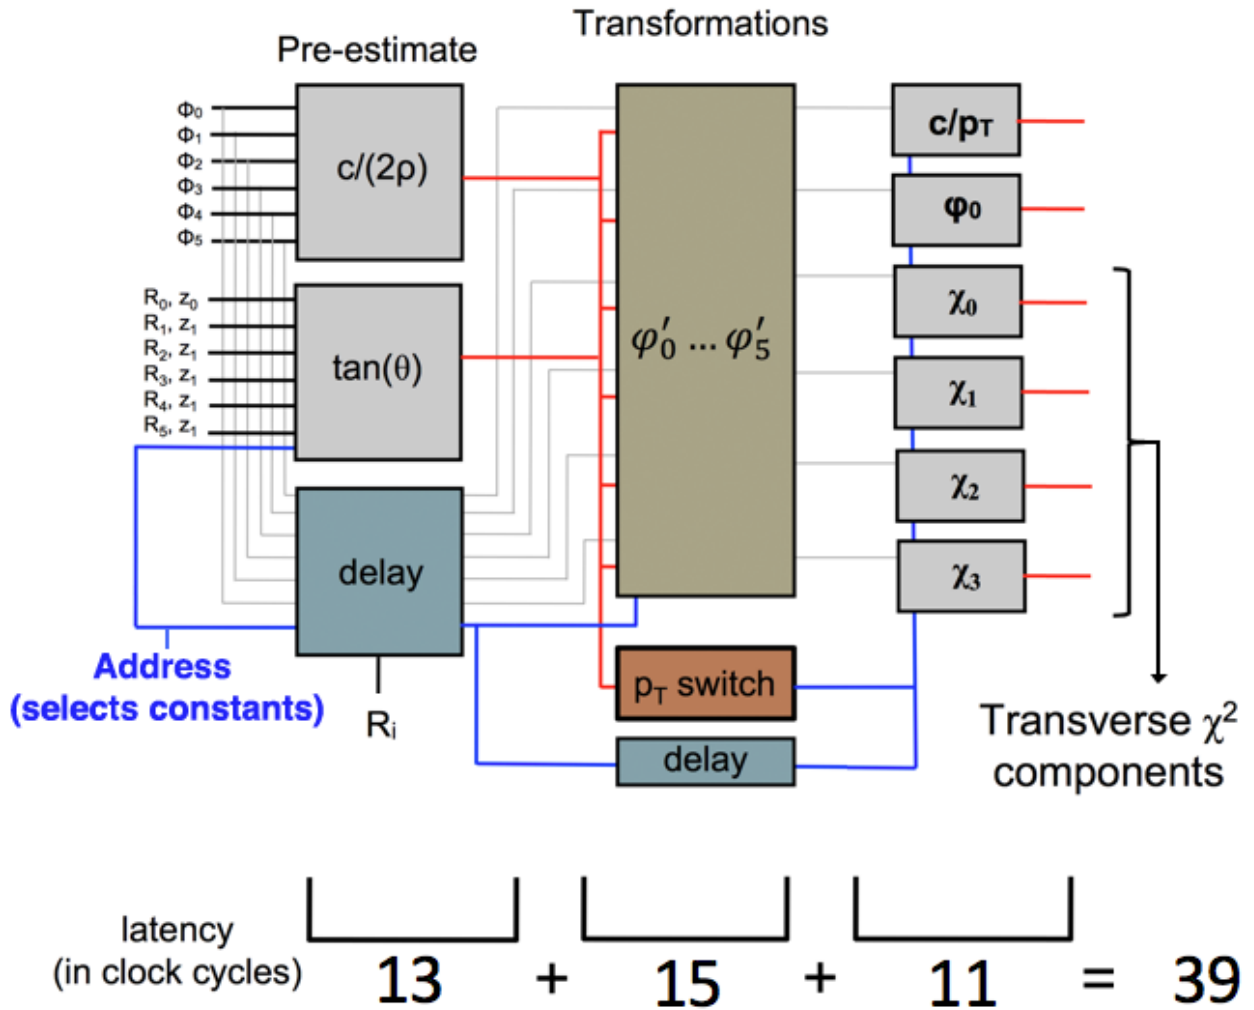
14

Figure 6: Block diagram showing the structure of the firmware implementation of this algorithm for finding the track parameters in the transverse plane. The latency of each step in clock cycles is shown on the bottom.

of detector geometry and the functional form of track trajectories. Unlike conventional techniques that reduce the region of application to contain the effects of the nonlinearities, our approach directly targets and reduces the nonlinearities of the system. This allows for much better track fit resolutions for a given region of application, or alternatively allows the region of application to be extended significantly for a given desired resolution, thus significantly reducing the required resources to implement the fit.

This approach is developed keeping in mind typical hardware resources and limitations and is demonstrated on a Kintex Ultrascale KU060 FPGA based on the geometry of the CMS upgrade tracking detector. It is shown to have excellent performance in terms of track fit resolutions, hardware latency, and resource utilization requirements. This approach is

therefore applicable in the context of very fast track fitting, such as found in hardware triggers for high-energy physics experiments.

8. Acknowledgments

References

[1] R. Frühwirth, "Application of Kalman filtering to track and vertex fitting", *Nucl. Instrum. Meth. A* **262** (1987) 444, `doi:10.1016/0168-9002(87)90887-4`.

[2] I. Heinze, "Development of a Hough Transformation Track Finder for Time Projection Chambers". PhD thesis, Hamburg University, 2013. `doi:10.3204/DESY-THESIS-2013-055`.

[3] L. Bugge and J. Myrheim, "Tracking and track fitting", *Nucl. Instrum. Meth.* **179** (1981), no. 2, 365, `doi:10.1016/0029-554X(81)90063-X`.

[4] CMS Collaboration, "The upgrade of the CMS trigger system", *JINST* **9** (2014), no. 08, C08002.

[5] J. Chaves, "Implementation of FPGA-based level-1 tracking at CMS for the HL-LHC", *JINST* **9** (2014) C10038.

[6] G. Hall on behalf of the CMS TMTT Team, "A time-multiplexed track-trigger for the CMS HL-LHC upgrade", *Nucl. Instr. Meth. A* **824** (2016) 292, `doi:10.1016/j.nima.2015.09.075`.

[7] L. Ristori and G. Punzi, "Triggering on heavy flavors at hadron colliders", *Annual Review of Nuclear and Particle Science* **60** (2010), no. 1, 595, `doi:10.1146/annurev.nucl.012809.104501`.

[8] ATLAS Collaboration, "Fast TracKer (FTK) Technical Design Report", Technical Report CERN-LHCC-2013-007. ATLAS-TDR-021, Jun, 2013. ATLAS Fast Tracker Technical Design Report.

[9] CDF Collaboration, "The Upgraded CDF Detector (Run II) (CDF)". oai:inspirehep.net:1110316.

[10] CDF Collaboration, "The CDF silicon vertex trigger", *Nucl. Instrum. Meth. A* **518** (2004) 532, `doi:10.1016/j.nima.2003.11.078`.

[11] CDF Collaboration, "The silicon vertex trigger upgrade at CDF", *Nucl. Instrum. Meth. A* **572** (2007), no. 1, 361, `doi:10.1016/j.nima.2006.10.383`.

[12] S. Amerio et al., "The GigaFitter: A next generation track fitter to enhance online tracking performances at CDF", in *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, p. 1143. 2009. `doi:10.1109/NSSMIC.2009.5402410`.

[13] CMS Collaboration, "Technical Proposal for the Phase-II Upgrade of the CMS Detector", CERN-LHCC-2015-010, LHCC-P-008, CMS-TDR-15-02 (2015).

[14] CMS Collaboration, "The Phase-2 Upgrade of the CMS Tracker", Technical Report CERN-LHCC-2017-009. CMS-TDR-014, CERN, Geneva, Jun, 2017.

[15] C. D. Jones, M. Paterno, J. Kowalkowski, L. Sexton-Kennedy, and W. Tanenbaum, "The New CMS Event Data Model and Framework", Proceedings of CHEP 2006, volume 1, editor S. Banerjee (India: Macmillan), 248.