CERN-SPSC-89-005

CERN/SPSC 89-5
SPSC/P244
13.01.1989

PROPOSAL TO THE SPSC

Research and Development Program on a Fast Trigger System for Heavy Flavour Tagging in Fixed Target and Collider Experiments

S.R.Amendolia[1,*], F.Bedeschi[2], G.Bellettini[3], S.Galeotti[2],
A.Lusiani[3], A.Menzione[2], F.Morsani[2], D.Passuello[2], L.Ristori[4].

1) Sezione INFN, Pisa and Universita' di Sassari, Italy
2) Sezione INFN, Pisa, Italy
3) Sezione INFN and Dipartimento di Fisica, Pisa, Italy
4) Sezione INFN and Universita' di Catania, Italy

*) Spokesman and Contactman

Introduction

We would like to place a request for the use of a CERN SPS beam during 1989-1990. The project is the development of a trigger technique, based on state-of-the-art technology, to be applied in fixed target and collider experiments for the study of heavy flavour production.
This document will briefly cover the physical and technological goals, the design study and the resulting requests of support from CERN.

Physics goals

Our work originated from the need of a very fast trigger in fixed target applications [1]: the trigger had to be extremely selective in tagging heavy flavour events, such as beauty production. At CERN SPS energies, in fixed target, the ratio of beauty events to QCD background is of the order of 10^{-6}-10^{-7} . The most promising way of disentangling the few beauty events from background is to make a full pattern recognition of the event structure, and then proceed to vertex counting, the number of secondary vertices being one of the discriminating quantities (see below).

Such a technique, which will be overviewed in this document and proposed for a test, is particularly well suited to high energy, fixed target set-ups. It could be used in any forthcoming spectrometer at CERN, or easily exported to one of the existing apparatuses at FNAL Tevatron. However, the existence of powerful P-Pbar collider experiments, and the much more copious production rate for beauty events, has naturally steered our attention towards finding an application in a collider set-up. Although our studies are not yet conclusive, we are convinced that some of the features of the trigger that we propose here can be to a good extent tailored to a collider detector environment. To name a few, the exploitation of the associative memories for pattern recognition and the extensive aplication of distributed parallel processors could make the task affordable.

Finally, the study of hardware techniques and software algorithms capable of making refined event selections at a rate well in excess of 10^5/sec makes the work particularly interesting in view of applications at the future supercolliders.

Technological breakthroughs

As it will become clear in the following, we plan to deploy state-of-the-art tools to realize the trigger, and a significant progress in detector technology would have to be achieved for us to reach our physics goals.

The detector (a silicon microstrip telescope) will have to cope with a very high incoming flux. The large amount of electrodes to be read and the speed requirement demand for an effective on-line zero suppression facility in the read-out electronics. The read-out chips, developed at Berkeley [2] and suitable for both fixed target and collider applications, provide an on-chip sparse data scanning which fulfills our needs.

One of the most difficult and time consuming tasks is the pattern recognition of the event, namely the association of the hits on the layers into tracks: this will be achieved via the use of associative memories [3], a custom VLSI chip which we developed and which allows the recognition of all the charged tracks in the telescope in a virtually null time. A high speed, high density version of this chip will be employed in a real physical environment for the first time in this experiment.

Finally, the complexity of the selection algorithm to signal secondary vertices and the required high degree of parallelism impose to develop of a powerful yet fast processor, which has to stand rates of more than 10^5events/sec with a limited pipelining of the data acquisition chain. Such a tool, that we plan to develop in collaboration with the Advanced Computing Group of CERN DD, would constitute a pioneering work easily exploitable in other similar applications.

Test experiment layout

The layout of the basic telescope and related electronics reflects the intent of optimising the device for fixed target and collider applications. However, the choice of some geometrical parameters was constrained by the availability of tools already developed for other applications, such as the silicon microstrip layers for the CDF vertex detector. By exploiting these tools we achieve a significant saving in costs without appreciably affecting the performance of the trigger.

Silicon detector layout

The telescope is made of 12 layers of silicon microstrips, with 5 planes oriented along y, 5 planes alonz z, 1 along u (+45°) and 1 along v (+135°). We use a coordinate system with x along the beam and z along the vertical. See Fig. 1.

The silicon detectors are those being used in the CDF vertex detector. The thickness is 280 μm. The sensitive area is 30.7 mm x 85 mm (only a fraction of this area will be used for the trigger). The microstrip pitch is 60 μm: it has been proved by simulation that this is sufficient for the vertex reconstruction at a test level. A pitch of 25 μm would bring little improvement in the resolution, since a larger number of strips is hit by inclined tracks.

The test of a prototype system performed at FNAL [4] has given a detector efficiency above 99.7%. In our digital read-out mode, this will turn into an efficiency of 99.1%. In order to reduce the number of hits in each layer, the central detector zone has to be made insensitive, to allow the beam to pass through. This will be done drilling a hole of 1-2 mm diameter. The hole is kept so small in order not to miss many forward produced tracks.

Read-out electronics

We will use the Berkeley SVXC chip, a 50 μm pitch, 128-channel device directly bonded to the silicon detector. Four chips per plane will be used (512 instrumented strips), they will be daisy chained. For the trigger purpose, only the digital information (address of the strip hit) will be used, thus the output of the read-out electronics for each plane is a sequence of 9-bit numbers. The SVXC chip has a built-in sparse data scan capability, the output addresses can be clocked at 10 MHz. The chip, originally designed for collider applications, can be equally well used in a fixed target environment [5].

The signal to noise ratio depends on the associated read-out electronics, each strip of the detector contributing about 1.1 pF/cm to the capacitance. The noise for a 8 pF input capacitance (SVXC quadruple sampling) is $\leq 1200\ e^-$, the resulting S/N ratio for a m.i.p. in 280 μm of silicon is better than 15. The electronics chain will allow to obtain the analog value of the energy release associated with a hit on a silicon plane, though at a lower speed than if pure digital read-out is used (~2 MHz). This will be used for monitoring purposes. It is still under study if this more refined information (which via charge division technique allows to reach a precision of 8 μm) is needed for the collider application.

Associative memories

The output of the read-out electronics is fed directly into the associative memories which compare the measured hit configuration with all the possible patterns corresponding to straight tracks coming from the interaction region. The associative memories have been described elsewhere.[3]

The chip which will be built for this test experiment has a "word length" of 5 x 10 bits, and a depth of 128 words, for a total of ~6000 bits. The chip can accept input at 10 MHz.

It is planned to perform the pattern recognition of the tracks separately in the two xy and xz projections, each one featuring 5 silicon layers.

The number of possible physical patterns is:

$$N = BW * NSTRIPS * NPLANES$$

where BW is the beam (or interaction region) width in units of strip pitch
NSTRIPS is the number of strips on the last plane

Thus per each view, assuming an effective BW of 2 mm (35 pitch units), N~90000.
The associative memory chips will be housed in VME double-Eurocard boards. Each memory bank (xy view and xz view) can be accomodated in 11 memory boards, housed in one crate.

Since the comparison process inside the associative memories takes place in few nanoseconds, the time needed to carry out this phase of the pattern recognition process is essentially the time of feeding the associative memory banks with the strip addresses from the read-out electronics, plus the time of reading out the valid patterns' information. With an average of 15 charged tracks per event in the sensitive area of the detector, this time is of the order of 2.5 μsec. For collider applications, it is important to reduce this number to a minimum. This will be done within the limits allowed by this fast developing electronics technology.

The number of patterns to be stored in the associative memory banks can be reduced if not all the strips enter separately in defining the coordinate of the hit on the silicon plane. In the outer regions, where the track density is lower, we plan to cluster 2 to 4 strips together into so called "superstrips". Also, the granularity of the detector can in first approximation ibe nversely proportional to the distance from the target. The resulting number of patterns can consequently be reduced. The size and number of the superstrips is optimized in such a way to keep reasonably low the number of ambiguities in the subsequent track finding phase. Thus, instead of defining tracks, in this case the associative memory banks define "roads" which can include one or more tracks.

Track finding

If superstrips are used, the problem of finding the exact track parameters, and of solving the ambiguities inside the "roads" defined by the associative memories, is left to a second phase of the pattern recognition process. This will be performed by a dedicated hardware [6], developed on purpose together with the associative memories electronics (Fig. 2). The number of possible patterns inside a "road" can be 1000 to 10000 times smaller than the total number of patterns allowed for the whole detector. Thus the comparison process in each "road" can be run in a sequential way, but many CPUs can act independently on different "roads".

The amount of time needed to output the reconstructed charged space-tracks' parameters depends on track density within the roads. In our fixed target test application, we estimate that this time will be of the order of 3 μsec.

Vertex finding process

The algorithm which counts the secondary vertices of the event, after track reconstruction, will be run on a dedicated processor, which is still in the study phase in collaboration with the Advanced Computing Group of DD.

We are presently analyzing the structure of the algorithm, which now runs on IBM, in order to determine the amount of parallelism which can be achieved in some crucial parts of the computation, and to find the relative weight of the different fundamental computer operations (load, sum, branch,..) inside the algorithm itself. This work will be of great importance for the most effective choice of the hardware structure of the processor.

We are considering in particular two approaches to the problem. One solution is the use of a data driven processor, in the spirit of the FNAL E690's one [7], which couples the high computing power with an intrinsic pipelining of the events. A drawback is the limited flexibility in programming (mostly achieved via hardware connections). The second solution is an array of commercial processors (e.g. DSPs), possibly sharing a common data memory. The most time consuming parts of the algorithm seems to be easily parallelized, which fits with a distributed intelligence structure. The ease of programming is however balanced by the problem of optimising the data flux and access inside the array of processors.
A hybrid solution, which joins the computing power and programming flexibiility of the DSPs with the high throughput of a data driven device, is a promising way which we shall investigate.

Data acquisition

Since most of the monitoring and control tasks for the read-out chain have already been developed on VAX/VMS systems, we will base the Data Acquisition and detector control system on a similar computer, namely a microVAX II interfaced to VME and CAMAC.

MonteCarlo study

We briefly report on a MonteCarlo simulation of the trigger. The work has concentrated to date mostly on a fixed target set-up.

Generation and tracking

We assumed p-Nucleus interactions at 450 Gev/c incident beam momentum. The p-p interaction was handled by PYTHIA [8] (prerelease version). The shower particle generation (multiple p interaction in target nucleus) and the grey track (knock-on recoil protons) and black track (evaporation) production (typical of interactions on nuclei) were added according to experimental data [9]. The effect of this nuclear interaction was to increase by 70% the average charged track multiplicity of p-p interaction (Cu target considered).

The program could generate beauty (all decay channels allowed), charm and minimum bias events. Once generated, the tracks were handled by GEANT3. Major secondary interactions and decays were included, namely γ-ray conversion, δ-ray production, nuclear interaction in silicon and air, short and long lived particles decays (K's, π's, Λ's). Multiple scattering was simulated, and its effect on the χ^2 of the tracks reconstruction was found to be almost negligible for silicon planes' thickness of up to 400 μm (12 planes).

Geometry

The telescope structure as from the MonteCarlo simulation was 12 planes, 280 μm thick, 4 x 4 cm^2 sensitive area, 50 μm pitch, strip orientation 5y, 5z, 1u, 1v, 1 cm separation between any two planes, the first plane positioned at 4 cm from the 1 mm thick Cu target. The beam spot was gaussian in the transverse yz plane, with $\sigma_z = \sigma_y = 0.5$ mm.

In order to evaluate the effect of the error in the positioning of the strips, due to mechanical displacements of the planes in y and z, random displacements of the silicon layers were introduced, with gaussian distributions with $\sigma_z = 10$ μm (vertical) and $\sigma_y = 5$ μm (horizontal). These figures are at least twice as much what has been measured in similar existing set-ups [10], and in a real experiment they can only affect the trigger, since the actual displacements can be on-line monitored and made known to the offline analysis. No effect was found on B events acceptance (possible loss of tracks) and on the trigger background rejection (possible generation of fake vertices), despite a little degradation of the track fit.

The performance of the silicon detectors was derived from the test data of CDF. The detection efficiency was however pessimistically taken as 97%. Charge division between adjacent strips was considered (a valid hit can affect 1 to 4 strips depending on the impact angle of the track on the plane). A circular zone of 1 mm diameter at the center was assumed to be insensitive.

Analysis cuts

A momentum cut of 2 Gev/c was used in analyzing the events, i.e. only events where

$$P_t = (1/N \, \Sigma \, p_t^2)^{1/2} > 2. \text{ Gev/c} \qquad (N = \text{charged multiplicity})$$

were considered for the trigger. In our MonteCarlo, B events exhibit a p_t distribution which drops below 2 Gev/c (Fig. 3). Accordingly, in a real experimental set-up we would exploit a pre-trigger capable of selecting events obeying the above cut. Similar devices, based on scintillation counter techniques, are being tested and can operate at a rate of $\sim 10^6$/sec [11].

Trigger algorithm

It runs in two steps.
The first step consists of performing a χ^2 fit of the tracks in space, and in computing for any two tracks the minimum approach distance (d_{min}) in space, the relative angle θ, and the coordinates of the point of minimum approach (x_m, y_m, z_m).

This requires $N^2/2$ computations (N being the number of tracks reconstructed by the associative memory logic), and can be naturally parallelized. Note that a track is considered for this computation only if it is not originated after the first silicon plane (to avoid contributions from secondary interactions and long lived particles decays), if it has at least 3 points in each one of the xy and xz projections (to have a meaningful track fit), and if the unweighted χ^2 of the fit (Fig. 4) is less than $(23\ \mu m)^2$.
A two-track approach point is considered a "secondary vertex" if it satisfies the following cuts:

$d_{min} < d_{cut}$ $\qquad$ $d_{cut} = 35\ \mu m$

$\theta > \theta_{cut}$ $\qquad$ $\theta_{cut} = 0.07$

$0.5\ \text{cm} < x_m < 3.2\ \text{cm}$ $\qquad$ (fiducial region, ref. Fig. 5)

An event where more than 3 track pairs exhibit such a "vertex" is retained as a B event candidate. All the real B events which we can retain (even with the second refined step of the algorithm - see below) pass this first cut. The background rejection is of the order of 1:2000 at this level (Fig. 7).

For events passing this filter, one proceeds to a second step, where one finds a "primary vertex" clustering together all track pairs which have similar minimum approach points, consistent with being inside the target $(-1\ \mathrm{mm} < x_m < 1\ \mathrm{mm})$.

Since a track can contribute more than once in defining this primary vertex (because it can belong to several track pairs), the precision of the primary vertex" definition is not improved if one weighs the contribution to the x-coordinate of each track pair with a weight inversely proportional to the angle between the tracks (this would give more weight to large angle tracks). This peculiar method of defining the primary interaction point has a good precision (Fig.6), and is intrinsically independent of the knowledge of the transverse coordinate of the beam particle.
Once the primary vertex has been found, all the remaining vertices are considered "secondary vertices". An event is retained if it has at least 4 such secondary vertices.

In this way, we are able to retain 55% of the B events, 1% of the charm events, and to reject the minimum bias background to better than $2 : 10^5$ (Fig. 8).

An impact parameter approach does not give better results, and in any case the impact parameter has to be computed with respect to the primary vertex, whose computation is the time consuming part of the algorithm.

Research program

Some silicon detector planes and read-out chips are already under test in Pisa.

We plan to install the telescope and the read-out electronics (complete) on a CERN beam by summer 1989, then proceed to real beam condition tests and acquisition of data for comparison with MonteCarlo expectations.

The associative memories chips will be ready and tested by end 1989, and then mounted on the beam set-up to be operated early 1990. High rate performances of the whole trigger requires a fast processor for the selection algorithm. We plan to test this by mid 1990.

When this step will be achieved, the trigger telescope should be capable of selecting B event candidates at a rate of 10^5/sec. We shall then take data in association with a real detector, and analyse them in order to verify the consistency of the trigger

If no other possibility will exist, the best way to date seems to move the apparatus at the Ω spectrometer, and acquire data in parallel: the slow acquisition rate of Ω will not test the speed capability of the apparatus, but will allow a full reconstruction of the event and a consequent check of the acceptance figures quoted above. The cuts would have to be modified to run on charm. It is conceivable to be ready for this verification step in the second half of 1990.

Test beam requests

We would like to a proton beam of up to 450 Gev/c (alternatively, a π beam of similar energy could do, provided that the μ halo can be stopped).

Ultimately, the beam has to reach a high intensity (10^6/sec), though this is not needed for the 1989 tests. The spot should be as small as possible, with a gaussian transverse projection of $\sigma < 0.5$ mm. If a microbeam can be obtained (at least in one transverse axis) this would greatly help. However, the beam should still be contained in a spot of 1-2 mm diameter at about 16 cm from the focus.

While running, a 1mm Cu target will be put before the telescope. No other sources of background are acceptable in front of our apparatus, which on the other hand will be mounted on a solid basement and can be displaced out of the beam line.

Conservatively, a total of 4 weeks of running time towards end 1989 seems appropriate.

We need a barrack where to house the electronics (~6 racks), the on-line computer, and a minimum of working space. Default Index (2) and Ethernet (1) connections are also essential.

Requests for further support

The CDF Group (as well as other groups using Si detectors) has expertise and tooling to design and build a suitably precise support structure for the telescope and we intend to use it. Still, we ask for a limited technical support by CERN, possibly from EF Division personnel, for the mounting of the system on the floor and for its remote control.

Since only two of us share other interests in CERN, we need some amount of office and lab space, following the standards adopted by CERN in similar cases. Our immmediate needs are of the order of one or two offices and a small lab, even located in the Labo 1 area.

Index or Decserver connections should however be provided. Easy access to an Ethernet line for a forthcoming VAXStation would be appreciated.

Most of the lengthy computations which lead us to this proposal were carried out in Pisa (INFN and Scuola Normale Superiore IBM and VAX computers). Though our work will be in coincidence with the first data from LEP, we would like to ask for a limited amount of computer time at CERN. This amounts, for 1989, to 1000 IBM168 hours on the CERN VM System, plus 150 cylinders fixed disk space, and a limited amount of space (3-4 accounts x 5000 blocks) for software development on VXCERN.

For the tests in Pisa, we are now exploiting a local microVAX II which we cannot think of displacing to CERN before 1990. It would be of great help if we could find, on a temporary loan of 6-9 months starting June 1989, a microVAX II with a 71 or 159 MB disk.
We hope that this can be found in the DD Pool, or someway provided by DD as part of the common effort in developing an advanced processor.

Finally, we will need the online software support of the DD OC Group for the running in and operation of the Data Acquisition System.

References

[1] J.M.Crespo et al. - "High statistics study of charm and beauty flavoured particles using the NA14 Spectrometer and a data-driven processor in a Hyperon beam" - Letter of Intent - CERN/SPSC-I/166 10 Oct. 1987

[2] S.A.Kleinfelder - "The LBL Silicon Vertex Detector Readout Integrated Circuit: an introduction" - LBL Nov. 16,1986

T.Zimmermann, R.Yarema - "Test results of two LBL SVX silicon strip detector readout chips and the LBL RADTEST chip" - CDF/DOC/SEC_VTX/CDFR/593 Dec. 1987

[3] M.Dell'Orso, L.Ristori - "The pattern matching machine" - Proc. of the "Workshop on High Sensitivity Beauty Physics at Fermilab", FNAL, Nov. 11-14, 1987

[4] F.Bedeschi et al. - "CDF Silicon Detector Prototype Test Beam Results", INFN - PI/AE 88/11

[5] J.S.Russ, R.J.Yarema - "Study of the LBL CMOS integrated amplifier/discriminator for randomly timed inputs from fixed target experiments" - presented at NSS/IEEE-88, Orlando, Fla., Nov. 1988

[6] P.Battaiotto, M.Budinich et al. - "A processor array dedicated to a fast track finding algorithm" - to be published in Proc. of Computimg in HEP 89, Oxford, April 10-14, 1989

[7] H.Cumitz, R.Hsiung, B.Knapp, W.Sippach - Proc of the Conf. on "Instrumentation for Colliding Beam Physics", SLAC-250, June 1982

W.Sippach, G.Benenson, B.Knapp - IEEE/Trans. on Nucl. Science 27, 578 (1980)

[8] H.Bengtsson, T.Sjostrand - "The LUND Montecarlo for Hadronic processes - PYTHIA version 4.8" - UCLA-87-001 Jan. 1987

[9] I.Otterlund et al. - "Nuclear interactions of 400 Gev protons in emulsion" - Nuclear Physics B142 (1978) 445-462

[10] M.Dameri, L.Rossi, WA82, private communication

[11] V.Flaminio, WA84, private communication

Figure captions

Fig. 1 Schematic lay-out of the set-up

Fig.2 Schematic structure of the track-finding machine

Fig.3 Pt distribution of events with beauty production (no cuts).

Fig.4 Unweighted χ^2 distribution (in cm) for the fit of the straight lines. Detector resolution is about 15 μm in digital mode read-out.

Fig. 5 Longitudinal distribution of the secondary vertices (B and C decay). The fiducial region is between 0.5 and 3.2 cm (see text).

Fig. 6 Error in the reconstruction of the longitudinal coordinate of the primary vertex.

Fig. 7 First step of the algorithm, acceptance cut is: ≥ 4.

a) distribution of the no. of vertices (track-pairs' origins) inside fiducial region for events with beauty production.

b) distribution of the no. of vertices (track-pairs' origins) inside fiducial region for background events.

Fig. 8 Second step of the algorithm, using reconstructed "primary vertex". Acceptance cut: ≥ 4.

a) distribution of the no. of reconstructed "secondary" vertices for events with beauty production.

b) distribution of the no. of reconstructed "secondary" vertices for background events.

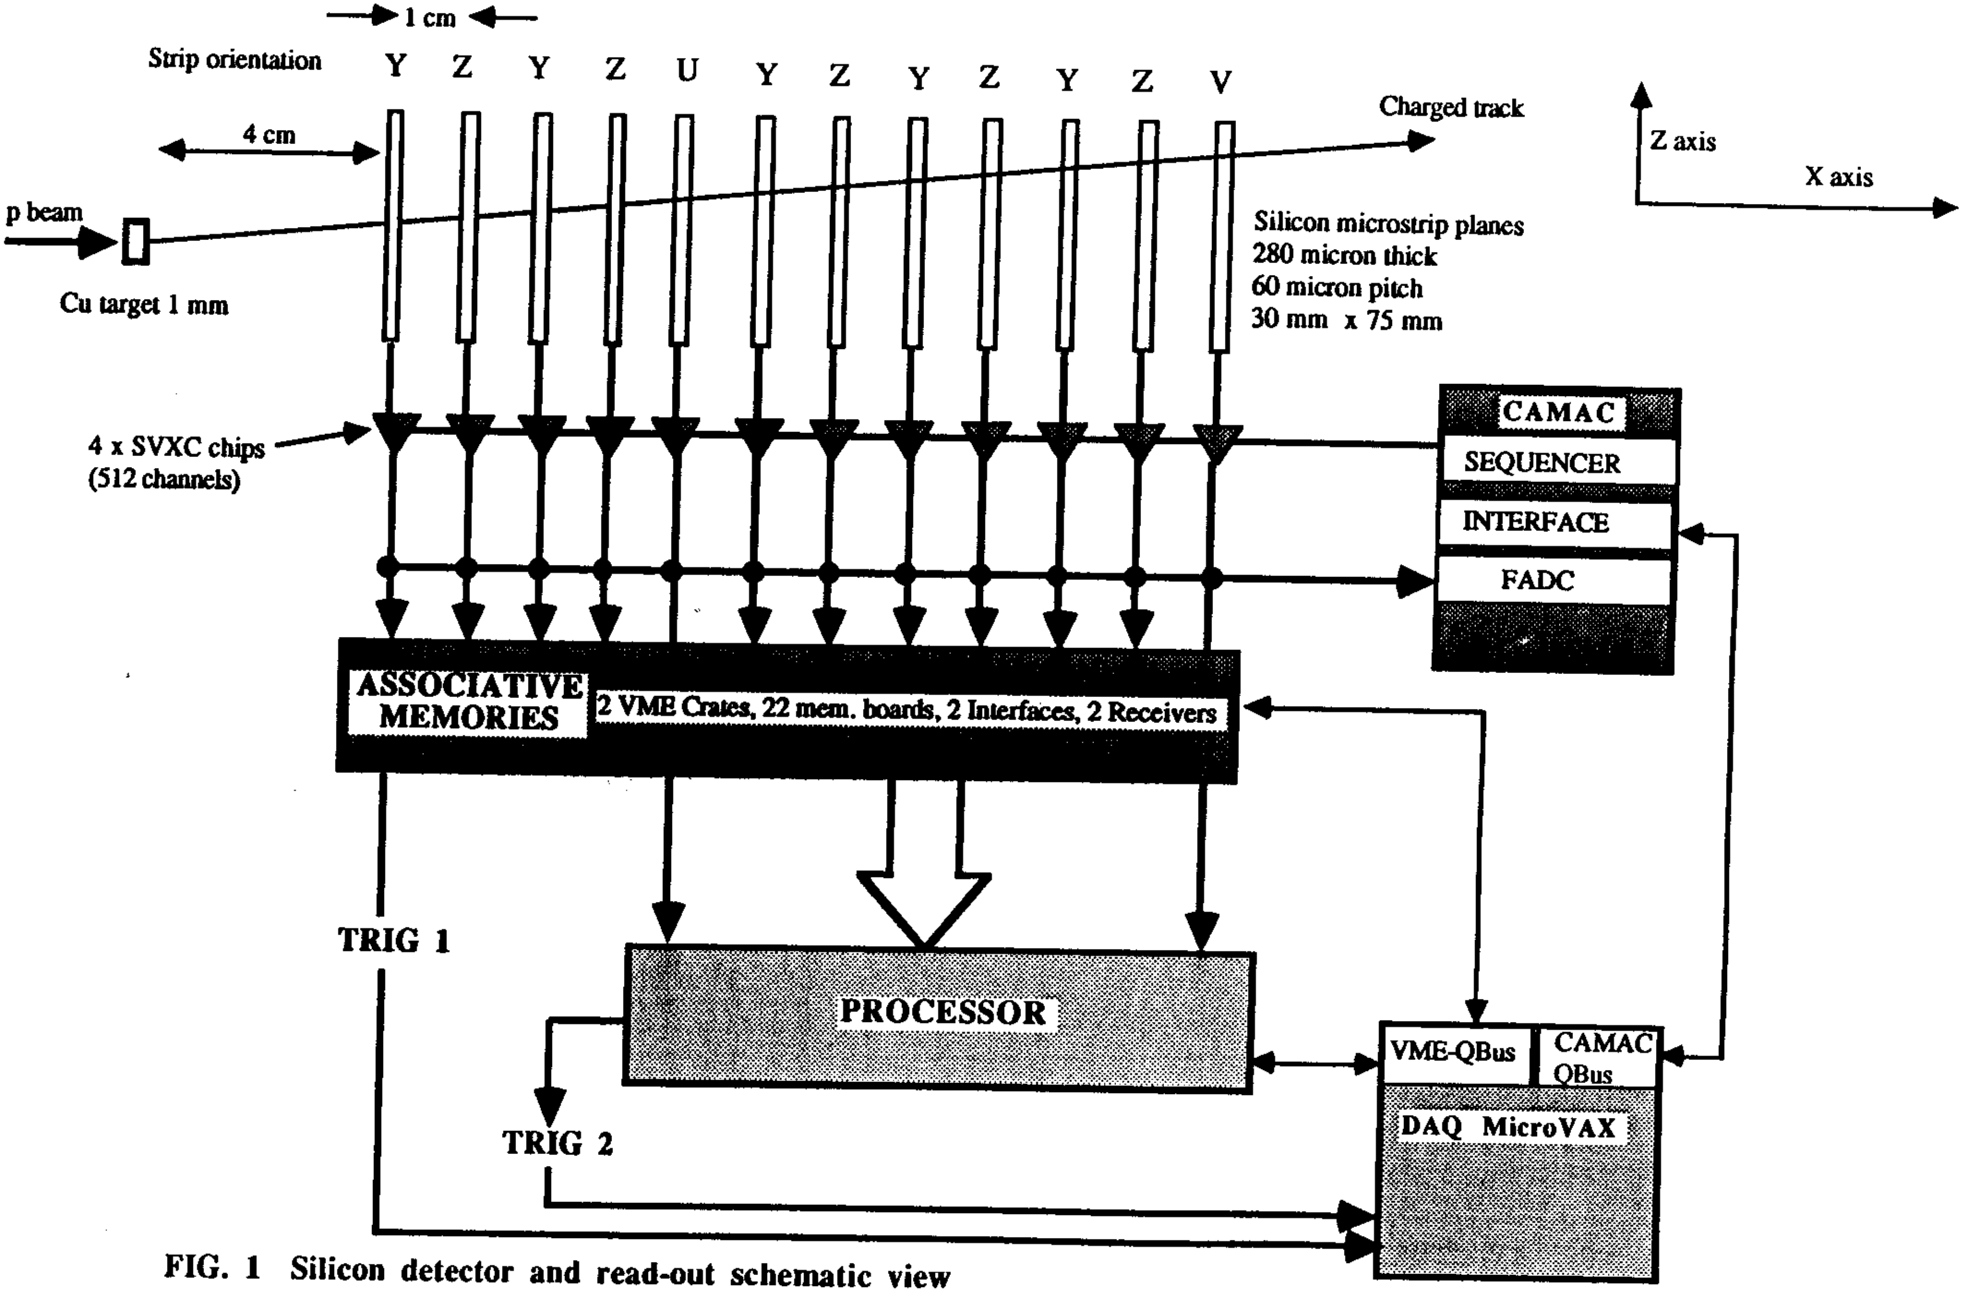

FIG. 1 Silicon detector and read-out schematic view

Fig. 1

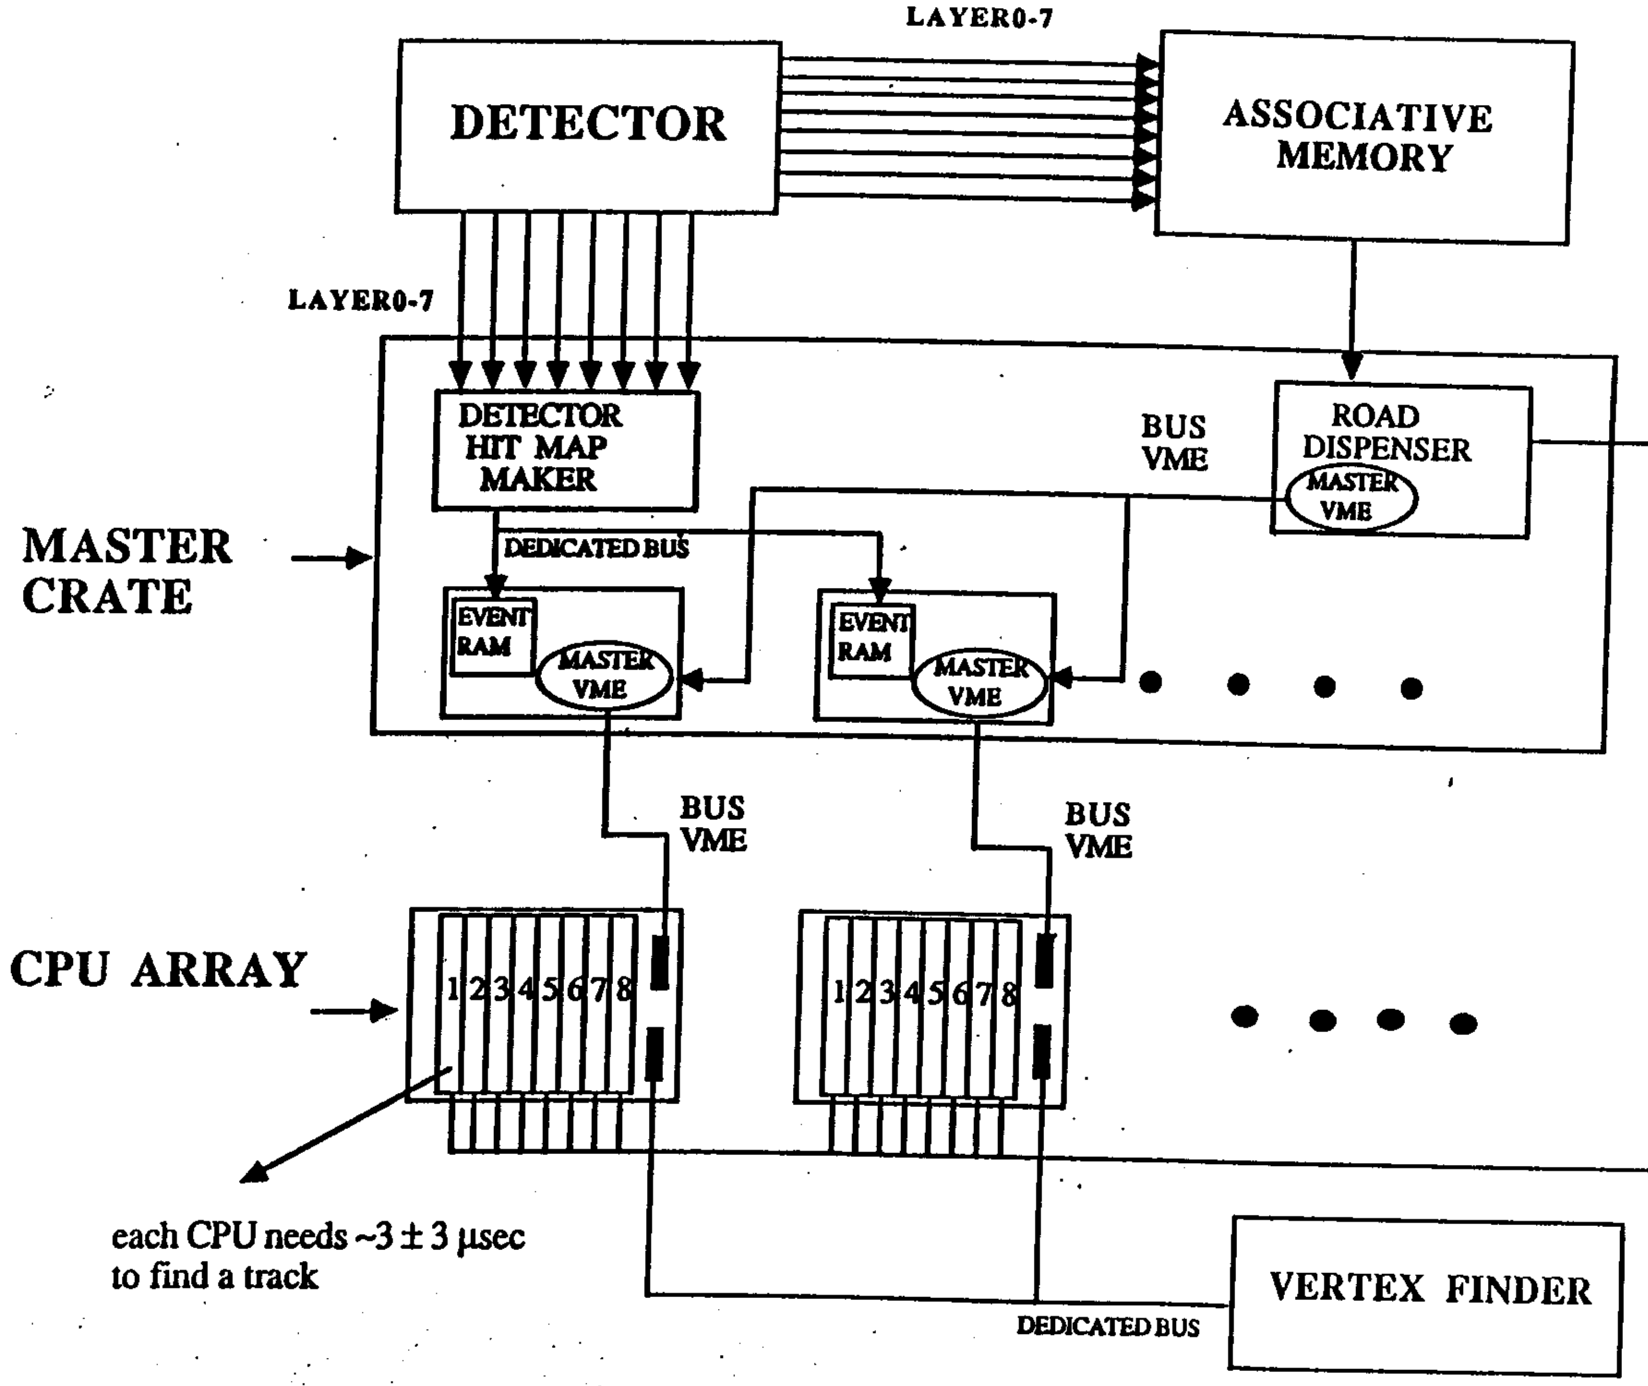
LAYER0-7
DETECTOR
ASSOCIATIVE MEMORY
LAYER0-7
DETECTOR HIT MAP MAKER
ROAD DISPENSER
MASTER VME
BUS VME
MASTER CRATE
DEDICATED BUS
EVENT RAM
MASTER VME
EVENT RAM
MASTER VME
BUS VME
BUS VME
CPU ARRAY
1 2 3 4 5 6 7 8
1 2 3 4 5 6 7 8
each CPU needs ~3 ± 3 µsec to find a track
VERTEX FINDER
DEDICATED BUS

Fig. 2

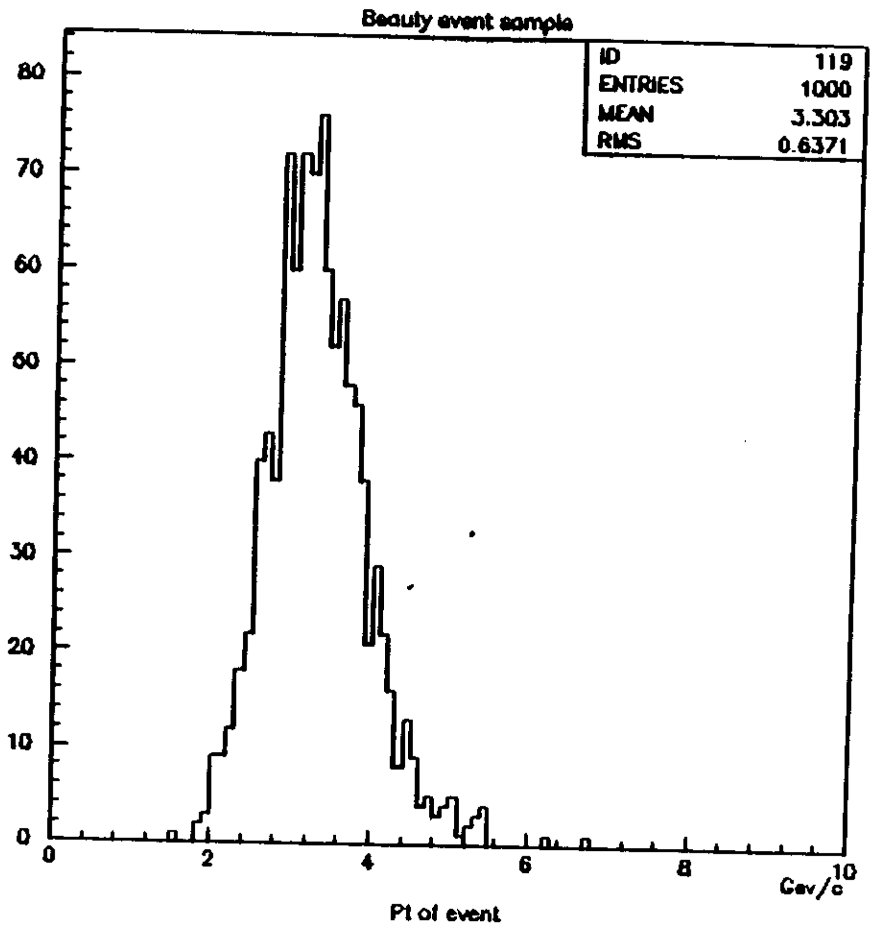
Beauty event sample
ID 119
ENTRIES 1000
MEAN 3.303
RMS 0.6371
80
70
60
50
40
30
20
10
0
0
2
4
6
8
10
Gev/c
Pt of event

Fig. 3

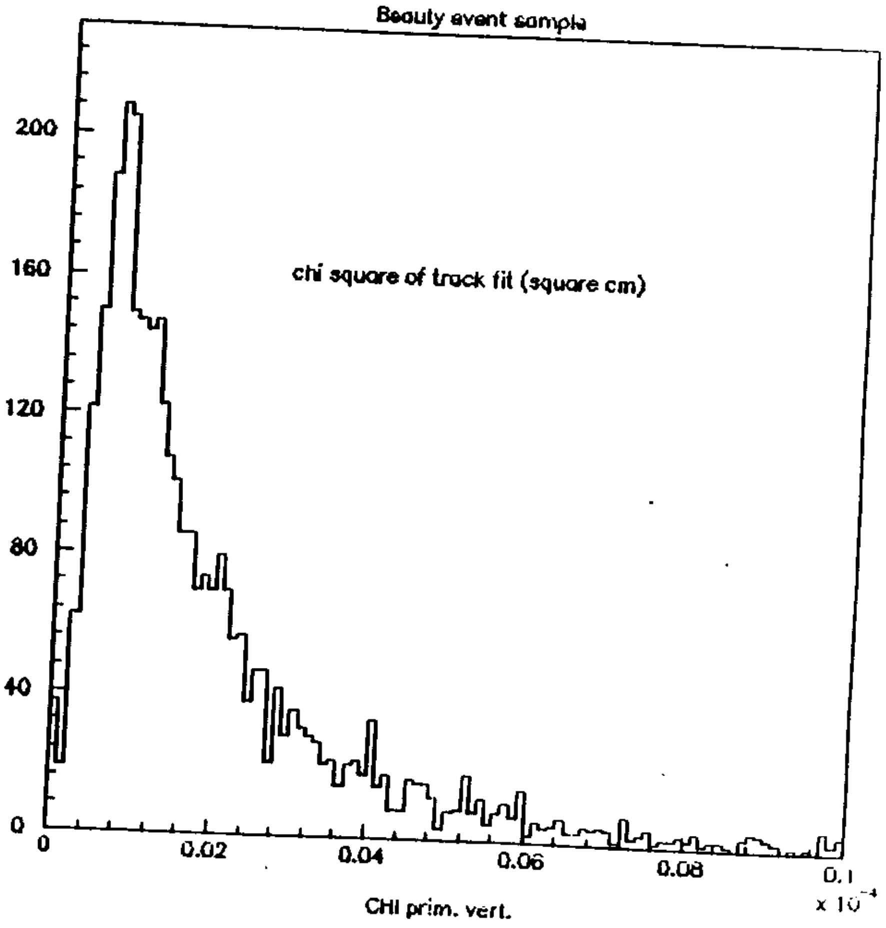
Beauty event sample
chi square of track fit (square cm)
200
160
120
80
40
0
0
0.02
0.04
0.06
0.08
0.1
x 10^-4
CHI prim. vert.

Fig. 4

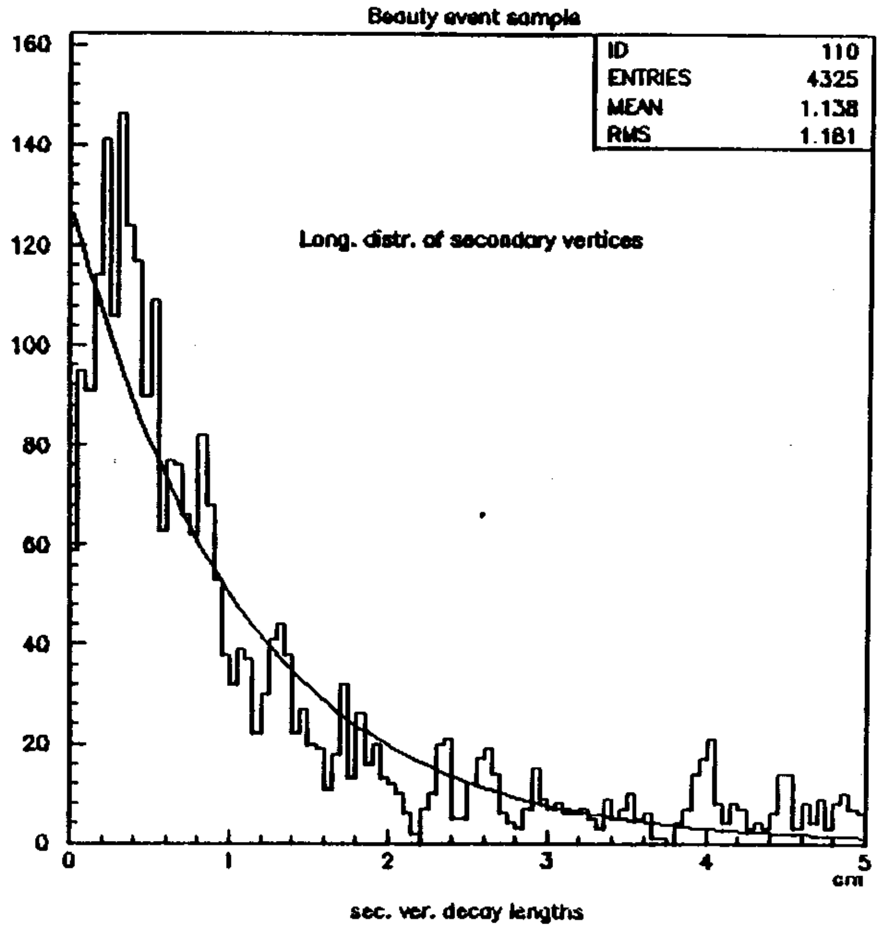
Beauty event sample
ID 110
ENTRIES 4325
MEAN 1.138
RMS 1.181
Long. distr. of secondary vertices
160
140
120
100
80
60
40
20
0
0
1
2
3
4
5
cm
sec. ver. decay lengths

Fig. 5

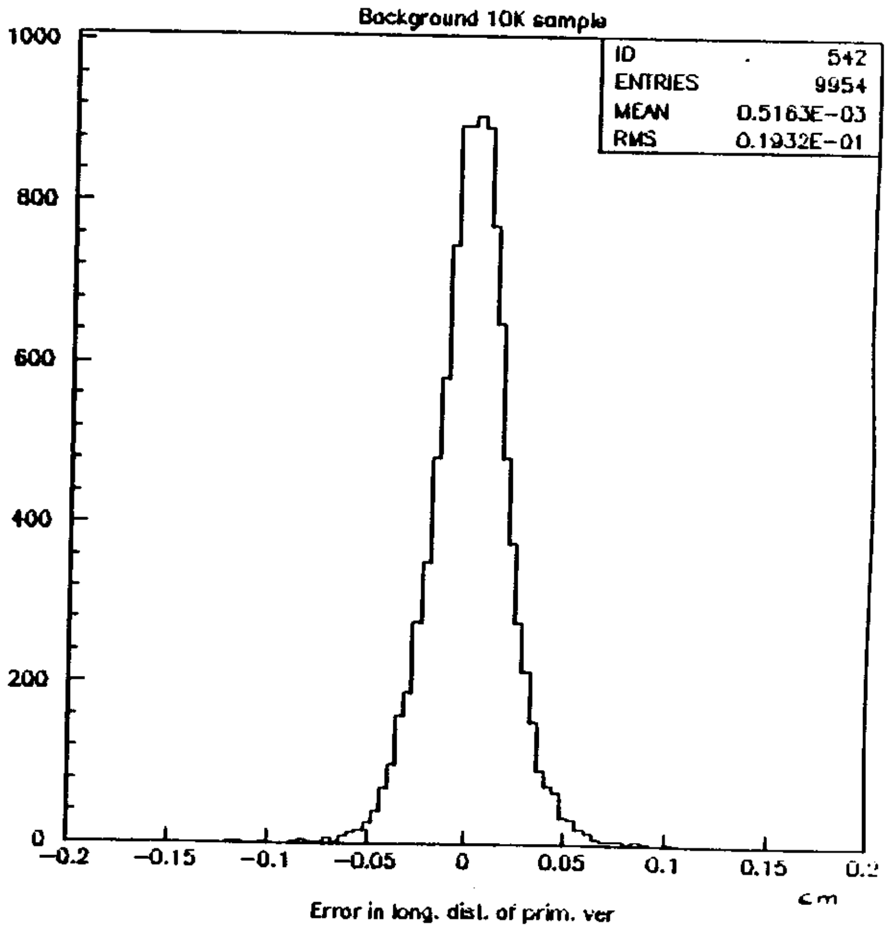
Background 10K sample
ID 542
ENTRIES 9954
MEAN 0.5163E-03
RMS 0.1932E-01
1000
800
600
400
200
0
-0.2
-0.15
-0.1
-0.05
0
0.05
0.1
0.15
0.2
cm
Error in long. dist. of prim. ver

Fig. 6

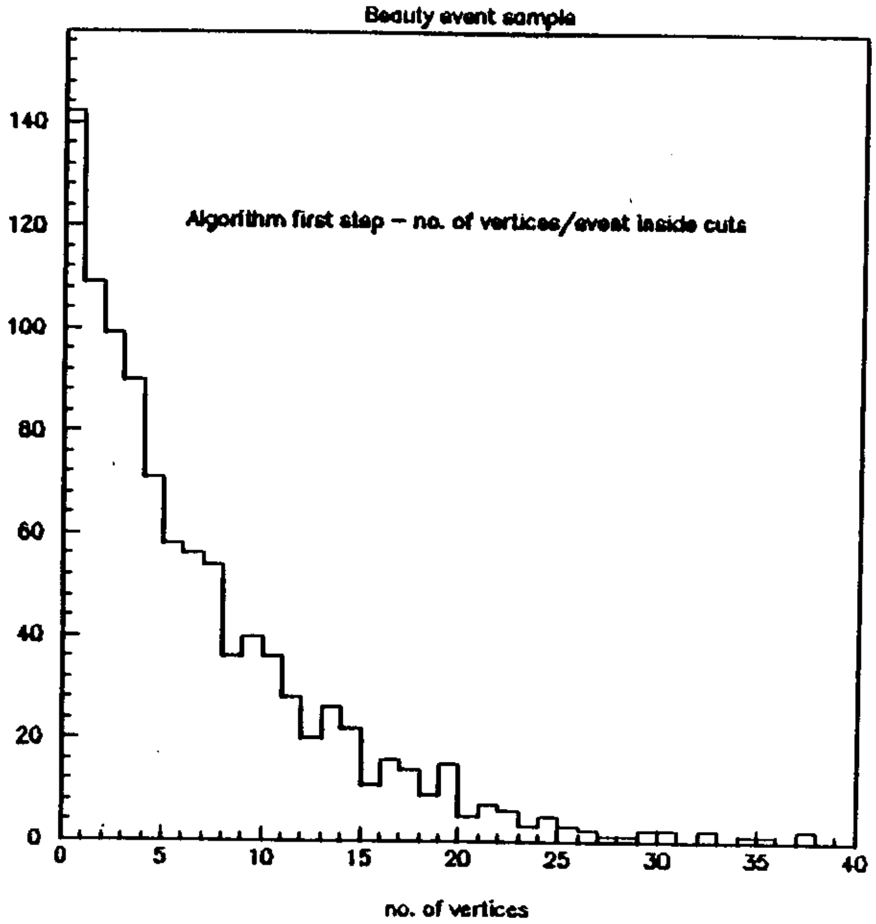
Beauty event sample
Algorithm first step – no. of vertices/event inside cuts
140
120
100
80
60
40
20
0
0
5
10
15
20
25
30
35
40
no. of vertices

Fig. 7a)

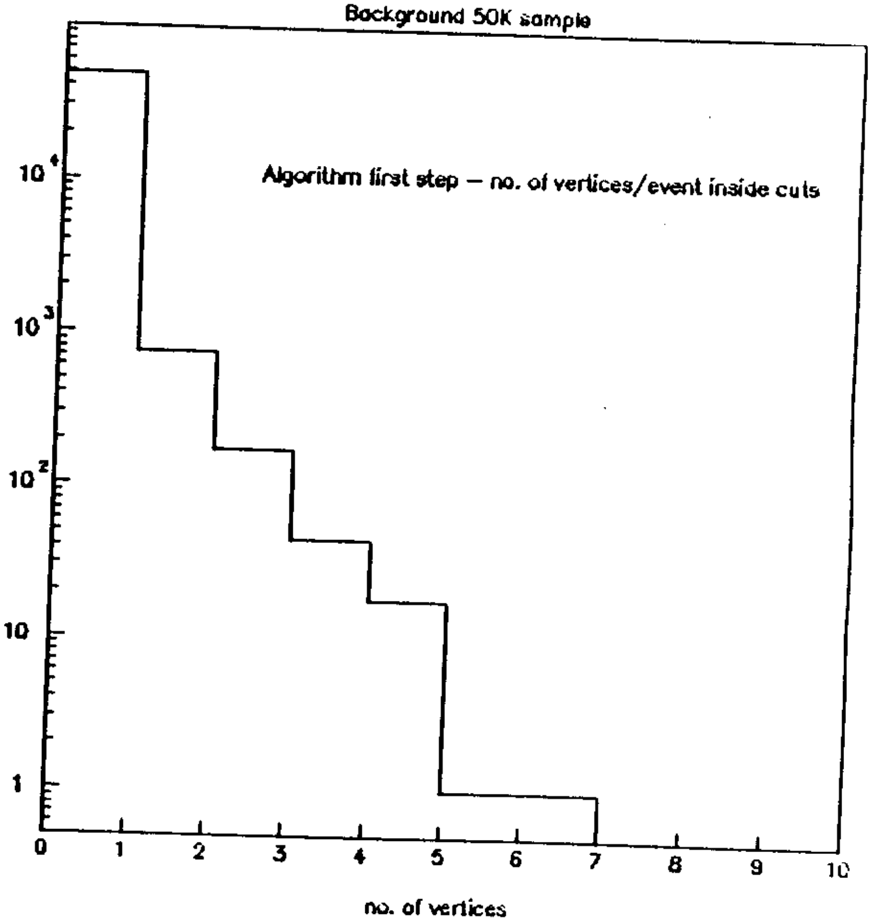
Background 50K sample
Algorithm first step – no. of vertices/event inside cuts
10^4
10^3
10^2
10
1
0
1
2
3
4
5
6
7
8
9
10
no. of vertices

Fig. 7b)

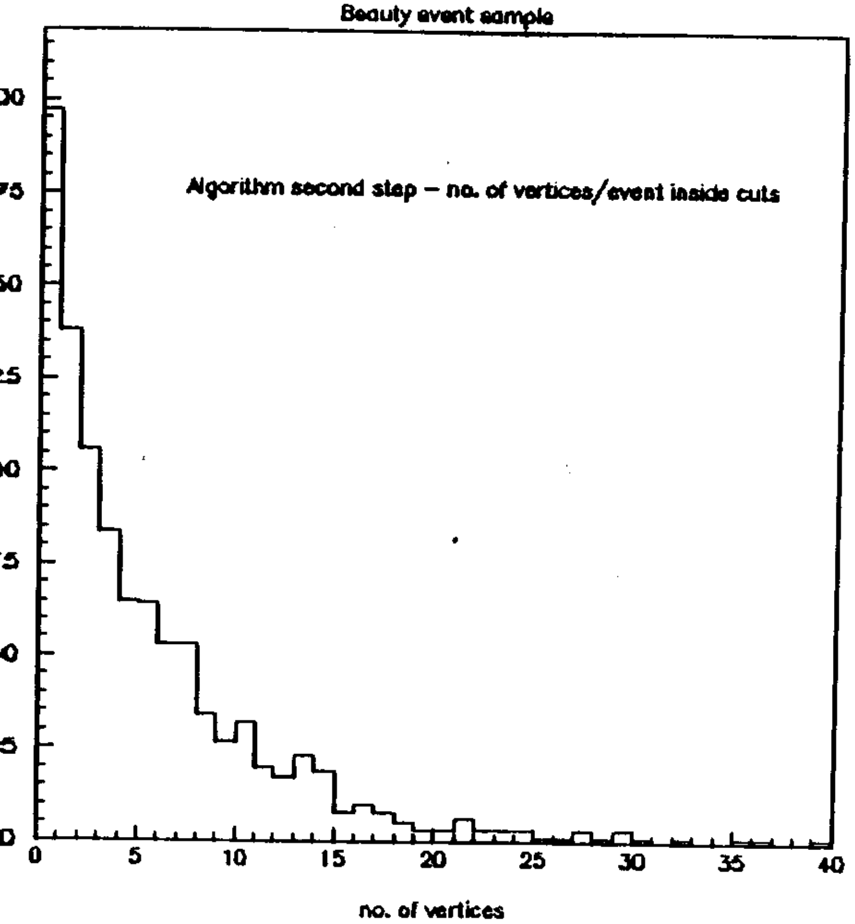
Beauty event sample
Algorithm second step – no. of vertices/event inside cuts
0
5
10
15
20
25
30
35
40
no. of vertices

Fig. 8a)

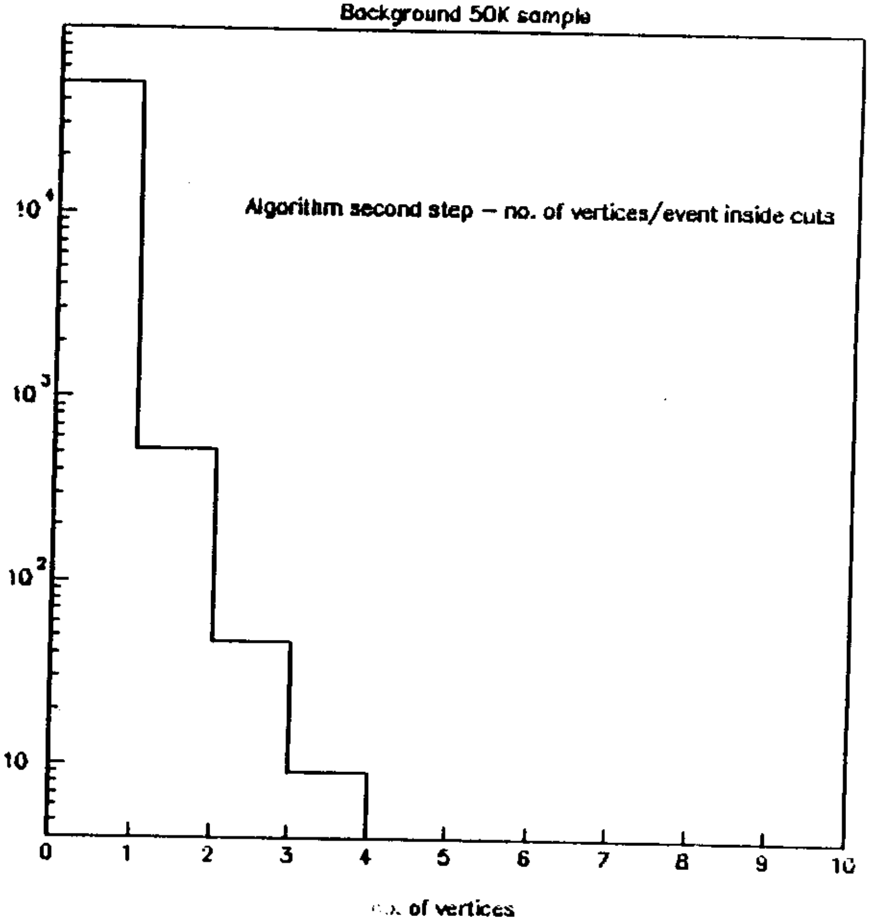
Background 50K sample
Algorithm second step – no. of vertices/event inside cuts
10^4
10^3
10^2
10
0
1
2
3
4
5
6
7
8
9
10
of vertices

Fig. 8b)