

Atlas-DAQ-Note 33

RD27 Note 39

16 December 1994

FIRST LEVEL CENTRAL TRIGGER PROCESSOR FOR LHC EXPERIMENTS

C. Jacobs, J.P. Vanuxem, H. Wendler

O. Buyanov¹

(ECP Division, CERN, Geneva, Switzerland)

¹ On leave from IHEP, Protvino, Russia

1. Introduction

The Central Trigger Processor acts on Subtrigger Data Patterns (SDP) which are generated by the calorimeter and muon trigger processors. Each bit of information contained in an SDP pattern typically indicates the presence in the detector of at least a given number of electrons, muons or jets having a transverse energy higher than a given value, or a missing transverse energy greater than a certain threshold, etc.

The Level-1 Central Trigger Processor (CTP) forms the final Level-1 trigger decision from the SDP's for each beam crossing, and may be seen as a fully synchronous, 40 MHz pipeline processor. As the SDP's from different trigger processors are available at different times due to specific subdetector and trigger processor delays (latencies), the first task of the CTP is to equalise these latencies in such a way that all subtrigger data belonging to the same bunch crossing arrive in phase before a decision can take place. The second task of the CTP is, of course, the decision making process itself, which determines which of the many different combinations of SDP's should be retained to form the final Level-1 trigger decision. The algorithm applied is a combinatorial expression of coincidence, veto or "don't care" of SDP inputs (bits). Several of these expressions may be required simultaneously (OR). An additional feature of the CTP is to allow for filtering of high-rate and/or less interesting (but partially retained) combinations; this can be implemented by prescaling. A deadtime circuit is included which reduces instantaneous high peak rates to a max. tolerable mean rate. For reasons of testability and flexibility, it should be possible to monitor all parts of the CTP and also to count the frequency of occurrence of various signals (input patterns, internal combinations, final output decision); the latter measurement can be done by means of scalars.

All these features are included in the CTP architecture (Figure 1). At a later stage, they will be implemented with the following concerns:

- The CTP should introduce only a minimum of latency.
- The CTP should be fully programmable and allow for parameter entry, monitoring and testing.
- The CTP should be modular and permit some level of scalability. It should be possible to operate several CTP "slices" (with e.g. 32 SDP bits and 32 trigger combinations per "slice" = module) in parallel in order to handle any reasonable number of SDP's, and also any reasonable number of possible trigger combinations (extension of trigger combinations via additional modules) derived from them.

We initially plan to design the CTP with a top-down approach, using the VHDL language for modelling and simulation at several levels: behavioural level, system level (architecture) and interface level (components). The merit of this approach is to assist us with the definition of the project, in the sharing of the work to be done at various stages, and in the final definition of components as FPGA's (Field Programmable Gate Arrays) and/or semi-custom ASIC's (Application Specific

Integrated Circuits). Also, the VHDL description of the CTP could later be integrated in a more general modelling of the experimental environment and be used for the description and simulation of the complete trigger system.

2. General Description of the CTP

Initial thinking on the Central Trigger Processor is described in more detail in another RD27 note². The block diagram of Figure 1 shows the currently proposed architecture for the CTP.

2.1 Synchronisation Circuits

Subdetector Data Patterns (SDP's) arriving from various trigger processors are for the moment assumed to be either 4-bits or multiples of 4-bits wide. The SDP's belonging to the same bunch crossing arrive at the level of the CTP at different times depending on a number of parameters such as the time-of-flight of the particles, the response time of the detectors, various electronics and cable delays, and the intrinsic latencies of various trigger processors. Therefore, these data patterns must be put in phase before any decision can take place. We plan to implement this synchronisation by means of two circuits.

The first circuit is a "Phase Adjust" (PA) circuit which will tune the timing of the internal CTP clock (simply called "BC" in the diagram as it is related to the Bunch-crossing Clock) to each individual SDP arrival time, such that the SDP data can be safely strobed into a register. It is currently felt that programmable steps of the order of 2.5 ns should be adequate to perform this first level of synchronisation and that a maximum range of a little more than a bunch crossing period, for example 30 ns, would suffice. Available commercial components for this circuit may well prove to be adequate. The design of a similar circuit (higher requirements and more complex) is under progress by the Microelectronics group of CERN-ECP division as part of the "Timing Receiver" chip being developed by them for RD12.

The second circuit is a Variable Length Pipeline register (VLP) which can be thought of as a FIFO with a variable, programmable depth. This circuit is required to add to the arrival time of the SDP a whole number of bunch crossing periods. A suggested maximum length for this pipeline is 24, allowing compensation of latency differences of up to 600 ns between different SDP patterns. Due to the large number of VLP's required, an integrated solution is desirable, and an FPGA implementation of this circuit is being developed.

² RD27 note 9: Central Trigger Processor for the First-Level Trigger for LHC Experiments, June 1993




Figure 1: General Block Diagram of the Central Trigger Processor

2.2 Decision Logic

After having passed through the synchronisation circuits, all subtrigger data patterns (SDP's) are in phase. They can then be strobed at an appropriate time into the alignment register (AR) which is N bits wide, and be presented to the decision logic itself, which is simply represented in the diagram by a Look-Up-Table (LUT). The output of the LUT is a pattern of M bits, where each bit indicates that a predetermined combination (coincidence, veto or "don't care") of SDP input signals have been recognised, and that a candidate for a Level-1 trigger signal is present. These M trigger combinations are first masked with an equivalent number of individual veto signals supplied by external control logic before being stored in the trigger register (TR). All outputs of the trigger register are ORed together, and the resulting signal may itself be ANDed with a global veto signal. The output of the AND is finally synchronised by the Level-1-Accept flip-flop.

It seems possible at this stage to implement all elements of the decision logic with commercially available parts such as RAM's, registers, gates, etc. However, a more integrated design (e.g., using FPGA's) might still be interesting from the point of view of reliability, speed and space. The realisation of the decision logic with LUT's presents technological problems considering the large number of SDP bits (e.g., $N = 128$) and the large number of trigger combinations (e.g., $M = 96$) proposed. The large address space needs to be partitioned in order to fit available parts and modules (see also 2.6).

2.3 Prescalers and Scalers

Programmable prescalers are necessary to reduce trigger rates by filtering high occurrences of specific trigger combinations, thus allowing only a fraction of them to generate a Level-1 trigger.

The 40 MHz clock rate required for the prescalers and the density of interconnections will probably direct us to an integrated solution. At this point in time, an FPGA implementation is being studied.

Scalers are an important ingredient of the CTP as they allow for real-time monitoring of the rates of various internal trigger patterns, which will probably prove to be essential for fine tuning of the Level-1 trigger system. The block diagram indicates that they could be used throughout all stages of the CTP. They are in principle read out and cleared once every LHC turn. Subsequent logic combines (in the monitoring system: either by software or hardware) the scaler data, maintaining totals integrated over longer periods of time.

Many scalers are foreseen throughout the CTP at various levels. The high number of scalers suggested and the requirement to read them out after each LHC turn would make an array of scalers (L scalers per package) an ideal candidate for the development of a gate array circuit (semi-custom ASIC).

2.3 Deadtime

A deadtime circuit is included in the architecture of the CTP in order to reduce instantaneous high peak rates and provide a max. tolerable mean rate (the parameters are not yet fully defined). The deadtime circuit should provide for the following (list not fully exhaustive yet):

- *Blanking period*: after having generated a Level-1-Accept, the CTP should introduce a minimum deadtime of L bunch crossing periods before considering generating another Level-1-Accept; dependent on detector parameters; L can be programmed,
- *Frequency reduction*: to allow only K Level-1-Accepts during a programmable time window I (e.g.: 10 μ sec); K and I can be programmed,
- *Predicting level-2 pipeline "full"*: the Level-2 trigger processor could transmit to the CTP the number of empty buffer places (BUF-FREE) left in its input pipeline; BUF-FREE can be used by the CTP to calculate the number of possible Level-1-Accepts until the Level-2 pipeline becomes full (pre-emptive blocking of further Level-1-Accepts); BUF-FREE must be constantly updated (each burst clock cycle) from the Level-2 processor; this method will avoid Level-2 pipeline overflows.

All the above-mentioned possibilities are presently being studied; the implementation of all or only some of them will be decided at a later date.

2.5 Miscellaneous Logic

Apart from the above-mentioned circuits, additional logic ("glue logic") will be required in the CTP such as registers, gates, flip-flops, etc. Although not explicitly shown on the diagram, various links have to be included for the control, monitoring and readout of the CTP. Suitable receivers and drivers will be needed for fast and reliable signal transmission with the trigger processors, the Level-2 trigger processor and the overall Timing, Trigger and Control distribution system (TTC system)³. Last but not least, as the CTP will be a vital component in the complete trigger system, reliability will be of major concern; testability aspects will also require extensive studies.

2.6 Scalability and Modularity

The CTP will have to comply with a certain number of SDP inputs (N) and a certain number of trigger combinations (M). As until now these numbers are not fully defined yet, some level of scalability has to be included in the architecture of the CTP.

At present it is planned to group 32 bit SDP's resulting with 32 trigger combinations into a SDP-slice (Figure 2), which could be implemented in a 9U⁴ high module. Likewise, additional 32 trigger combinations of the same 32 bit SDP's, could be grouped into a TC-slice (Figure 3) -- this could also be a 9U high module. The whole CTP structure in scalable form (Figure 4) combines several SDP-slices and TC-slices, and the final decision part. The final decision part -- conjunction of horizontal SDP-slices, prescaling, final OR gate, internal GEN_VETO, deadtime circuit and final synchronisation -- combines all SDP-slices and TC-slices in order to provide the Level-1-Accept signal. Mechanically this could also be a 9U high module.

³ Timing, Trigger and Control Distribution for LHC Detectors, B. G. Taylor: CERN/ECP, Rev 1.7, 4 December 1994.

⁴ e.g.: 9U high and 400 mm deep (9U VME or Fastbus).

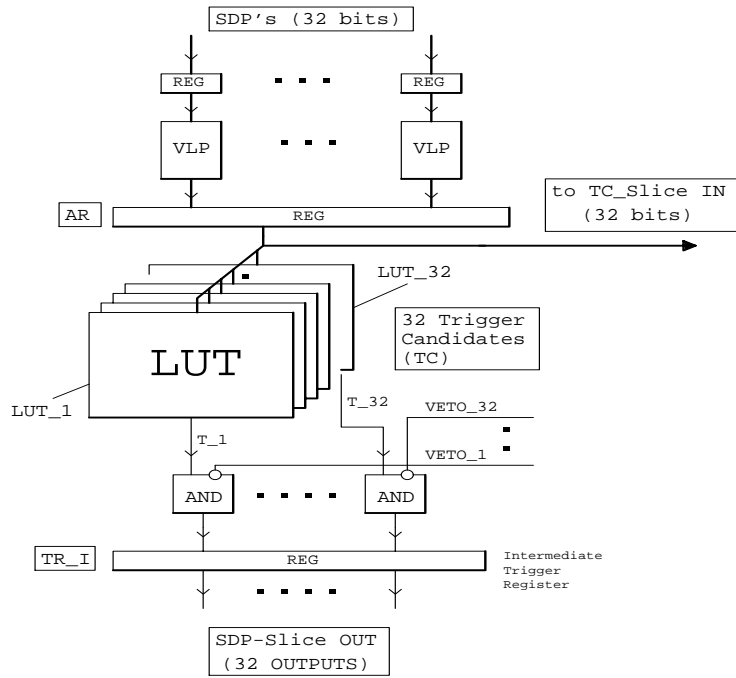


Figure 2: SDP-Slice

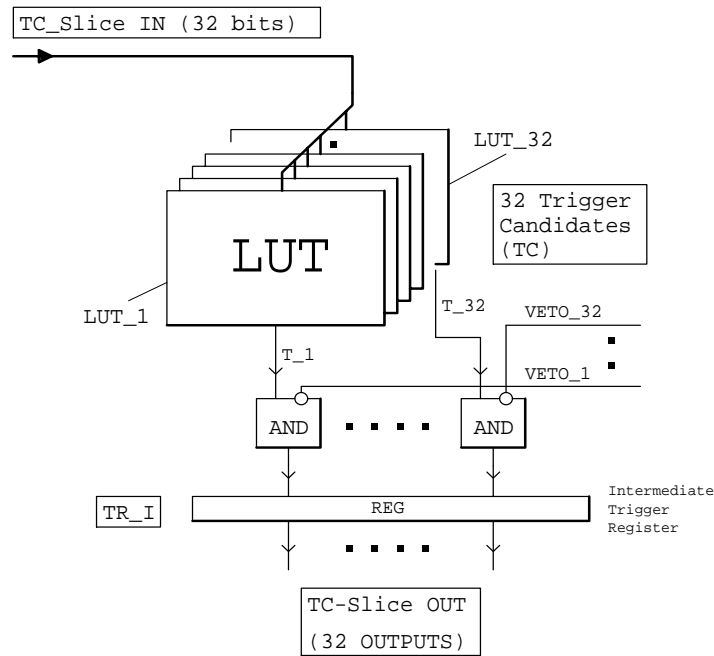


Figure 3: TC-Slice

A possible practical implementation of 128 SDP inputs and 96 trigger combinations is shown in Figure 4. It consists of four SDP-slices expanded via two TC-slices for each of the SDP-slices (32 SDP inputs). The final decision part consists of three sets of 32 AND gates (4-inputs), prescalers, trigger register, final OR gate, veto, deadtime and final synchronisation.

The CTP in scalable form will have to be constructed with several modules as mentioned above. Interconnections between SDP-slices, TC-slices and the final decision part are numerous. The fact that signals will have to propagate from one module to the next will obviously introduce delays. Therefore, an additional synchronisation stage (TR_I) at each SDP-slice and each TC-slice will be needed in order to preserve signal integrity, increasing the intrinsic latency of the CTP by one additional bunch crossing period. Extensive studies are needed in order to determine the feasibility of the above proposed packing style in respect to resources of interconnections needed, density of physical components on each module and the possibility to house the CTP in a single crate.

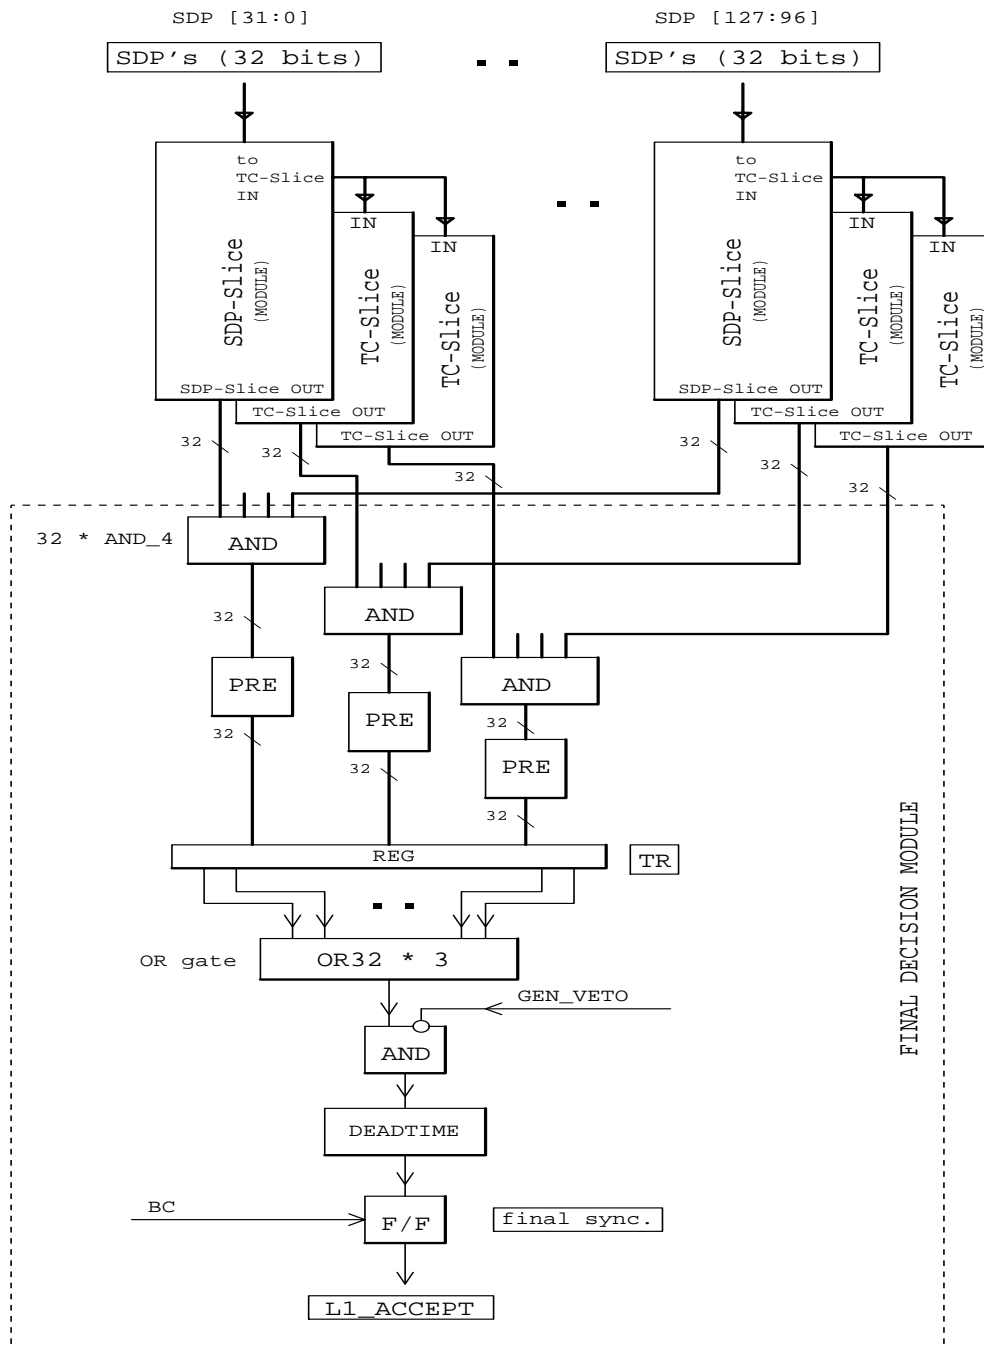


Figure 4: Possible Implementation of the CTP Structure (scalable form)

3. VHDL Description (Specification)

As mentioned above, a VHDL description (model) for the data path of the CTP has been developed⁵. For the simulation of the full VHDL model a minimum part of the control circuitry had to be added to the description in order to provide for the loading of parameters such as values of the prescalers, etc. Placed at the top of the design hierarchy, a VHDL test-bench featuring full testing and monitoring capabilities has been developed in order to permit the simulation of the data path.

3.1 Concept of the VHDL model

The VHDL model has been developed using a top-down structured design strategy and providing scalability. The design is expandable in order to provide for different applications: it consists in fact of three models representing, in increasing degree of complexity, the description of the CTP:

- Behavioural model: System Level Model (SLM); it is based on the decomposition of the CTP space into classes of functions and shareable procedures (Figure 5). It is free of technological constraints.

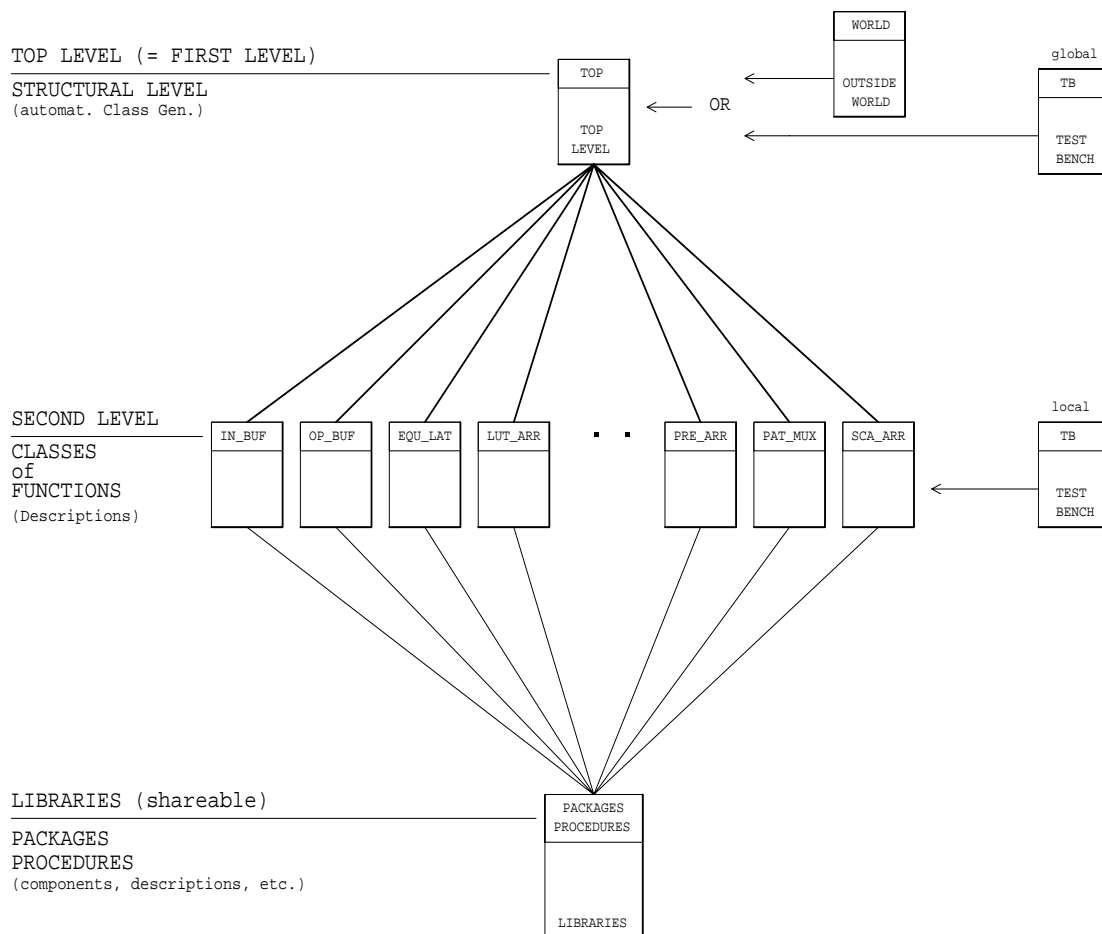


Figure 5: System Level Model (SLM)

⁵ C. Jacobs, CERN/ECP, EDE-Note 01-94, Design Methodology for the Modeling of the CTP using VHDL.

- Intermediate model: Interface Level Model (ILM); it describes the architecture of the system by using physical component entities as effectors (via a set of procedures) of class functions and their connections (Figure 6).

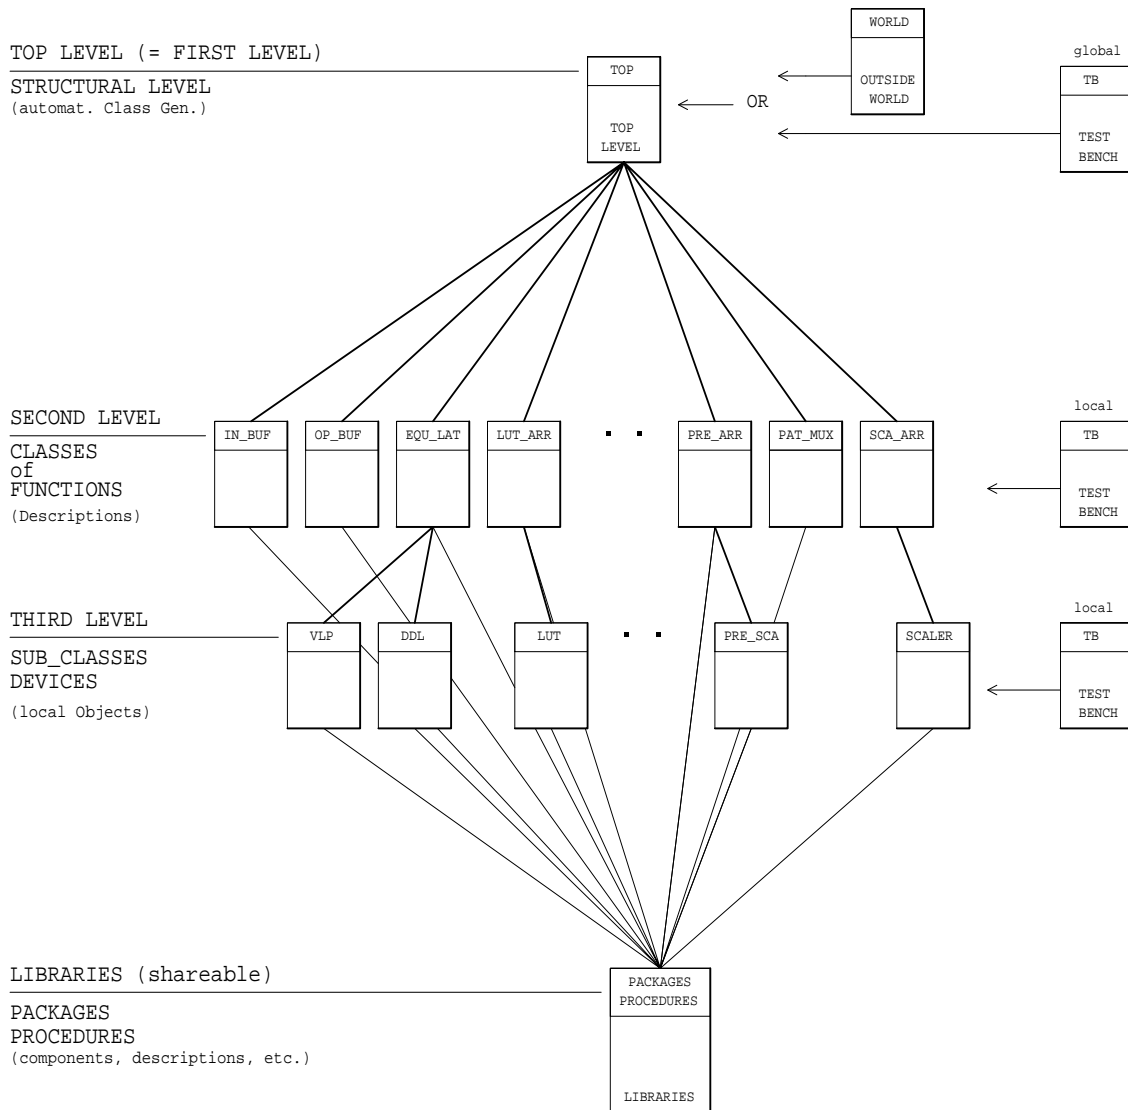


Figure 6: Interface Level Model (ILM)

- Final model: Functional Level Model (FLM); based on the ILM, it is expanded by importing to the design the final component models -- also called hardware metrics.

The top-level design unit of the CTP is a structural description and is common to the three models: SLM, ILM and FLM. The SLM version is the least complex description. It can be inserted in the full model of the trigger (and also of the whole experiment) for simulation providing that this model is also written in VHDL (portability). It permits higher simulation speeds.

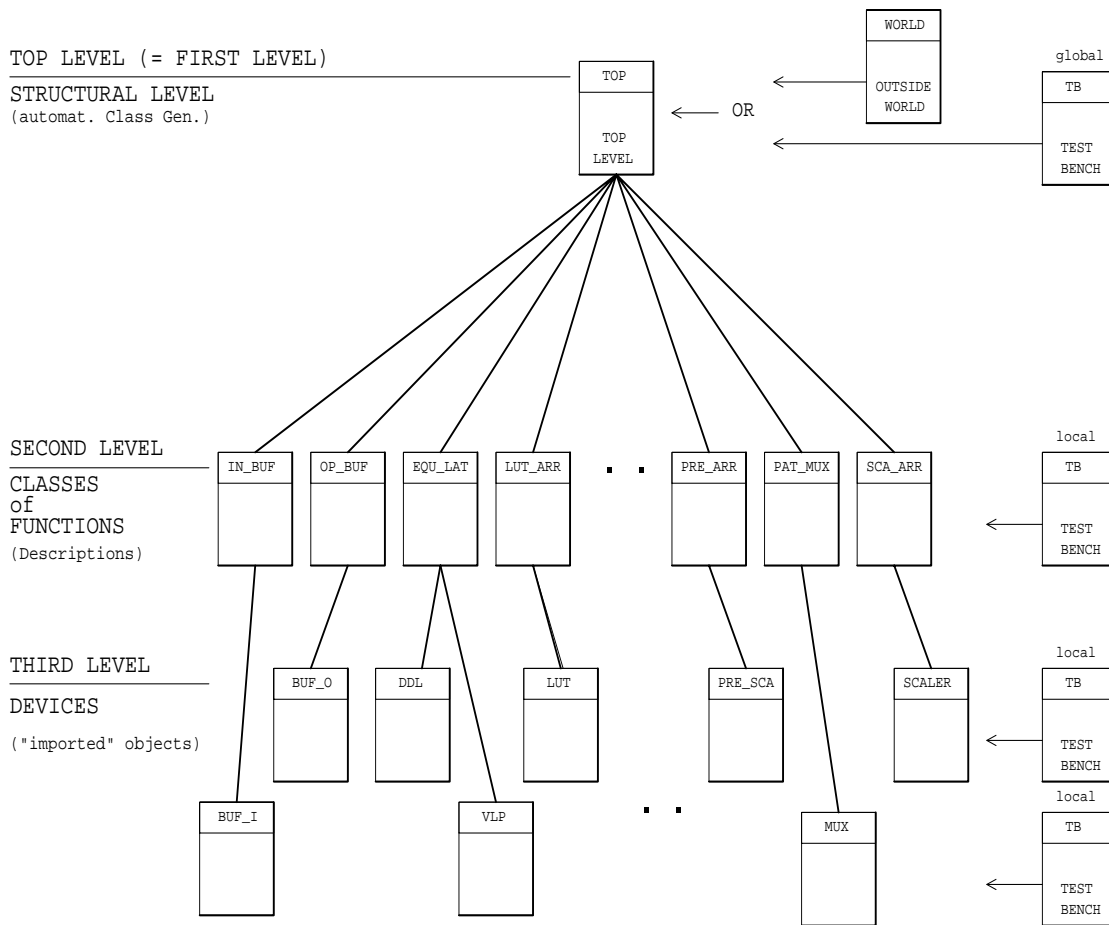


Figure 7: Functional Level Model (FLM)

The outside world -- subdetector trigger processors, Level-2 trigger processors, control and monitoring, etc. -- is connected to the top-level design unit. This structural level description generates and interconnects classes of functions -- second level of description -- such as input buffer, register, latency equalising circuits' prescalers, LUT arrays, deadtime, etc. (Figures 6 and 7) -- depending on definable parameters that provide for scalability. Classes of functions are present at the second level of the model and define the functionality of objects. Sub-classes (devices) represent the third level of the description and define local objects (Digital delay lines, VLP, scalars, etc.). Simulation can be performed by specific test-benches at any of the above-described levels, including the top-level design unit.

3.2 Realisation

This VHDL model permits insertion of lower level descriptions down to the detailed description of the final components with their associated (actual) timing parameters. Scalability is implemented by using the very powerful "generate" statement of the VHDL language. Generic parameters are used throughout the design. At the interface level rather elementary component descriptions and procedures have been written for the moment in order to permit a preliminary simulation of the complete model. It will be easy to replace these preliminary component descriptions with the final and detailed descriptions of the real circuits as soon as they are developed. Actual timing parameters from these components will also substitute the generic parameters defined

at the structural level by that time, thus permitting to simulate the CTP with accurate and up-to-date timing.

The structural level description incorporates the full description of the CTP.

3.3 Simulation

Simulation runs have been performed using the basic component descriptions with satisfactory results. Obviously many more simulations will have to be done: first with the development of FPGA demonstrators, and second when real circuits (eventually semi-custom ASIC's) will be available.

A full simulation of a single SDP bit yielding a Level-1-Accept is shown in Figure 8.




Fig. 8: Sample Simulation Result

4. Critical Components

In the context of an LHC experiment, the CTP itself is “critical” in the sense that it fulfils an extremely crucial function. It should be able to run at sustained data rates without any need for intervention. Speed and reliability are therefore two essential objectives of this work. With adequate circuit integration, we think that it should be possible to maximise both the speed and the reliability of the CTP.

Three critical components have been identified for which an integrated solution (either an FPGA, or semi-custom ASIC) may be appropriate: variable length pipeline (VLP), scaler and prescaler.

4.1 Variable Length Pipeline (VLP)

A circuit having the full functionality of the VLP not being commercially available, a special study is required. Figure 9 suggests a possible implementation of a cascadable (24 clock cycle length) 4-bit VLP slice ASIC.

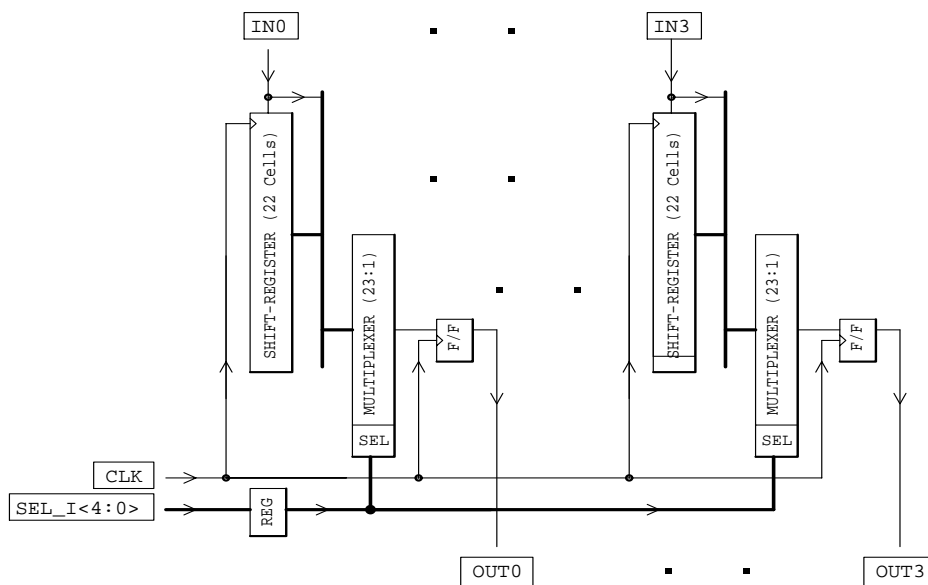


Figure 9: Four-bit Variable Length Pipeline Slice

4.2 Scaler

A possible ASIC design for a scaler circuit is shown in Figure 10. With this design, the available time for scaler readout is long and actually equal to almost a whole LHC revolution period -- 88.924 μ s -- this is achieved by copying the value of each scaler to an associated register. Registers corresponding to an array of scalers will be arranged to form a large shift register in order to permit a simple sequential readout.

Several of these scaler arrays (chips) can be cascaded so that the sequential readout is extended over several of them, this method having the obvious advantage of minimising the number of connections needed. The scaler arrays have a 12-bit dynamic range, allowing them to count up to any number of bunch crossings within one LHC turn.

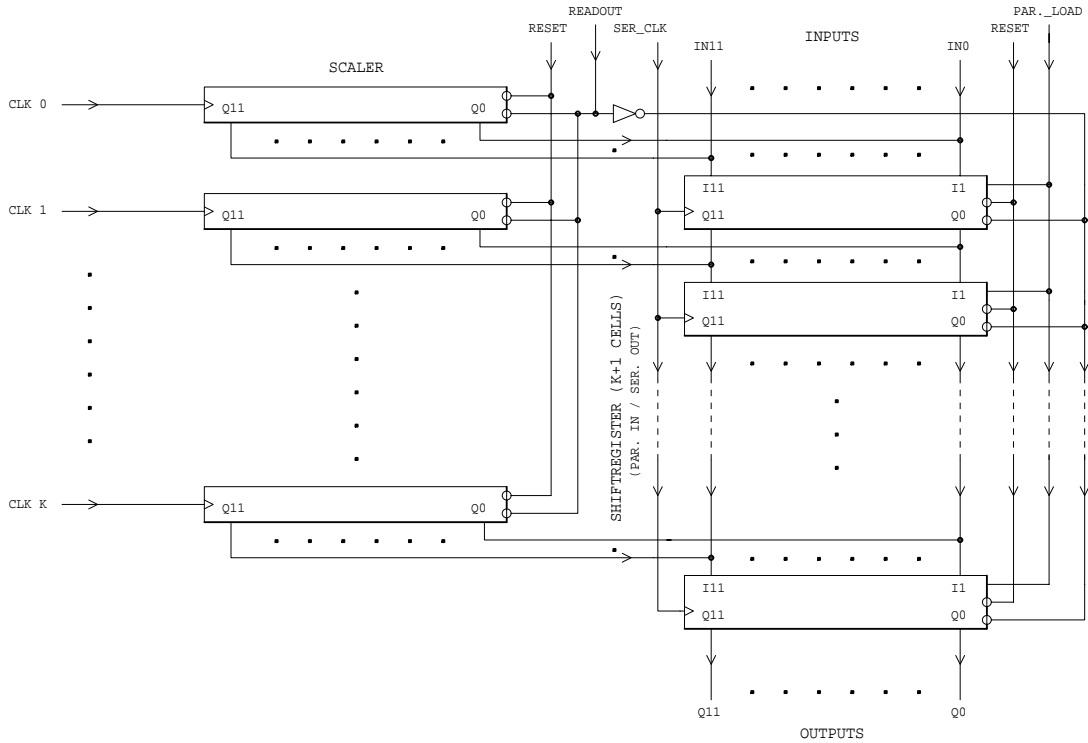


Figure 10: Scaler

4.3 Prescaler

The architecture of a possible ASIC design for a cascaded prescaler is shown in Figure 11. For similar reasons as for the scaler circuit, the dynamic range of the prescaler has been fixed to 12 bits, although a lower number of bits may be sufficient. The prescaler value could be preset once per turn or once every N turns, allowing full flexibility.

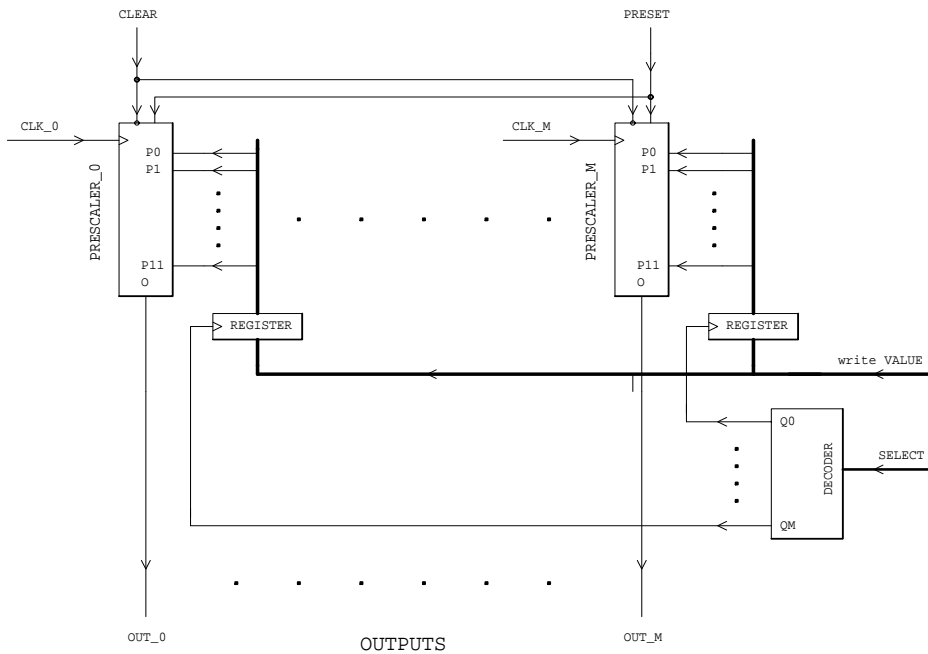


Figure 11: Prescaler

5. Control, Monitoring, Readout and Testability Aspects

5.1 Control

The Control part will consist mainly of the writing and reading of various parameters needed for the functioning of the CTP. Some of these parameters are listed below (from top to bottom following Figure 1):

- Phase Adjust (PA) settings,
- VLP lengths,
- LUT contents,
- VETO's for Trigger Candidates (TC),
- Deadtime parameters,
- Multiplexer setting (see 5.2),
- Slow general reset, e.g. at the time of a start of a new run,
- etc.

All these parameters listed above can be considered as “static” values, meaning that they are not subject to rapid or frequent changes -- obviously not at every bunch crossing but rather on a run-by-run basis.

The source of these parameters will be the Control System which needs to be defined but will be most likely based on a processor either within the context of the CTP or elsewhere.

Nevertheless a few signals will arrive at shorter intervals, for example the internal GEN_VETO signal (see below), a fast general reset signal received at the end of each turn of the LHC machine (via the TTC system) and prescaler count settings. Other parameters will be the Bunch Crossing IDentification number (BCID, also via the TTC system).

The internal GEN_VETO signal is an OR signal of the T1_Inhibit signal (Level-2 trigger system Supervisor) and a FIFO_FULL signal that indicates that any of the detector front-end derandomising FIFO's is unable to accept any further Level-1-Accept signals. The FIFO_FULL signal will be generated inside the CTP by applying a mechanism similar to method described in 2.4 (predicting Level-2 pipeline “full”) in order to determine early enough such a situation. More investigation is needed.

5.2 Monitoring

The Monitoring part will be more complex and extensive due to the massive amount of information to be collected and transmitted. Some of this information is listed below (from top to bottom following Figure 1):

- Scaler values from every SDP bit and from several levels -- synchronised SDP's after PA circuits and the alignment register (AR),

- Scaler values from every trigger combination (TC) -- after LUT's and those surviving the prescalers taken from the trigger register (TR),
- Scaler values from the OR gate output, after the AND gate (with the internal GEN_VETO) signal, after the deadtime circuit and the Level-1-Accept signal,
- etc.

The above listed scaler values will be -- in principle -- transmitted only once per LHC turn (during the next LHC turn) to the monitoring system. This monitoring system needs to be defined but will be most likely based on a processor either within the context of the CTP or elsewhere; e.g. a dedicated monitoring module (with on-board processor) within the CTP crate could build histograms from all scaler values and transmit them to the global monitoring system upon request. Important design aspects of the monitoring system are the bandwidth problem and the definition of buses and their sizes. Links to other part of the experiment requiring this information have to be defined too. Further investigation is needed.

5.2 Readout

Three pipelines (PL) are foreseen which will provide supplementary information with each BC clock about the following:

- Alignment register (AR),
- Status of all prescalers,
- Trigger register (TR).

This pipeline information will be used in the “time-window” circuit (Figure 12). All information must be brought in “phase” with the Level-1-Accept signal. The circuit should provide history information relative to a Level-1-Accept (- 5 BC clocks to + 5 BC clocks) thus permitting to identify a possible “pile-up” problem. The resulting information will be transmitted to the readout system. This complex circuit will need further investigation.

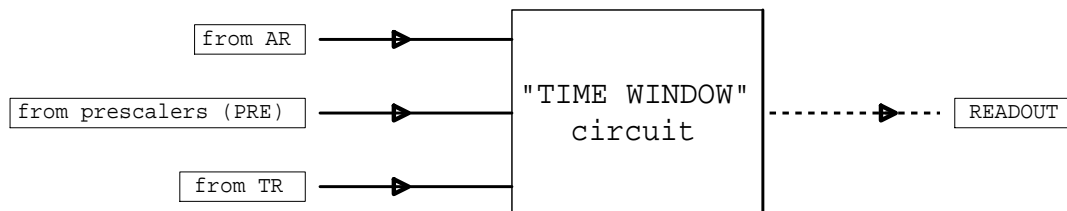


Figure 12: “Time-Window”

Upon each Level-1-Accept the alignment register (AR) and trigger register (TR) values will be transmitted to the Level-2 trigger processors. The information will be used to guide the process for Level-2 decisions. Additionally it will be sent to the monitoring system.

5.4 Testability Aspects

The total latency of the Level-1 trigger decision system -- subdetectors and associated trigger processors, central trigger processor and sending the Level-1-Accept signal back to each of the subdetectors -- should be less than a given value (currently in the order of 2.0 μ s). This latency includes physical cable runs (optical or copper cables) over the whole system. It seems therefore appropriate to place the CTP physically as close as possible to the subtrigger processors.

The Monitoring System is also a crucial part of the CTP. Any anomalies in the behaviour of the CTP should be detected by this system. Nevertheless a sudden change in the recorded average rate of Level-1-Accept signals could indicate the possibility of some incorrect behaviour of the CTP. In order to distinguish between real hardware errors and errors originating from the reception of “incorrect” SDP’s we propose to insert a multiplexer above the Alignment Register (AR) with the following inputs (Figure 13):

- Synchronised SDP’s,
- Data from a fast data PATtern MEMory (PAT_MEM),
- Data generated by a hardware random number generator (RND)⁶.

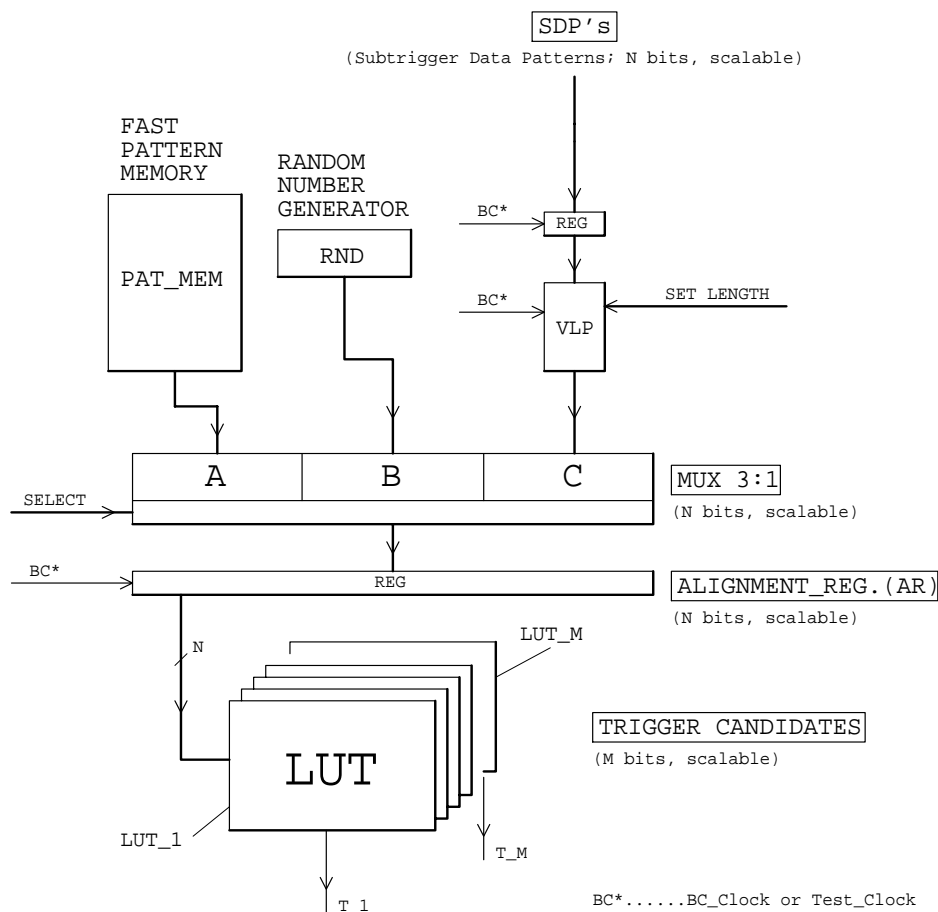


Figure 13: Detail of Insertion of PAT_MEM and RND

⁶ First studies by Miguel Angel Pantoja Molina (Summer Student 1993).

Two test facilities which permit verifications at the full speed of the CTP are described below:

The fast data pattern memory (PAT_MEM) will be a large and fast random access memory (RAM) -- arranged as a cyclic memory -- which can be loaded with SDP patterns generated by Monte Carlo simulations. Other possible SDP patterns:

- Continuous succession of only “good” SDP’s,
- Continuous succession of invalid SDP’s,
- “Good” SDP’s interspersed with several empty BC clock periods,
- SDP’s from previous data taking periods,
- etc.

The size of the PAT_MEM will have to be sufficiently large to contain a succession of SDP’s to cover at least one revolution of the LHC machine (2835 BC clocks). It can be extended to cover a larger period but precaution has to be taken in order to ensure the correct functioning of a large memory with the clock frequency of the LHC machine (40 MHz)⁷.

The PAT_MEM will also be used for the testing (commissioning) of the CTP without SDP’s from the detector. These tests can be done with either the BC clock from the LHC machine or with external clock frequencies ranging close the BC clock therefore permitting the detection of potential timing problems within the CTP hardware.

The hardware random number generator (RND) provides a quick first level test method for investigating Level-1-Accept rate deviations. The principle behind the use of a RND generator is the following:

Providing that the CTP has been tested with the PAT_MEM with positive results we can assume that the CTP functions correctly from the alignment register (AR) to the Level-1-Accept signal. We now use data from RND and record SDP bit occurrences and Level-1-Accept occurrences during a fixed time period (to be defined). This will provide us with a set of references of individual SDP bit-rates and the Level-1-Accept rate. These recorded values will provide a yardstick for the performance of the CTP with the specific RND generator and its initial settings. If the performance (rate) of Level-1-Accepts is doubtful, we can commute to the RND generator and run for an identical fixed time period in order to verify the CTP. Verification of now obtained Level-1-Accept rates can be compared to the initially obtained ones. In case of exact agreement between these two measurements we can determine that the error is not generated by the circuitry implemented between the Alignment Register (AR) down to the Level-1-Accept signal. The error must therefore be located between the arrival of the SDP’s (from the subtrigger processors) and the Alignment Register (AR) -- the link between the subdetector trigger processors (transmitters, cable and receiver), the input register, the VLP and the multiplexer of the CTP. Thus debugging could then concentrate only on those parts of the CTP.

⁷ The size and implementation of the PAT_MEM is a subject of further studies.

The above-stated possibilities for testing are tailored for high speed tests and verifications, and are obviously not exhaustive. Additional test facilities should be foreseen. Most of these methods will nevertheless make full use of monitoring facilities and will therefore not be executable at full speed. Further and extensive studies are required.

6. Current Status

For the second year of the RD27 project, the objectives for the design of the CTP had been defined as follows:

- 1) Write a high-level VHDL specification of the CTP architecture.
- 2) Write a low-level VHDL description of the three critical circuits -- VLP, scaler and prescaler.
- 3) Implement these three critical circuits with FPGA's (e.g., XILINX) at a speed as close as possible to the target speed of 40 MHz.
- 4) Build a test bench for these FPGA realisations.

At this point in time, items 1 and 2 are almost completed. Items 3 and 4 are well advanced -- a sample prescaler demonstrator has been developed, targeted for an XILINX FPGA (XC3142-3, PLCC84) on an APTIX prototyping system⁸. The corresponding logic analyser screen plot is shown in Figure 14.

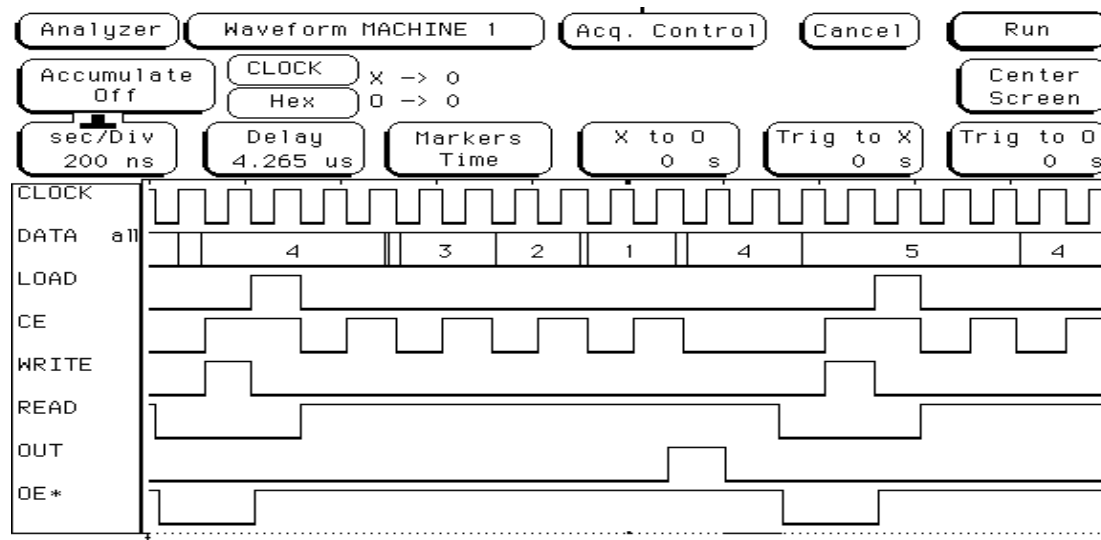


Fig. 14: Logic Analyzer Screen Plot for Sample Prescaler

Parameters:

- 12 bit resolution,
- Clock frequency = 10 MHz (limit of the APTIX prototyping system \leq 15 MHz).

⁸ APTIX prototyping system: Field Programmable Interconnect Component (FPIC) on an AXB-GP4 general purpose Field Programmable Circuit Board (FPCB).

7. Future Plans

During the first phase of this project we have concentrated our efforts on the top-level architecture -- data path -- and its VHDL description. Additional effort will be necessary to review the top-level VHDL description in order to include special circuits as the phase adjust circuit, the “time-window” circuit, details of the deadtime circuit, etc. We also continue to search for more optimised architectures and refined VHDL descriptions for the critical components concentrating on synthesisable code generation. Implementations of these components will be prototyped using FPGA’s. This will become more important in respect to a future realisation as new technologies become available that offer higher integration with improved speed parameters. Nevertheless the possibility of an even higher integration and potentially higher speed of Gate Arrays will lead to the implementation of some of the critical components using this technology.

The second phase of the project -- the lower level architecture -- comprising the control, monitoring and readout part will start soon. This part of the CTP is rather large and complex, and needs careful studies in order to achieve a manageable and compact system. At present it is not envisaged using VHDL for this part as such a description will be large and time consuming. We also think that it is not really necessary for the correct functioning of this part of the CTP. Nevertheless some part like the readout that will be connected to the whole trigger system of an experiment might need to be described using VHDL. We will also study the physical implementation of the CTP and the connections to the front-end trigger processors, the TTC system, the Level-2 trigger processors and trigger supervisors as well as connections to the overall monitoring system of future experiments.