

A data driven high performance Time to Digital Converter

J. Christiansen, A. Marchioro, P. Moreira, M. Mota, V. Ryjov, CERN/EP
S. Debieux, Geneva Engineering School
jorgen.christiansen@cern.ch

Abstract

A data driven multi-channel Time-to-Digital Converter (TDC) circuit with programmable resolution (25ps - 800ps bin) has been implemented in a 0.25um CMOS technology. An on-chip PLL is used for clock multiplication up to 320MHz from an external 40MHz reference. A 32 element Delay Locked Loop (DLL) performs time interpolation down to 98ps. Finally, finer time interpolation is obtained using an on-chip R-C delay line. Time measurements are processed and buffered in a data driven architecture based on time tags. This results in a highly flexible triggered or non-triggered TDC, which can be used in many different applications.

I. INTRODUCTION

Time to Digital Converters are needed extensively in high-energy physics experiments. In drift based tracking detectors a time resolution of the order of 1ns is normally sufficient. In Time Of Flight (TOF) detectors a time resolution of a few tens of Pico-seconds are often required.

The microelectronics group at CERN has previously developed several TDC's but these can unfortunately not be produced currently, as the IC technologies used have been phased out. The ALICE TOF and the CMS muon detector both need a highly integrated multi-channel TDC. The requirements from these two applications are quite different, but it was decided to make one common TDC. Significant savings can be obtained, when making one common circuit, if all requirements are taken into account from the beginning of the project. Savings in prototyping, testing and qualification outweigh the additional complications in the design. Merging the two projects also increases the total production volume (~20.000) giving a significant reduction in production costs. The features required to support two so different applications have resulted in a very flexible TDC which can be used in many other applications.

II. REQUIREMENTS

The requirements were defined from the CMS/ALICE needs plus an additional set of generally useful features. The high flexibility will enable the function of the TDC to be adopted to different working conditions and also ease debugging and commissioning of detectors. A short form summary of requirements is as follows:

- 1 ns – 25 ps resolution
- Self calibrating
- Reference: 40 MHz clock
- Dynamic range: up to one LHC machine cycle
- 32 channels low resolution
8 channel high resolution
- Leading and/or trailing edge or leading + width
- Hit rates: few hertz to few MHz
- Triggered and non triggered mode
- Programmable trigger latency
- Trigger rate up to 1 MHz
- Support for overlapping triggers
- Parallel and serial readout with token passing
- JTAG boundary scan + testing features
- SEU detection

III. ARCHITECTURE

Two completely different architectures have initially been considered. In a pipelined TDC a time measurement can be injected into a synchronous pipeline each clock cycle. Such an architecture is conceptually quite simple, but is not useful for non-triggered applications and has a problem supporting overlapping triggers in large time windows. The trigger latency is also limited directly by the depth of the pipeline buffer. An advantage of this architecture is its insensitivity to channel occupancy as it can accept one hit per clock period.

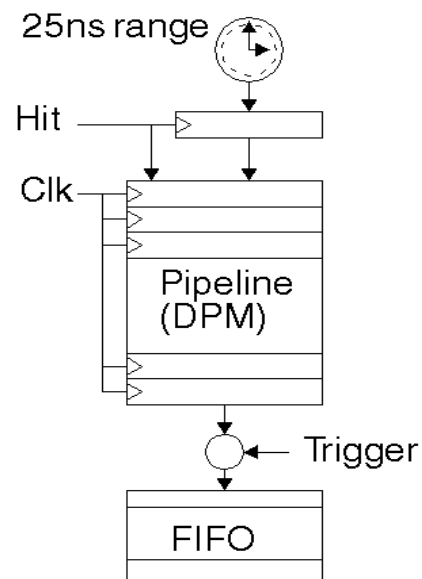


Figure 1: Pipelined TDC

A data driven TDC only stores hits in a buffer when a hit has been detected. This approach can be used in a triggered and non-triggered mode. In the non-triggered mode the buffer is used as a simple FIFO. In the triggered mode, time measurements from the buffer are compared to a trigger time tag and hits located inside a given time window are extracted. Overlapping triggers can be supported, by using a dual port RAM and a slightly extended matching mechanism. The basic time measurement must cover the complete dynamic range of the TDC in contrast to the pipelined architecture that only needs to cover one clock cycle. The maximum trigger latency is given by the dynamic range, and not directly by the buffer depth. The buffer occupancy must though be carefully evaluated for the hit rates expected. In the data driven architecture it is also possible to merge hit measurements from several channels into one buffer using channel derandomizer buffers.

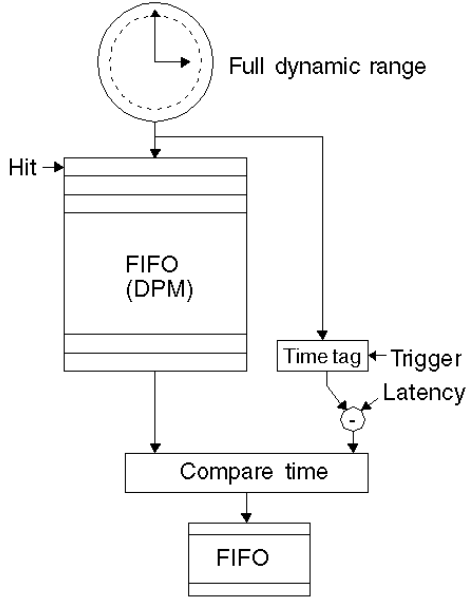


Figure 2: Data driven TDC

A data driven architecture offers higher flexibility at the cost of increased complexity. This makes it more sensitive to failures caused by Single Event Upsets (SEU). Buffer overflows in the data driven architecture must also be handled carefully to prevent the loss of event synchronisation in the system.

IV. HPTDC ARCHITECTURE

In the High Performance Time to Digital Converter (HPTDC) a time measurement is performed by storing the state of a clock synchronous counter and a 32 tap Delay Locked Loop (DLL). The position of the rising edge of the clock in the DLL gives an interpolation within the clock cycle and is also used to resolve uncertainties from the asynchronous sampling of the counter. By driving the counter and the DLL from an on-chip PLL, running with different frequencies, a time binning of 798ps/195ps/98ps is obtained. An additional high-resolution mode is also

available (see later). Time measurements are stored in a four deep asynchronous channel derandomizer buffer before being merged into a common latency buffer, shared by 8 channels. Measurements from four latency buffers are optionally trigger matched to trigger time tags from a 16 deep trigger FIFO before being passed to a 256 deep readout FIFO. Data is finally read out via a 32 bit parallel, 8 bit parallel or a serial readout interface. Merging event data from several TDC chips is handled by a token passing mechanism.




Figure 3: HPTDC architecture

A. R-C interpolation

A 98ps binning from the DLL is limited by the technology used. To obtain a 25 ps binning, a distributed R-C delay line with four taps covering 98ps can be used to drive four normal TDC channels as shown in Fig. 4. R-C delay lines have large process variations, but can be considered sufficiently stable over a temperature range of +/- 20°C. Compensation for process variations is performed with banks of capacitors that can be added as additional loads. This scheme has previously been demonstrated to give good results [1].

The PLL and the DLL are self-calibrating using the 40 MHz reference. The R-C delay line though requires calibration. A simple calibration procedure is a code density test with a sufficiently large population of random hits. A histogram of the number of hits in each of the four time bins from the R-C delay line is a direct measure of their relative size. The hits used for calibration only need to be random over a small time window covering the R-C delay chain (100ps) and can therefore in most cases be performed directly on physics data. For very high

precision measurements an additional option exists to adjust the integral linearity of the DLL.




Figure 4: R-C delay line interpolation

B. SEU detection

For the applications in mind the total radiation dose is assumed to be below 10krad and a standard technology and library can be used. Single event upsets must though still be taken into account, especially for a data driven architecture having rather complicated control logic. It was considered sufficient to implement SEU detection, as SEU immunity would carry a too large overhead. All internal memories have parity check on data and all programmable parameters have one common parity. Finally all state machines have been implemented with a one-hot encoding scheme, which allows the detection of any illegal state transition.

C. JTAG

JTAG boundary scan is required to perform efficient board testing. JTAG is also used to load programming parameters and to perform extensive production testing. The “raw” data from the time interpolation is accessible to perform special timing tests. All internal memories have BIST (Build In Self Test) controlled via JTAG. In a special test mode all internal flip-flops in the TDC are configured as a large shift register controlled from JTAG. Detailed status information (errors, buffer occupancies, etc.) can also be read via JTAG.

D. Architectural limitations

Merging measurements from 8 channels into one latency buffer limits the hit rates that can be sustained. The channel derandomizers together with a “fair” arbitration insure the best possible use of the bandwidth available (40MHz). For hit rates up to a few MHz per channel the loss of hits are insignificant. Above this limit the hit loss increases sharply as shown in Fig. 5. This is illustrated in Fig. 6 for 1 MHz and 4 MHz hit rates. The histograms show the number of hits in the individual channel buffers (max 4) and in all channel buffers in a group (max 8 x 4). If a channel buffer is full (4), when a new hit arrives, it is simply ignored. Higher hit rates can

be accommodated if fewer channels per TDC are used. The TDC also has the option of running the internal logic with an 80 or 160 MHz clock from the PLL. This gives a performance increase of a factor of 2 or 4 respectively, at the cost of increased power consumption.

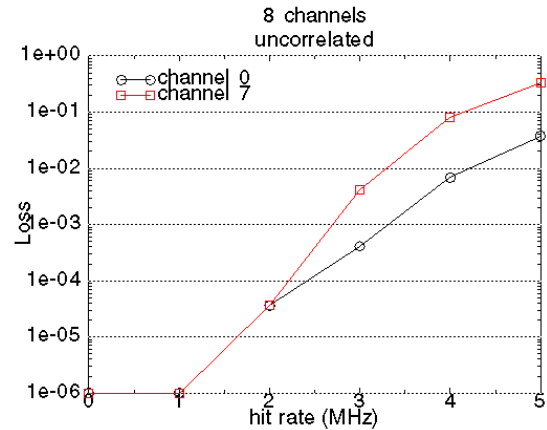


Figure 5: Hit loss as function of hit rate with (40 MHz clock).

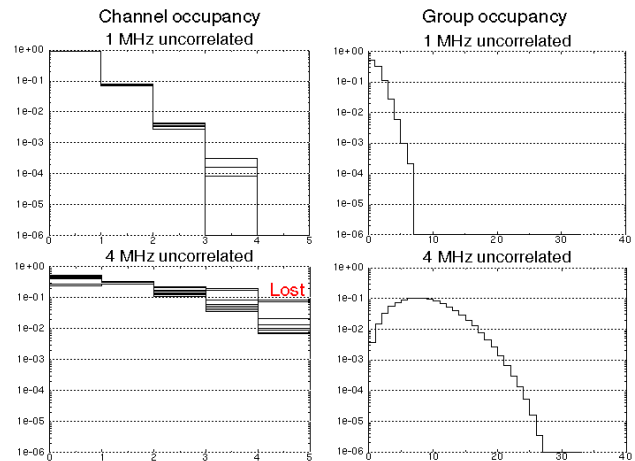


Figure 6: Derandomiser occupancies at 1 and 4 MHz hit rates.

In a data driven architecture, the effective trigger latency that can be accommodated depends on the hit rates. The average occupancy of the latency buffer can be calculated from: number of channels (8) x hit rate x trigger latency. To prevent buffer overflows from statistical fluctuations, it is recommended to keep this average below half the buffer size (256/2). The actual buffer occupancy is shown in Fig. 7 for 1 MHz hit rates and latencies of 10, 20 and 30 us. It is clearly seen that the buffer overflows in the last case, as the distribution is truncated at 256. The HPTDC handles this kind of overflows by marking any event that has lost hits.

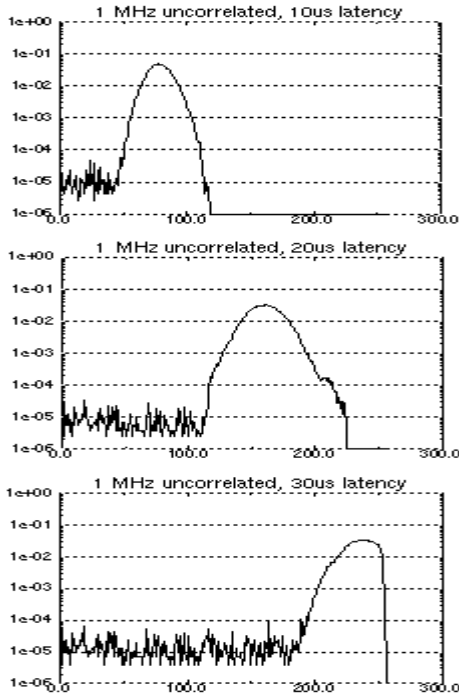


Figure 7: Latency buffer occupancies at 1MHz hit rates.

E. Programmable features

As previously mentioned, a high level of flexibility has been incorporated into the HPTDC. 700 bits of programming data are required for a complete configuration of the TDC. The main programmable features are summarised below.

- Resolution
- Integral error correction
- Channel offsets
- Leading/trailing/pair
- Channel enable/disable
- LVDS/TLL hit inputs
- Channel dead time (10 – 100ns)
- Encoding of Triggers and resets
- Trigger matching or no trigger matching
- Trigger latency
- Matching window
- Reject latency
- Limiting number of hits per event
- Readout FIFO size (<256)
- Readout of buffer occupancies per event
- Buffer back propagation scheme
- Serial, byte or parallel readout
- Readout via JTAG
- Serial readout speed
- Use of headers and trailers
- Token passing scheme
- Low power mode
- Test modes

V. IMPLEMENTATION

The HPTDC has been implemented in a 0.25 μm CMOS technology using a commercial standard-cell library and a memory macro. The PLL, the DLL and the channel buffers have been implemented as full custom, to obtain the required timing performance.

The logic behaviour of the design has been extensively verified at the register level. Simulations of realistic conditions (and extreme conditions) for different detector configurations, have shown that it is sufficiently robust. The large number of programmable features has required a significant effort to verify different combinations of features. A special design verification environment has allowed a large number of conditions and options to be verified in an automated way on large sets of hits with different statistical properties. The design has finally been mapped into gates using logic synthesis and the automated verifications have been re-executed to confirm its correct function and performance. The final design is 6.5 x 6.5 mm using ~ 1 million transistors and is packaged in a 225 pin BGA.




Figure 8: Picture of HPTDC

VI. MEASUREMENTS

The HPTDC has just recently been received from prototype production and is currently under test. The logic part of the chip has been found fully functional. The PLL and DLL have been found to be fully working with a combined jitter below 20ps (limited by test setup). A first preliminary measurement of the effective time resolution has been performed with time sweeps generated by a commercial IC tester with a timing resolution of 50ps. A RMS resolution of 236ps, 74ps and 49ps have been measured with the DLL running at 40 MHz, 160MHz and 320MHz respectively. This should be compared to the best possible RMS resolutions of 231ps, 56ps and 28ps.

For the high resolution modes the measurements are limited by the time resolution of the IC tester. Improved results are expected, when more accurate measurements, using a motor controlled trombone, will be performed shortly. The RC-delay line has been found to work, but sufficiently precise time measurements are not yet available to determine its detailed behaviour. As soon as more detailed measurements are made, they will be available on the web [2].




Figure 9: Error histogram with DLL working at 40MHz.




Figure 10: Error histogram with DLL working at 160MHz (limited by test equipment used)




Figure 11: Error histogram with DLL working at 320MHz (limited by test equipment used).

VII. CONCLUSION

A highly flexible high-resolution TDC has been implemented in a modern 0.25um CMOS process. The use of a data driven architecture has given a TDC that can be used in a large variety of high-energy physics experiments.

VIII. REFERENCES

- [1] J. Christiansen, M. Mota. IEEE Journal of solid-state circuits, pp 1360 – 1366, VOL. 34, NO. 10, October 1999.
- [2]: <http://pcvlsi5.cern.ch/MicDig/hptdc.htm>