REMOTE CONTROL OF A STREAK CAMERA FOR REAL TIME BUNCH SIZE MEASUREMENTS IN LEP

A. J. Burns, J. J. Savioz, J. C. de Vries
CERN, SL Division, CH-1211 Geneva 23

Abstract

A double sweep streak camera, built by industry according to CERN specifications, has been used for a number of years to provide real time three-dimensional measurements of bunches in LEP, by means of a dedicated synchrotron light source. Originally requiring local manipulation in an underground lab close to the LEP tunnel, the camera can now be fully operated via the control system network. Control functions such as the adjustment of lens and mirror positions, the selection of camera sweep speeds and the setting of 12 ps resolution trigger timing, are handled by various networked VME systems, as is real time image processing. Bunch dimension averages are transferred every few seconds via the control system to the LEP measurement database, and a dedicated high bandwidth video transmission allows the streak camera images and processed results to be viewed in real time (at 25 Hz) in the LEP control room. Feedback control loops for light intensity, trigger timing and image tracking allow the setup to provide useful bunch images and logged measurements over extended periods, without human intervention. An X-Window based control application (GUI) allows LEP machine operators to select different bunches for display and measurement. The same application permits the specialists to control all parameters of the system.

1. INTRODUCTION

Synchrotron light pulses are produced by the passage of e^+ and e^- bunches through small wiggler magnets placed at 67 m on either side of intersection point 1 of LEP [1]. The light is extracted by two thin beryllium mirrors in the vacuum chamber and focused via two evacuated optical lines on a double sweep streak camera [2] in an underground optical laboratory 15 m from the tunnel [3]. The optical setup allows the simultaneous observation of the side and top views of any photon bunch from both LEP beams within the same fast sweep. Measurements can therefore be made of the instabilities and sizes of the particle bunches in all three dimensions [4]. The photon bunch length and longitudinal density distribution corresponds to that of the particle bunch that emitted it. The absolute calibration of the transverse measurements is more difficult, being influenced by the optics of the source and the light collection and transport, and is better done by other instruments that are dedicated to emittance measurements [5]. The slow sweep allows up to 100 fast sweeps to be recorded on one image, which can be used to follow successive bunch passages.

The streak camera system can be decomposed into a number of major blocks requiring control, as shown schematically in Figure 1. First, there is the synchrotron light path that brings light pulses from the LEP tunnel to the entrance of the streak camera. This includes a light intensity feedback system to maintain the light incident on the streak camera at an optimum level, as well as a number of mechanically controlled lenses, mirrors, shutters, and camera rotary switches. This subsystem is described in Section 2. Next, there is the synchronisation between the beam signals, the streak camera itself, its associated CCD camera and the image processing system. This covers the range from milliseconds down to picoseconds and is described in detail in Section 3. Some features of the image processing system, based on the Datacube MaxVideo20 board [6], are described in Section 4. Finally, aspects of the control software, including the X-Window control application being prepared for this entire system, are covered in Section 5.

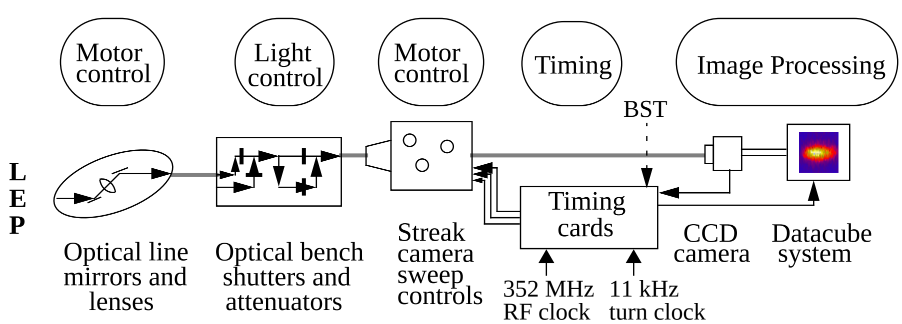

Figure 1 : Principal streak camera subsystems

2. THE SYNCHROTRON LIGHT PATH

2.1 Lenses and mirrors in optical lines

The optical lines that transport the synchrotron light from each beam to the optical laboratory are each equipped with an achromatic lens that can be moved longitudinally over a range of 75 cm in order to bring the light to a focus on the streak camera at a distance of about 20 m. These movements are motorised via stepping motor interfaces standardised for LEP beam instrumentation [7], and resolvers provide measurements of the achromat positions. Also motorised are the inclinations in two orthogonal planes of various mirrors used to centre the light beam in the optical lines, including the beryllium mirrors used to extract the light from the LEP vacuum chamber.

2.2 Synchrotron light intensity control

A feedback loop maintains the synchrotron light intensity received by the streak camera within a limited range. The aim is to have as bright as possible an image on the CCD camera without saturation, thus optimising the signal to noise ratio. The optimal incident light level of course depends on the streak camera sweep settings. Slower sweeps concentrate the bunch image on fewer CCD pixels and require more attenuation. Taking this into account ensures that there is no risk of damaging the streak camera phosphor screen by concentrating too much light on it.

A fixed small proportion (4%) of the incident synchrotron light is deflected towards photomultipliers (PMs), one for each beam. Analogue electronic circuitry, including a peak detector and fast comparator, processes the PM pulses generated by the individual LEP bunches, producing a level that is proportional to the highest pulse height in the last few tens of milliseconds. This level is digitised by an 11 bit ADC every 100 ms; availability of a value triggers an interrupt service routine (ISR). If the measured signal level falls outside the presently required range, then the ISR calculates the change in the attenuation required to re-establish the central value of the range.

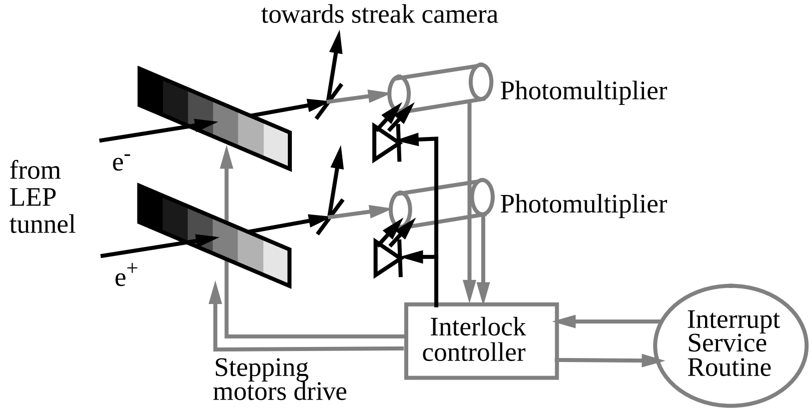

Figure 2 : Synchrotron light intensity control

Optical attenuators, 150 mm long, are mounted on sliders driven by stepping motors at a speed of 40 mm/s, with a step size of 63 µm. The light attenuation obtained increases exponentially with distance up to a maximum value of 1000 (i.e. from 100% to 0.1% of incident light transmitted). The 4 mm movement that can be produced in the 100 ms between interrupts therefore corresponds to a 20% change in attenuation. The ISR derives the number of motor steps required, assuming that the attenuation changes linearly within the 4 mm interval. The obtained value is limited to 64, the maximum number of steps that can be executed in 100 ms, the excess being treated in subsequent intervals. To compensate for variations in the PM gains and ensure that the reference signal level always corresponds to the same incident light level, a calibration procedure is used to adjust the PM bias voltages to obtain the required reference level from a light emitting diode (LED) included in the PM bases.

In addition to the software-controlled light attenuation feedback, which can take up to 4 seconds to reach maximum attenuation, there is a fast analogue detection of preset maximum and minimum light levels. The maximum level is particularly important as it has been determined to ensure that the photocathode of the streak camera is not damaged by very intense light pulses. As the streak camera receives synchrotron light from a region of LEP where an intense laser beam is brought into collision with the particle beam in order to measure polarisation, there is always a possibility that a reflection from this laser is directed towards the streak camera. The analogue detection circuitry generates an interlock that closes a shutter in front of the streak camera when the maximum light level is exceeded for more than 20 ms. In addition, an error condition detected on either PM high voltage supply generates an interlock. However, in the case of normal operation of the polarimeter this level of protection should not be necessary, as there is also a hardware veto signal generated by switching on the laser that produces an interlock. In addition, there is a software call made to the low level streak camera light control software that closes down the streak camera system elegantly when the LEP operators request that the polarimeter laser power supply be turned on. This shutdown consists of closing the main shutter, disabling the light intensity control system and setting the attenuation to maximum. Once the laser has been turned off another call generates a warm start of the system. The light intensity control system is first calibrated and then enabled. Once it is stable the main shutter is reopened. Even once the conditions that caused it have disappeared, the interlock always remains active until reset remotely by software.

A VME module called the Interlock Controller [8] has been developed for measuring the pulses from the PMs and controlling their bias voltages and test LEDs, controlling the attenuators via stepping motors, detecting the high and low light levels, handling the interlock and operating the view shutters described below.

2.3 Particle beam and view shutters on optical bench

Before being combined into one beam on the optical bench, the light emitted by the e^+ and e^- beams passes through electro-mechanical shutters that can be remote-controlled via the Interlock Controller. The combined e^+/e^- light beam is then split into two paths, one containing a Dove prism that rotates the photon bunches by 90° about their longitudinal axis, the other containing a manually variable optical delay that provides a time separation between the rotated and unrotated bunches after recombination in front of the streak camera. On the final image, one light path provides the view of bunches seen from the side, while the other provides the top view. Shutters in the two paths allow the selection of either or both views. The time difference of a few 100 ps introduced by the optical delay and the fixed 500 ps e^+/e^- difference makes it possible to display both top and side views of both beams on the same image. In the case where it is required to change from one view of different bunches of the same beam (such as shown in Figure 7) to another beam/view combination of the same bunches, the control software is able to use the known time separations to adjust the fast trigger delays to keep the bunch images in the same positions on the screen.

2.4 Streak camera sweep controls

At the time of construction of the streak camera it was not considered necessary to have the possibility of controling the sweep speeds remotely. A special mechanical assembly therefore had to be constructed later to allow the 3 main control rotary switches to be turned remotely by stepping motors. These motors are controlled by the same stepping motor interfaces used for the mirror and lenses of the optical lines [7]. The software must keep track of the rotary switch settings as there is no measurement available of the motor positions. There are 2 discrete position rotary switches, one for the fast sweep settings of 0, 90, 150, 200, 300 and 1200 ps/mm (measured values) and the other for the slow sweep settings of 0, 4, 40 and 400 µs/mm. The sweep values correspond to displacements measured on the image incident on the CCD chip, which has a pixel size of 23 µm. In addition there is a continuous slow sweep control rotary switch, which in 10 revolutions varies the slow sweep from any set non-zero value to that 10 times slower (e.g. from 4 to 40 µs/mm). Is it thus possible to obtain any value from 4 µs/mm to 4 ms/mm, corresponding to a time range on the 288 pixels of the CCD in the direction of the slow sweep from 26 µs to 26 ms (horizontal in streak camera but vertical on final image -- see Figure 7). However in practice the largest useful time is limited to about 10 ms, as discussed in Section 3.1. The CCD chip has 384 useful pixels in the direction of the fast sweep resulting in a set of effective screen widths of 0.8, 1.3, 1.75, 2.6 and 10.6 ns. Changes of sweep frequency must be communicated by the control software to the light intensity control system, to the timing control software and to the image processing software.

3. SYNCHRONISATION

Successful operation of the streak camera system clearly requires synchronisation between the different entities involved, i.e. the streak camera itself, the CCD camera that records the image produced on the streak camera phosphor screen after passage through a multi channel plate (MCP) image intensifier and finally the commercial image acquisition and processing system from Datacube that digitizes the CCIR composite video signal from the CCD camera (of frame frequency 50 Hz, line frequency 15.625 kHz, and pixel frequency 7.375 MHz). In addition one has to synchronise with the synchrotron light pulses being received from the two circulating beams in the LEP machine. However it should be noted that light pulses are available from the beam at a high frequency (11.2 kHz for any individual circulating bunch) and thus in general only the streak camera needs to be precisely synchronised with the LEP beams. A special mode of operation, referred to as the “BST mode”, is under development for synchronising bunch visualisation with particular machine events such as the instant of injection. This mode, which also allows the variation of the image frequency and the interval between visualised turns, is described in Section 3.3 below.

3.1 CCD mode

In the normal mode of operation, referred to as “CCD mode”, it is the CCD camera 50 Hz frame clock that controls the synchronisation. As shown in Figure 3, it is the “CCD ready” signal produced by the CCD controller a few ms after the start of each alternate frame that launches the sequence of streak camera sweeps and also triggers the Datacube acquisition. The image from the streak camera is integrated on the image zone of a Thomson TH 7863 frame transfer CCD chip **i**n the milliseconds following the “CCD ready” pulse in one 20 ms frame, and is available to the Datacube in the following frame of the composite video output.

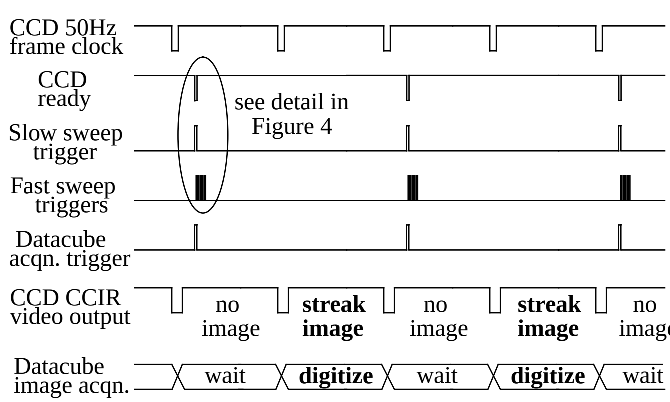

Figure 3 : Synchronisation of streak camera, CCD camera and Datacube system

The integration of the light from the entire burst of bunch images must be complete before the end of the frame and the transfer of the image into the memory zone of the CCD, from which the charges are read out. Taking into account the relatively slow decay time of the P20 phosphor at the exit of the MCP intensifier, the burst cannot cover more than about 100 LEP turns. The Datacube MaxVideo20 module operates in a continuous loop in which once it has completed digitizing one video frame it waits for the arrival of the trigger before starting the digitization of the next complete frame. This triggering scheme thus provides a new image for display and analysis every 40 ms.

The synchronisation of the streak camera with the synchrotron light pulses produced by the e^+ and e^- bunches circulating in LEP is done in various stages. The streak camera is equipped with a fast internal pulse bias of the photocathode that acts as an electro-optical gate. This gate must be triggered once for each LEP bunch passage to be measured. The time the photocathode remains sensitive can be varied between 0.2 and 2 μs, but is normally set to the maximum value for convenience. It is clear that to measure a given bunch on successive revolutions one needs to generate a trigger with a fixed delay with respect to a clock train synchronous with the circulating beam. This 11.25 kHz clock, referred to as the "LEP turn clock", is derived from the LEP RF system and transmitted via optical fibres to the underground laboratory containing the streak camera. It has an rms jitter of about 1 ns. A general purpose delay module (LSD [9]) is used to produce pulses with a programmable delay and width in 50 ns steps with respect to the turn clock.

The bunch structure of the LEP beams in 1995 consisted of 4 equidistant bunch trains each containing 1-4 bunches separated by 87 RF buckets (87 τ_{RF} = 247 ns). One LSD output is used to produce the triggers for the optical gate, allowing the selection of 1-4 bunch trains per turn. As the full bunch train passes within 750 ns, all bunches in a given train are accepted by the optical gate width of 2 μs. It is the fast sweep triggers, generated with a special picosecond timing module (see Section 3.2) that select the particular bunches within a train and also determine their position on the phosphor screen in the direction of the fast sweep. The synchrotron light pulses from the other bunches that fall on the photocathode when the optical gate is open do not fall on the visible part of the phosphor screen**.** The streak camera slow sweep is used to separate the bunch images from successive LEP turns on the phosphor screen. To optimize the linearity, the initial and final 8% of the slow sweep ramp is off the visible part of the phosphor screen and therefore the slow sweep trigger must precede the first optical gate trigger by about 10% of the useful sweep time, a delay that varies with the slow sweep speed selected. A second LSD output is therefore used to generate one slow sweep pulse per turn that can be positioned to include the part of this relative delay corresponding to the fractional part of a LEP turn. In practice this level of refinement is not exploited, although it could be used to ensure that the first bunch image on the screen is optimally positioned.

The specific streak camera "Timing Adaptor" [10], takes as input the trigger pulses already correctly positioned with respect to the LEP turn clock and generates the burst of triggers for each full image. In particular it creates the integer part of the slow sweep to optical gate delay, by delaying the optical gate triggers by a fixed 8 LEP turns and the slow sweep trigger by a variable 0-7 turns. As shown in Figure 4, which gives as an example the case of 4 bunches acquired for 2 consecutive turns, these delays apply from the turn clock pulse received immediately following the "CCD ready". The variable slow sweep delay, as well as the number of LEP turns in the burst, are presently set manually on the front panel of the model, but will be programmable in the final version.

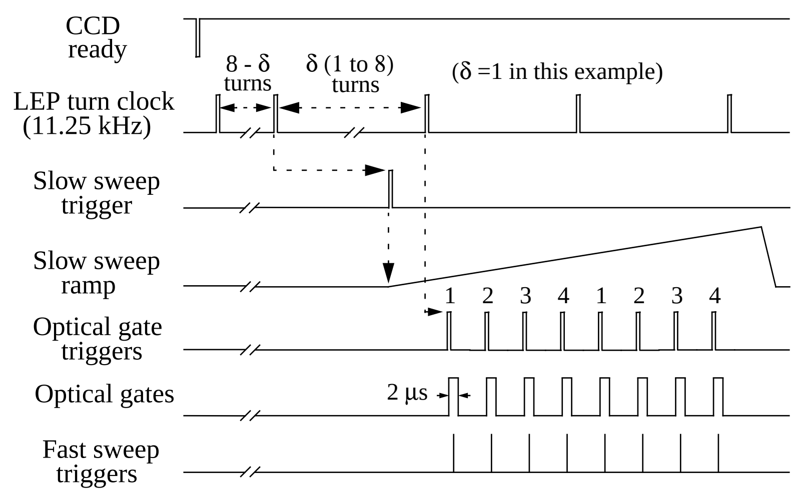

Figure 4 : Example of generation of triggers to measure 4 bunches on 2 consecutive LEP turns

The Timing Adaptor, also applies two constraints on the streak camera triggering necessary for the safe operation of the camera. The first is that the minimum time between triggers be 20 μs. In practice this fits well with the 22 μs between the bunch trains in LEP and simply limits the system to acquiring 1 bunch per train. The second constraint is on the average pulsing rate which should not exceed 200 every 40 ms. This limits one CCD image to containing, for example, 4 bunches per turn for 50 turns.

3.2 Picosecond timing

A picosecond timing module has been developed to provide beam synchronous trigger pulses with a resolution of 12 ps for LEP beam instrumentation. The pulses produced are synchronous with the 352 MHz RF clock train, with a jitter of less than 5 ps, and are positioned with respect to the first RF ticks detected after the arrival of the 11.25 kHz turn clock pulse. The delays of the pulses are specified in two parts, first the number of RF ticks since the previous pulse (or first RF tick in the case of the first pulse in a turn), then an 8 bit fine delay (1 bit = 12 ps) that covers one RF clock period. The module is able to generate up to 32 pulses per LEP turn, but for the streak camera only 4 pulses separated by 21-23 µs are used, corresponding to the selection of 1 bunch out of every train.

These pulses are used to drive the fast sweep of the streak camera via the Timing Adaptor, as discussed in Section 3.1 above. As in the case of the slow sweep, the fast sweep has to be triggered in advance of the arrival of the photon bunch in the streak camera by a time that decreases from 60 to 8 ns as the sweep speed increases. The fact that the pulse for each fast sweep has its own fine delay with respect to the nearest RF clock makes it possible to individually adjust the position of the corresponding bunch image in the direction of the fast sweep.

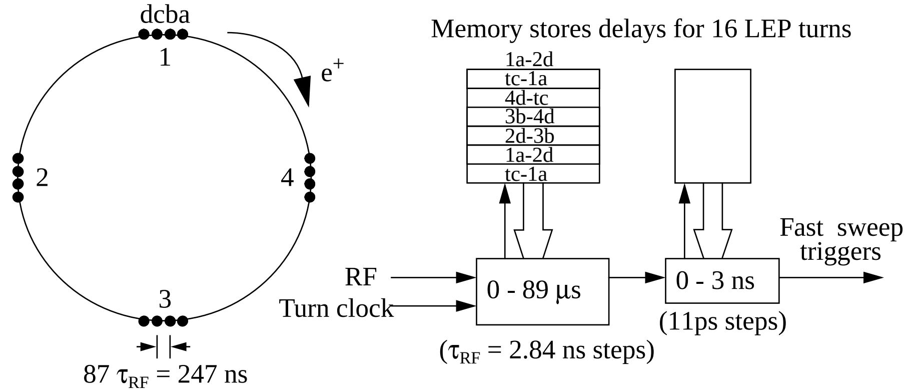

Figure 5 : Picosecond timing for bunch selection

When the streak camera fast sweep is triggered 4 times per LEP turn for a number of consecutive turns (i.e. a burst at 45 kHz), the bunch images of the first 16 or so turns appear displaced on the screen by up to 300 ps (see left part of Figure 6). This effect is reproducible and due to a small distortion of the fast sweep triggering in the first 1 - 2 ms of the burst. Rather than attempting to correct this distortion at the source, the LSD multi-page sequencer module has been programmed to present slightly different delays for the 4 trigger pulses in each of 16 consecutive LEP turns. Applying the same shifts for each trigger in a given turn produces the result shown on the right of Figure 6. An even finer correction has been implemented by having a variable shift for each of the 4 triggers per turn. This correction has only been tried with one fast sweep speed and it is quite possible that different corrections would be required for different sweep speeds. In the present implementation, the required delays are stored in the 16 memory pages of the LSD module and read each time the RF clock delay counter in the picosecond timing module passes through zero. It is also possible to disable the distortion correction by looping indefinitely on one memory page, which is the only way to generate more than 16 turns of triggers. The final version of the picosecond timing card, presently under development, will include similar functionality onboard but with 32 sets of 32 delays per turn.

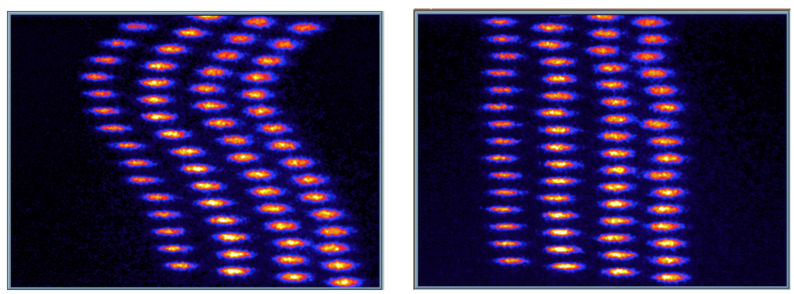

Figure 6 : Fast sweep trigger distortion before and after correction

3.3 BST mode

This alternate mode has been built into the present version of the Timing Adaptor, but has not yet been fully commissioned. The Beam Synchronous Timing system [11] is a dedicated message distribution system in which different beam instruments located at different points on the 27 km LEP ring can receive triggers within the same LEP turn. This means that exactly simultaneous measurements can be produced on different beam instruments. Different bits in the 56 bit message transmitted every LEP turn are allocated to different instruments; with some destined to trigger hardware directly, while others trigger interrupt driven software or contain data values. In the case of the streak camera, one “hardware” bit (which is either set or not during every LEP turn) is used to trigger sequences of measurements via the Timing Adaptor. It is possible to program the BST system to trigger any arbitrary sequence of turns, either continuously or linked to particular machine events, such as injection. For example, one could request images with 10 turns displayed with an interval of 8 turns between each displayed turn. Although the more exotic combinations of turns may not be very useful in practice, the possibility to vary the image frequency from 25Hz down to single shot is certainly an advantage of this mode.

In the BST Mode it is the arrival of the first BST streak camera bit, rather than the CCD ready, than enables the slow sweep trigger. The whole sequence of BST bits is then delayed 8 turns and serves to enable the appropriate sequence of optical gates and fast sweep triggers on the desired turns. The streak camera is then being driven by the BST, but asynchronously with the CCD camera, which would receive images from the camera phosphor screen at any moment in the CCD cycle. While this would mostly work for short bursts of turns (e.g. 11 turns = 1 ms), the results would not be reproducible, with some images lost and some with uneven light integration. To solve this problem, the manufacturer of the CCD camera was asked to provide a “CCD init” input that allows the restarting of the CCD frame clock at any moment. The first BST bit should then produce a “CCD init” pulse and then after an appropriate time start the sequence of streak camera triggers. Another problem with this mode of operation occurs in the Datacube image acquisition. When the CCD camera is free running, the phase of the video line frequency is maintained from one frame to the next, which allows the Datacube to use a PLL on the line synchronisation pulses to generate the pixel clock for digitization. However, the use of the “CCD init” breaks this phase (unless the “CCD init” is pulsed at an exact multiple of the line period) and the PLL becomes unusable. The Datacube permits synchronisation with an external pixel clock, but this should normally be continuous, i.e. even when there are no active pixels to digitize during blanking periods. However with some modification to the pixel clock of the CCD camera to add some extra pixel ticks between lines and appropriate programming of the Datacube input video parameters, it has been possible to drive the Datacube image acquisition reliably with the CCD pixel clock.

4. IMAGE PROCESSING

The MaxVideo20 pipeline image acquisition, processing and display card from Datacube forms the heart of the streak camera image processing system. Processing elements on the card can be reconfigured by software to form a data pipeline through which image data can then “flow”. In the present implementation the MaxVideo20 card is only used to (i) digitize the CCD camera image in CCIR standard video mode with an 8 bit flash ADC and (ii) generate the final high resolution video display for the LEP control room. All the image analysis is done on a 68030 processor card running Microware’s OS-9 operating system.

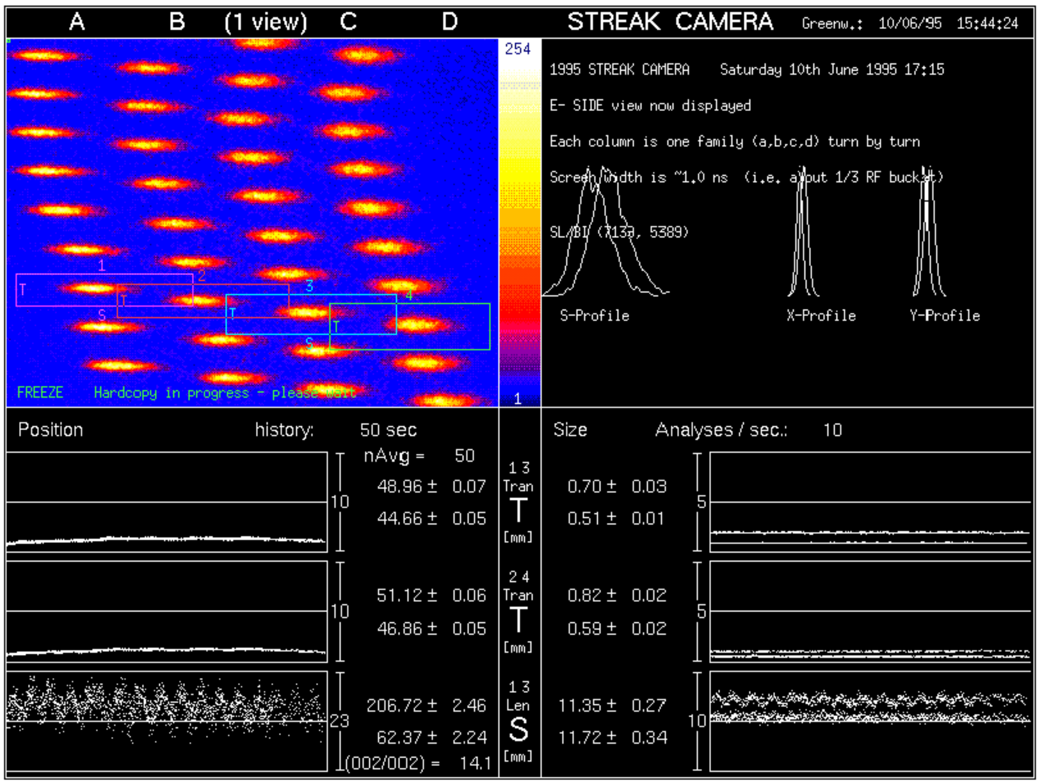

Figure 7 : High resolution video display transmitted to LEP control room

The image digitized from the CCD camera video output is refreshed at full speed (25 Hz in the case of normal “CCD mode” timing) in the top left part of the display. The intensity recorded in each CCD pixel is displayed in false colours, as indicated on the colour scale. The fast sweep appears to run from left to right and the slow sweep from top to bottom. In this example the side view of bunches 1a, 2b, 3c, and 4d of the e^- beam for 9 consecutive LEP turns has been selected. The width of the image corresponds to 1.3 ns, the height to about 0.8 ms. It can be seen that the distortion correction described in Section 3.2 has not yet been implemented.

The evaluation of bunch position and size is done on a copy of the image made by the MaxVideo20 card in 5 ms. This copy is only updated when the calculations on it are complete. When the screen copy indicated in Figure 7 was made, 4 bunches were being analysed at a frequency of 10 Hz. The rate is higher when the profiles shown in the top right part of the screen are not displayed. The 4 rectangles superimposed on the CCD image indicate the screen regions being used for the determination of the bunch positions and sizes. All rectangles are producing transverse measurements, while only rectangles 1 and 3 are measuring longitudinally. This was suitable for simultaneous measurements of the side and top views of the same bunch. For the measurement of bunch train families, all rectangles measure longitudinally and only 1 and 3 measure transversely. The bunch size calculation algorithm is quite complex but essentially consists of: (i) making projections of the bunch density distribution on the sides of the rectangle, (ii) estimating the background from the densities at the edges of the rectangle, and (iii) evaluating the bunch sizes from the standard deviations (σ) of the projections after background subtraction. The sides of each rectangle are at somewhat more than 4 σ from the centre of gravity of the enclosed bunch. The result of each evaluation of bunch position and size is converted to mm with conversion factors that depend on the sweep speeds and plotted on the corresponding continuously scrolling history graph. The time period shown in the history plot can be varied in steps of 10 s from 10 s to 24 h (50 s is selected in the figure). The individual values are also combined to produce an average and standard deviation for each measured quantity, and this is displayed numerically. The number of consecutive measurements used for this purpose can be varied from 17 to 3000 (50 is used in the figure). The average and standard deviation of each position and size is also available for reading out by a remote logging process (see Section 5.2). In the case of fewer than 4 columns of images on the screen it is possible to reduce the number of rectangles being used.

An important feature of the image processing is the continuous tracking of the bunch image by its enclosing rectangle. Each calculation of bunch position and size defines a new position and size for the rectangle. However in order that noise on the measurements does not generate oscillations in the rectangles, the calculated changes are strongly damped by limiting them to a one pixel change in position or size every analysis period. Nevertheless, this is sufficient for the rectangles to respond to real changes in the positions and sizes of the bunches. In the case of an abrupt change in the image, caused for example by a change in sweep speed, a given rectangle may lose the bunch it has been tracking. A complete screen bunch search algorithm enables it to locate the nearest one and lock onto it. This search may also be triggered at any time for all the rectangles by a software request to the system.

During acceleration in LEP the particle bunches change position in the RF bucket by a few hundred ps. This may result in some bunch images moving off the screen. An automatic procedure is available to keep the bunch in the first rectangle within a vertical band on the left side of the screen. When the bunch leaves this band in either direction a fine delay is changed that acts to bring it back. This fine delay is obtained by programming an 8 bit DAC on the MaxVideo20 card to add a DC level to the fast sweep trigger generated by the picosecond timing module, before sending it to the Timing Adaptor. The slope of the leading edge of the pulse is adjusted such that the pulse triggers about 1 ns earlier for full scale on the DAC. This automatic alignment procedure also compensates for any other effect that may change the bunch position on the screen, such as a slow drift in the 352 MHz RF clock distribution, or thermal effects in the streak camera or its associated electronics, as long as the effect is no larger than a few 100 ps.

The 800 x 600 pixel display produced by the MaxVideo20 at an interlaced frame frequency of 60 Hz requires a transmission bandwidth of 17 MHz, considerably in excess of the 5.5 - 7 MHz provided by the standard PAL transmission system. The separate analogue video signals for red, green, and blue (RGB), with synchronisation on the green channel, are therefore transmitted via 3 multimode optical fibre links over the 4 km to the LEP control room. Propagation delays on the 3 links are adjusted to be identical and industrial optoelectronic transmitters and receivers provide a bandwidth of 25 MHz.

5. CONTROL SOFTWARE

The streak camera system control architecture [12] is shown schematically in Figure 8, where the row of bubbles in the lower part of the figure represents the front-end control software resident in various VME computers, each controlling a part of the streak camera system described in the previous sections.

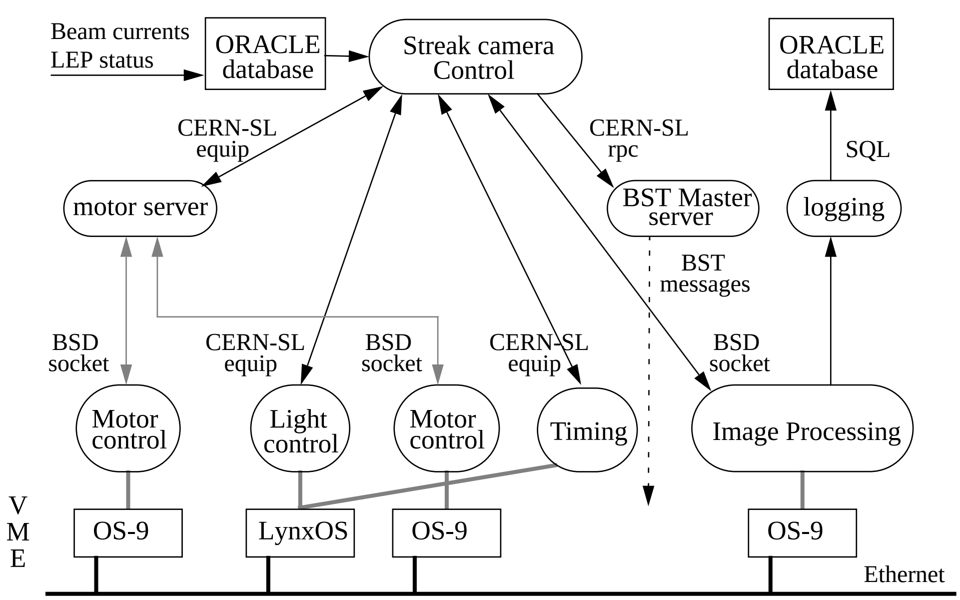

Figure 8 : Streak camera control architecture

These computers are diskless systems with 68030 processors directly connected to Ethernet, and their applications are loaded from a central file server via NFS (Network File System). The motor control software runs under Microware's OS-9 operating system and controls standard stepping motor interfaces. The light control and timing software runs in a single Lynx OS system, while the Datacube image processing system uses OS-9. For communication the "equip" package developed in the SL Division of CERN [13] has been used where possible. This is a layered data exchange protocol built on the CERN remote procedure call (rpc), which in turn uses a transport protocol handling datagrams. An "equip" server was developed for controlling the various VME modules of the LynxOS system, where it was installed as a thread sharing data with the acquisition thread that is used to run the light intensity control system. The motors were controlled via an "equip" server running on an HP-UX node that handles accesses to all such motor systems via TCP-IP using BSD sockets. TCP-IP is also used directly to communicate with the image processing system. In this case shared memory segments (MOPS [14]) visible to both the receiver process and the image analysis process are transferred, and a circular buffer is used to handle successive calls.

5.1 X-Window application

A Motif style graphical user interface (GUI) for control of the streak camera system is under development using the X-Window User Interface Management System (XUIMS) supported by the CERN SL Division Control Group [15]. This interface will replace the set of three presently used control applications, all of which have simple text interfaces. The main panel, as shown in the upper part of Figure 9, presents the current state of the view selection shutters, possible error conditions from the synchrotron light intensity control subsystem and relevant LEP machine parameters, in particular mini-wiggler magnet and beam currents which are retrieved from an on-line database. From this information one can determine at a glance whether there is a simple reason for the absence of images on the streak camera video display, e.g. no currents in the mini-wigglers used as the synchrotron light sources. Space has been reserved on this panel for a graphical representation of the evolution of the measured bunch lengths using the commercial plotting package XRT/Graph. This will allow monitoring of the correct functioning of the image analysis, even from places where the streak camera video display is not visible. If both top and side views are being measured, then the bunch lengths given will be corrected for the frontal spot size of the synchrotron light pulses [16].

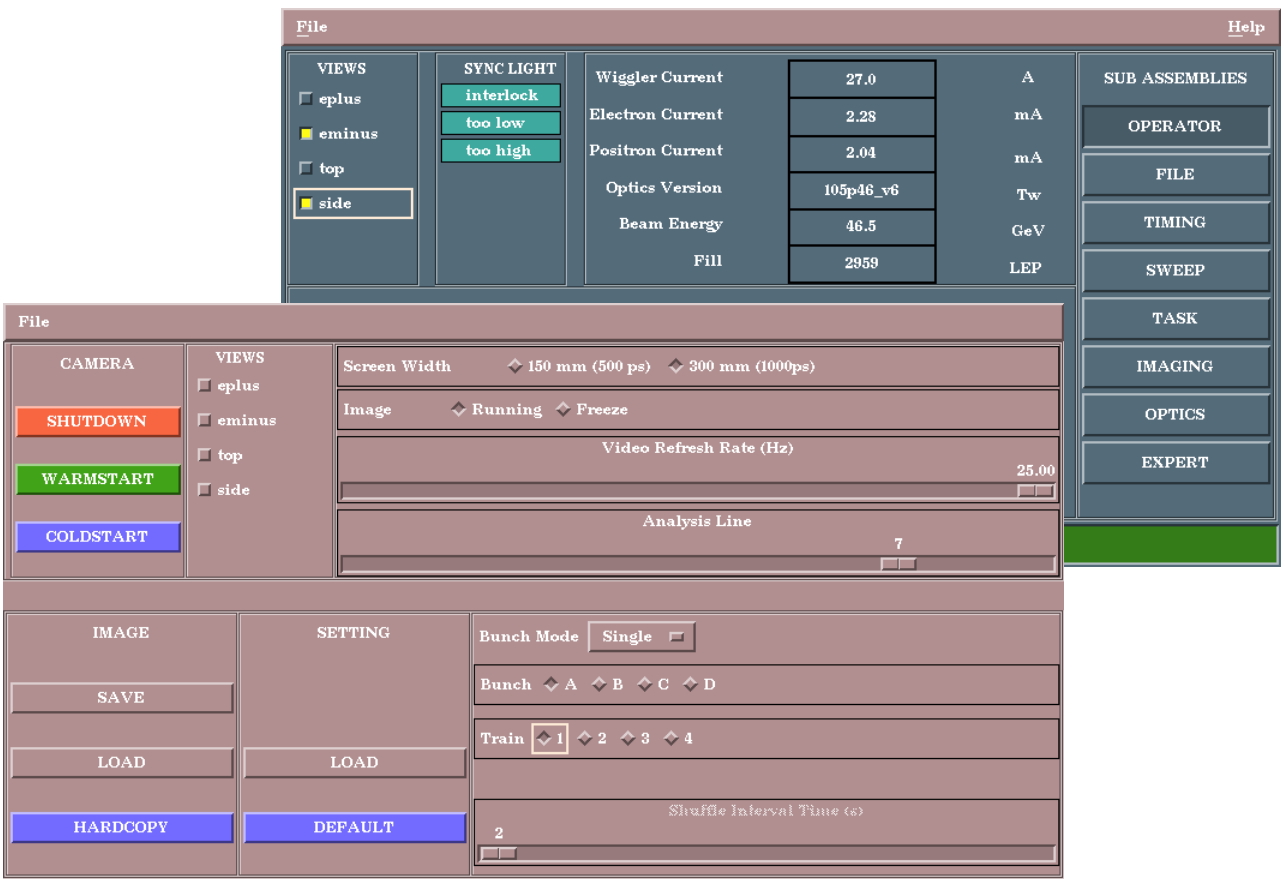

Figure 9 : Streak camera control main User Interface panel and sub-panel for LEP operators

Subsidiary panels group together all the controllable parameters according to type, e.g. timing, streak camera sweeps, motorised optical elements, etc. Control panels that are destined for use only by experts (e.g. adjustment of optical line elements) are password protected. A special panel is being prepared containing the limited set of functions considered to be necessary for the use of the streak camera by LEP machine operators, a first version of which is shown in the lower part of Figure 9. Starting in the top left corner of this panel, the first section is devoted to global operations on the system including “Shutdown”, that ensures that no light reaches the streak camera; “Warmstart”, that puts the camera once again into operation and “Coldstart”, that allows the rebooting of the various low-level VME control computers. Then a “Views” section permits the opening and closing of the e^+, e^-, top and side shutters on the optical bench. The next section allows one to choose from the two most useful fast sweep settings, to freeze/unfreeze the CCD image display, to change the image refresh frequency (once the BST mode is commissioned) and to set the analysis rectangles on a particular displayed turn. In the bottom left section of the panel there are controls to save and restore the CCD image area of the video display to and from disk files and to make a colour or monochrome Postscript file or printed copy of the entire video display (such as that shown in Figure 7). Then there is a “Setting” section to allow the loading of predefined configurations (including the most useful beam, view, bunch and sweep combinations). Finally there a section for changing the bunch selections, but the operator would normally use the preset configurations proposed in the “Setting” section.

One important task of the top-level control application is to distribute relevant changes in the parameters of one sub-system to another. For example a change of either the slow or fast sweep (*motor control*) implies changes in trigger delays (*timing*), light attenuation (*light control*) and bunch dimension pixel to mm conversion factors (*image processing*). Changes to the identity and type of bunch images selected via the beam and view shutters and fast sweep timing must also be communicated to the image processing system (to keep the identification on the video display correct) and the logging process (to ensure that the correct data is updated in the database).

The introduction of the bunch train scheme and the interest in observing simultaneously the different bunch “families” (i.e. the different members of trains, labelled from “a” to “d”) has considerably increased the number of beam/view/family combinations possible. Previously, only one bunch per turn was selected and both beams and views were displayed across the screen in the direction of the fast sweep. This produced 4 columns corresponding to e^+ top, e^+ side, e^- top, and e^- side. For bunch train operation in 1995 the default display consisted of 4 columns corresponding to one bunch family per train (usually 1a, 2b, 3c, 4d, as in Figure 7) corresponding to one beam and view. However combinations of e^+, e^-, top and side remain of interest, particularly with the likely use of only 2 bunches per train for LEP2 operation from November 1996, which reduces the number of beam/view/family combinations to 6 from the present 32 for 4 bunches per train.

5.2 Logging on measurement database

An Oracle database is used for storing the history of a large number of LEP beam and machine parameters. A logging process running on an HP-UX node reads the current particle beam, view and bunch selections and the results from the image analysis every 10 s. Embedded SQL calls then update the appropriate fields in tables (one for each beam) containing the last averages and standard deviations of position and size in all 3 dimensions for all possible 32 bunches (i.e. 384 values in all). This triggers another process which copies only the data which has been updated, tagged with a bunch identification, into the history tables. In this way only the measurements being generated by the streak camera system at any time (i.e. maximum 24 values, see Section 4) occupy space in the history tables, and data corresponding to a particular bunch or bunch family may be easily extracted for off-line analysis.

5.3 Configuration data handling

A system for handling configuration data stored in a single data file and required in different processes has been developed [17]. The identification of data items was inspired by the handling of resources in the X-Window System and consists of an ordered pair of alphanumeric tags. Routines are provided to retrieve, update and print the data. The data is stored in an ASCII text file and any comments added by a text editor are retained by the data update routine. As well as strings and decimal values, hexadecimal and boolean values can be included as required. Global variables of scalar type, structures and arrays can be transferred between the resource file and the program. In applications where data have to be stored in shared or dynamically allocated memory, data offsets are calculated during compilation and the absolute location of the start of the storage area is passed as an argument to the data transfer function at run time.

6. CONCLUSION

A streak camera is a rather specialised and delicate device which one is not normally required to control remotely. Nevertheless, a project was undertaken to make the LEP streak camera and its associated optics, timing and image processing available for use by operators in the LEP control room. A secondary goal was that the camera should provide regular measurements of bunch dimensions, in particular bunch length, for logging on a database. Several VME-based computers now provide all the necessary low-level functionality and the interface to the LEP control system. However, despite the successful adaptation of the timing, image processing and logging to handle the new bunch train scheme in LEP, serious hardware problems with the streak camera itself prevented the achievement of these goals in 1995. With the return to CERN of a renovated camera in spring 1996 and the completion of the X-Window control application, the 1996 LEP run should see the goal of a fully operational streak camera become a reality.

7. Acknowledgements

The project to convert the streak camera from the initial locally-operated experimental tool into a fully remote-controlled system usable from the LEP control room has covered a number of years. Although the authors of the present paper are now responsible for completing this project, many others have also contributed. Special mention must be made of G. Baribaud, who led the project in 1992-94 and contributed particularly to the light control, timing and video transmission [18] [19]. Two visitors, first Y. Solberg and then M. Werner, developed the image acquisition and processing software on the Datacube system. E. Rossa and a number of visitors have been responsible for the initial streak camera installation, optics and improvements to the original PC-based analysis software and have provided input throughout the project. E. Rossa has also supervised the design and implementation of the picosecond timing module by a visitor, P. Joudrier, and proposed its use for correcting the fast sweep trigger distortion. C. Beugnet provided the stepping motor hardware and G. Morpurgo developed the OS-9 motor control software, the HP-UX motor “equip” server and a library for programming the LSD module. P. Pivot contributed to the testing and development of the light control system. Finally, C. Bovet must be thanked for his continual support and encouragement of the project, as well as the modification of the synchrotron light source to accommodate the bunch train optics, in collaboration with M. Placidi.

8. References

[1] C. Bovet, M. Placidi, “A dedicated synchrotron radiation source for LEP diagnostics”, LEP note 532, April 1985.

[2] Manufactured by A.R.P, route de Hausbergen, 67309 Schiltigheim, France.

[3] E. Rossa et al., “Double sweep streak camera for LEP”, Proceedings of the 2nd European Particle Accelerator Conference, Nice, June 1990, pp. 783-785.

[4] E. Rossa et al., “Real time measurement of bunch instabilities in LEP in three dimensions using a streak camera”, Proceedings of the 3rd European Particle Accelerator Conference, Berlin, June 1992, pp. 144-146.

[5] R. Jung, “Precision emittance measurements in LEP with imaging telescopes, comparison with wire scanner and X-ray detector measurements”, Proceedings of the International Workshop on Particle Dynamics in Accelerators “Emittance in Circular Accelerators”, Tsukuba, Japan, November 1994.

[6] Datacube Inc., 4 Dearborn Road, Peabody, MA 01960, USA.

[7] C. Beugnet, “Description du module Stepping Motor Interface”, CERN LEP/BI-MP/Note 89-8, November 1989.

[8] J.J. Savioz et al., “Le module Interlock Controller”, documentation technique, CERN SL/Note 93-114 (BI), December 1993.

[9] J.J. Savioz, “Le module LEP Synchronous Delay”, documentation technique, CERN SL/Note 95-03 (BI), January 1995.

[10] E. Rossa et al., “BE 3D”, CERN SL/Note 91-01 (BI), February 1991.

[11] G. Baribaud et al., “The Beam Synchronous Timing Master”, CERN/AT-CO/91-10 & CERN SL/91-28 (BI), July 1991.

[12] J. de Vries, “Streak camera software: present status and future strategy”, CERN SL/Note 93-99 (BI), October 1993.

[13] P. Charrue, “Accessing equipment in the SPS and LEP controls infrastructure - The SL-EQUIP package”, Software User Manual, CERN SL/Note 93-86 (CO), September 1993.

[14] W. Herr, “M.O.P.S. User Guide for C programs”, CERN SPS/88-43 (AMS), January 1993.

[15] M. Vanden Eynden, “The X-Window User Interface Management System (XUIMS) at CERN”, presented at this conference.

[16] F. Tecker, “Evaluation of the performance of a streak camera in beam diagnostics at LEP”, diploma work PITHA 93/7, Physichalische Institute RWTH Aachen, Germany, November 1993, & CERN SL/Note 95-66 (BI), June 1995.

[17] J. de Vries, “Use of resource files for multiple clients sharing a single data file”, in preparation.

[18] G. Baribaud et al., “Specifications for the operation of the first version of the streak camera from the main control room”, CERN SL/Note 92-29 (BI), May 1992.

[19] G. Baribaud et al., “Remote control of the operational streak camera -The second phase”, CERN SL/Note 93-80 (BI), August 1993.