# The CERN approach in comparison with the TS

Bart Van De Vyver

May 8, 1998

**Abstract**

This note explains the philosophy of the CERN microscope project, and gives some details of the database used. We also present the properties of the new optics.

## 1 Introduction

Rather than giving a status report for a usual Collaboration Meeting, I will try to discuss the CERN project in more general terms, stressing the future plans. Of course, the current status is important in the sense that only large enough statistics and thorough understanding of the current techniques (as defined by Nagoya) allow one to identify problem areas and to test any proposals for improvement. The title of my talk deliberately refers to the Nagoya TS since the enormous progress made at CERN over the past year has only been possible thanks to having the TS. Not only did it provide us with a working system on a short timescale, it also served as a reference and example in setting up the new systems. With the Nagoya TS as starting point, the main differences at CERN are:

- the DAQ philosophy
- the on-line database
- the new optics
- the usage of commercial DSP hardware
- a (proposal for a) different scanning flow

Some of these points are already quite advanced, some only exist as ideas for the time being. I will discuss each of these in the following, focussing on the technical aspects. However, it is important to keep in mind that the underlying motivations are always driven by physics. In particular, the objective at CERN is to give priority to oscillation physics, considering (at least) all the tau topologies considered in Nagoya with (at least) the same efficiency. But, in doing so, we try to set up the system such that any data relevant to other physics is taken, either in parallel or in a second pass where we envisage to follow down all tracks found at the vertex to allow for momentum measurements and particle ID.

## 2 The DAQ philosophy

As the scanning capacity is continuously growing, it seems that for example in Nagoya the limiting factor is becoming more and more the logistics of running 6 microscopes in parallel, each producing about 30Mb/s of raw data. In particular, with the current speeds, a large fraction of time is spent in plate changes. Of course, problems at CERN are somewhat more modest, but still we felt things could be improved using the experience of the CERN group with DAQ techniques. All microscopes have a dedicated PC (winNT) which runs the acquisition program itself, however several tasks have been delegated to other computers :

- the user interface is a Tcl/Tk panel, communicating with the acquisition program through the dispatcher running on a UNIX machine (Johan Uiterwijk)
- the data is written into the database (again running on a UNIX machine) on-line, taking advantage of Fast Ethernet connections
- in the future, two additional tasks will be implemented, partly on the acquisition PCs partly on the Unix server : slow control and on-line image analysis.

In order to reduce the time spent in changing plates, we have developed a local oil system (Piero Zucchelli, Jean Dupraz, Guy Van Beek) which eliminates the cleaning of the plates after acquisition. Furthermore, every emulsion sheet is mounted in its own a single plastic holder.

## 3 The on-line database

The requirements of any database for the scanning are :

- the possibility to access data both per event (time ordering) and per plate (space ordering)
- the possibility to directly store both numerical/textual data and images
- the possibility to handle heterogeneous data (e.g. track segments found by track selector, manual scanning, image analysis have a common base, and fundamentally different additional information)
- the possibility to trace back any on-line scanning
- the possibility to easily interface to C++ (all acquisition programs)

To satisfy these requirements, we have chosen to set up an Objectivity database. In particular, the various problems are solved as follows:

- The elementary data is organized in independent 'objects', not part of any particular table. The individual objects contain bidirectional associations to other objects, thus providing the various access modes. For example a 'Segment' object is 'FoundOn' a particular plate, 'PredictedBy' a particular prediction, 'Composes' a particular track,... Of course, some of these associations are one-to-many or many-to-many.

One level higher are 'containers' which group a set of objects. However, this grouping is only important for the locking (concurrent access by different programs for both read and write operations) and clustering (physical storage is clustered to provide fast access to objects belonging together). The ordering in containers is transparent in the sense that any segment belonging to an event can be accessed directly from the event, independent of the fact whether it is residing in an 'Event' container or in a 'Plate' container. The former is convenient for analysis, the latter for on-line scanning.

- As far as images are conCERNed, we foresee two storage levels: pixels belonging to grains ('pre-clustered data') and compressed images in MPEG format ('raw data'). For pre-clustered data, Objectivity allows one to store STL containers directly, and this corresponds to the format we use on-line. For the raw data, the possibility to store binary data directly in Objectivity is under investigation. The alternative would be to use the file system and keep only pointers to the image files.

- The normal C++ class hierarchy of base classes and derived classes is supported also for the data storage classes. As such, we have for example a base class Segment containing information on slope, position, timetag, ... and a set of derived class SegmentCS, SegmentBulk, SegmentManual corresponding to the information typical to any particular algorithm: pulseheight, base angle, grain data,.. Any algorithm only interested in general quantities (calibration, most analysis,...) can access SegmentBulk, SegmentManual,.. objects exactly the same way the basic Segment objects are accessed.

- Since the on-line program writes directly into the database, all data available on-line can be kept for future reference. Two examples: the set of marks to fix the plate on the stage, are in a Marks object which is 'versioned'. By default, one obtains the latest objects, but any previous scanning can be rolled back in case any mistakes are suspected. The calibration of the plate is done on-line but does NOT overwrite the prediction data; the latest calibration is always used to transform predictions and segments at read/write time into the appropriate reference system ('stage' for scanning, 'Chorus' for analysis)

- Objectivity programming is hardly any different from C++ programming; furthermore, we have provided a layer between the on-line programs and the database using a mirror of all the database classes.

# 4 The new optics

The objective currently used in most scanning labs of the Collaboration, is a Tiyoda 50x lense with NA 0.85 working in the green (540nm), used in conjunction with a Nikon condenser with NA 0.85. Especially the large working distance (1.2mm) makes this objective very different from typical objectives which are commercially available for biology, metallurgy,.. It has been developed specifically for the task of emulsion scanning in a collaboration with the Nagoya lab (Prof.Niu) and Tiyoda. The choice of green light corresponds to the manual scanning prevailing at the time. We have decided to start the design of a new set of optics to satisfy the changing requirements:

- since most of the scanning is now automatic, blue light (430 nm) gives higher resolution
- a larger FOV (0.5 mm diameter) is required to fully profit from larger CCD sizes as well as to investigate the possibility of lower magnifications
- a higher NA (objective design value 1.05) not only increases the resolution, but especially reduces the depth of the focal field. A smaller depth of field translates into higher processing speed (less data to be treated) and better signal to noise ratio (same number of track grains, smaller number of fog grains), see also Piero's note on filtering
- to reduce the singal/noise ratio on the CCD camera, it is necessary to shutter the light source with on/off ratios of the order of 10

A new set of optics meeting these requirements are currently being designed by Jenoptik at Jena. This 'new optics' includes both the illumination device (light source + condenser) and the imaging device (objective). By design, the magnification will be variable, using a first objective focusing at infinity and a second objective (up to factors of two of both magnification and demagnification) forming a real image. A first prototype is set up on an optical bench in Jena. The mechanics is being finalized and we expect to have a first prototype at CERN by the end of June.

# 5 The usage of commercial DSP hardware

Each of the three new microscopes is equipped with a DSP board built by LSI around the Texas Instruments C81 chip. These boards have 32Mb on-board memory, a digital frame grabbing section and the chip itself (a C80), containing a RISC processor and four parallel processors. Our first aim with the DSP is to perform the filtering/grain finding of the image in real-time (30Hz frame rate for Mpixel images). The filtered image then forms the starting point of further processing which can go in several directions:

- a 'copy-in-principle' of the Nagoya TS, ie. to shift and sum a certain number of images according to a predefined offset, given by the prediction (Ioannis Papadopoulos), this shift-and- sum would be done typically on the host PC (Pentium) but variations are possible. One of the main advantages of the Nagoya TS algorithm is the fact that the actual clustering only needs to be done on the sum.
- more general tracking, commonly referred to as 'image analysis' (Johan Uiterwijk); depending on the time requirements, the clustering and tracking, could be done immediately on the PC or on another machine running in parallel while data-taking in TS mode continues For the setup of these new systems, many ideas are based on the existing (Nagoya) hardware and in particular the continuous cross checks at all levels (raw data $\rightarrow$ pixels belonging to grains $\rightarrow$ shifted sums $\rightarrow$ clusters corresponding to tracks) have been extremely useful. So far, the algorithms have not been fully tested on the digital CCD camera and DSP, but this will be done in the near future. However, a proof in principle of the technique has been given.

One of the main reasons for using commercial hardware rather than custom-made electronics, is the rapid progress in this domain. To profit from these developments, it seems best to concentrate on "portable" solutions, typically written in C, with some critical parts in Assembler. We already have an evaluation board based on a new generation of DSPs (the 6xx series from Texas). They use a very long instruction bus (vlwib) of 256 bits, which allows for 8 instruction executions per cycle running with a 200MHz clock. A board will be available in the near future with two of these DSPs. As an example, on this new board, one of the two processors could do a 3x3 convolution in 6 cycles/pixel. The board provides both COM ports and PMC slots, which allow us to use the current DSP as frame grabber card and then pass the data directly to the new DSP with a DMA transfer on the PCI bus. Digital frame grabbers in the PMC format are also available so that grabbing and analysis could take place on the same board, however these frame grabbers have not yet been ported to the 6xx series. A whole chain of image processing could thus be set up to treat the data in real-time.

# 6 Scanning flow

At the moment, the only production scanning running 24 hours per day is done on the Nikon, equipped with the old TS and changed to conform to European standards (motion controller, frame grabbing card, PC, operating system,...). Very little bulk scanning has been done so far and therefore it is highly speculative to envisage alternative schemes. The following is to be taken as a proposal which will be confronted with real data in the near future to judge its feasibility and scanning load.

Up to special sheet, procedures are identical to what is done in Nagoya, especially now that also in Nagoya all charged tracks are being scanned (and not only the candidates for a kink on the scanback track). The only modification may be in the rejection of low fit probability tracks, since we believe this is a very powerful criterium to identify bad ('ghost') tracks. In particular, the real efficiency calculated according to Gustaaf's method drops to 0% for tracks with $P < 0.001$. The scanback procedure is again very similar to the procedure in Nagoya. It is foreseen to scan only the last $100u\,\mu$m pstream in each plate for each scanback track using a TS-like algorithm, ie. an algorithm only sensitive to a predefined slope. However, we would like to make a somewhat looser 'vertex plate trigger'. In Nagoya, a vertex plate is identified when the scanback track has been missed in two consecutive plates. This means that any plate has to be mounted at least twice. We envisage to consider 'immediately' any track missed in the TS-like scanback procedure while the plate is still on the stage. For this, it is obvious the efficiency needs to be high enough.

From Nagoya data, 4% of the tracks are not found, 2% correspond to genuine neutrino vertices and 2% correspond to various sources of inefficiencies. Therefore, it seems feasible to consider a multi-level trigger.

- LEVEL 1: track missed → take image data over the full depth for one view (new optics with new camera: 300x300) $\mu$m displaced according to the slope for the downstream half. This data is to be transmitted to an image analysis program running in parallel on a different machine so as not to load the acquisition PC. The microscope can continue with the next event and the 'almost on-line' programming has both time (typically 1 out of 20 events need to be treated), CPU (this program could run on a powerful remote processor) and data (both the recent data as well as the data from the previous plate) available to reach a decision.

- LEVEL 2: depending on the outcome of the image analysis
  - $\rightarrow$ track was missed by TS: inefficiency, continue as before
  - $\rightarrow$ track shows a kink: use parent track as prediction for next plate, acquire image data in 9 views around the kink (1mmx1mm)
  - $\rightarrow$ track reaches neutrino vertex: acquire image data in 9 views around the kink (1mm $\times$ 1mm) or the vertex
  - $\rightarrow$ track not found, neither upstream or downstream: acquire image data for manual checks
- LEVEL 3: manual checks As for the manual checks, various ideas exist to try and make manual scanning independent of microscope hardware (so-called manual scan on images) but they have not been tried out so far.