

A data driven High Performance Time to Digital Converter HPTDC

- Requirements
- Pipelined or Data driven architecture
- HPTDC architecture
- Time measurement
- Data buffering
- SEU detection
- JTAG
- Implementation
- Timing performance
- Current status - History of bugs
- L1 buffer problem in latest version
- Users - Production planning


Email: JORGEN.CHRISTIANSEN@CERN.CH

TDC requirements

- High resolution for ALICE Time Of Flight (25 ps)
- Low resolution for CMS muon detector (1 ns)
- Multiple use of TDCs in HEP
- Dynamic range: One LHC machine cycle (12 bit 40Mhz counter)
- Hit rate: Few Hz - few MHz
- Leading edge and/or Trailing edge,
or Paired leading edge + Width (not in VHR)
- High integration level (32/8 channels per TDC)
- LVDS or TTL hit inputs
- Self calibrating using 40MHz clock reference (must be low jitter)
- Triggered or not triggered
- Trigger latency: 4.0, 3.2, 2.4, 1.2 us (programmable)
- Trigger rate: Few KHz to few hundred KHz (1MHz)
- Radiation: Total dose below 10Krad, SEU detection
- Low power
- JTAG boundary scan
- **High flexibility**

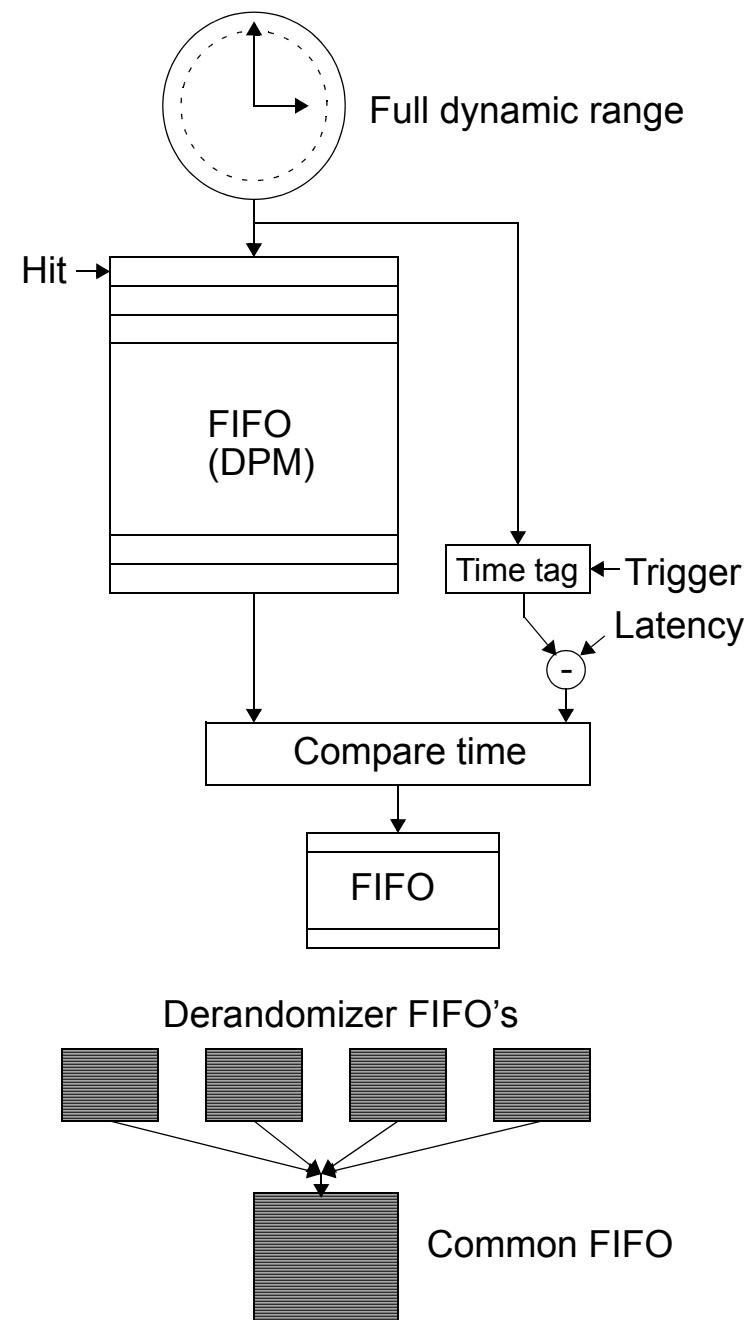
Traditional Pipelined TDC

- Stores hit data every clock cycle
- High hit rates (one per clock cycle)
- Fixed dead time
But limited double pulse resolution
- Fixed trigger latency (limited by buffer size)
- Only useful in triggered mode
- Difficult to support overlapping triggers
- No problem with buffer occupancies
- Narrow latency buffer (covers 25ns)
- Little sensitive to SEU in control part
- Simple architecture -> quick implementation
- Limited flexibility

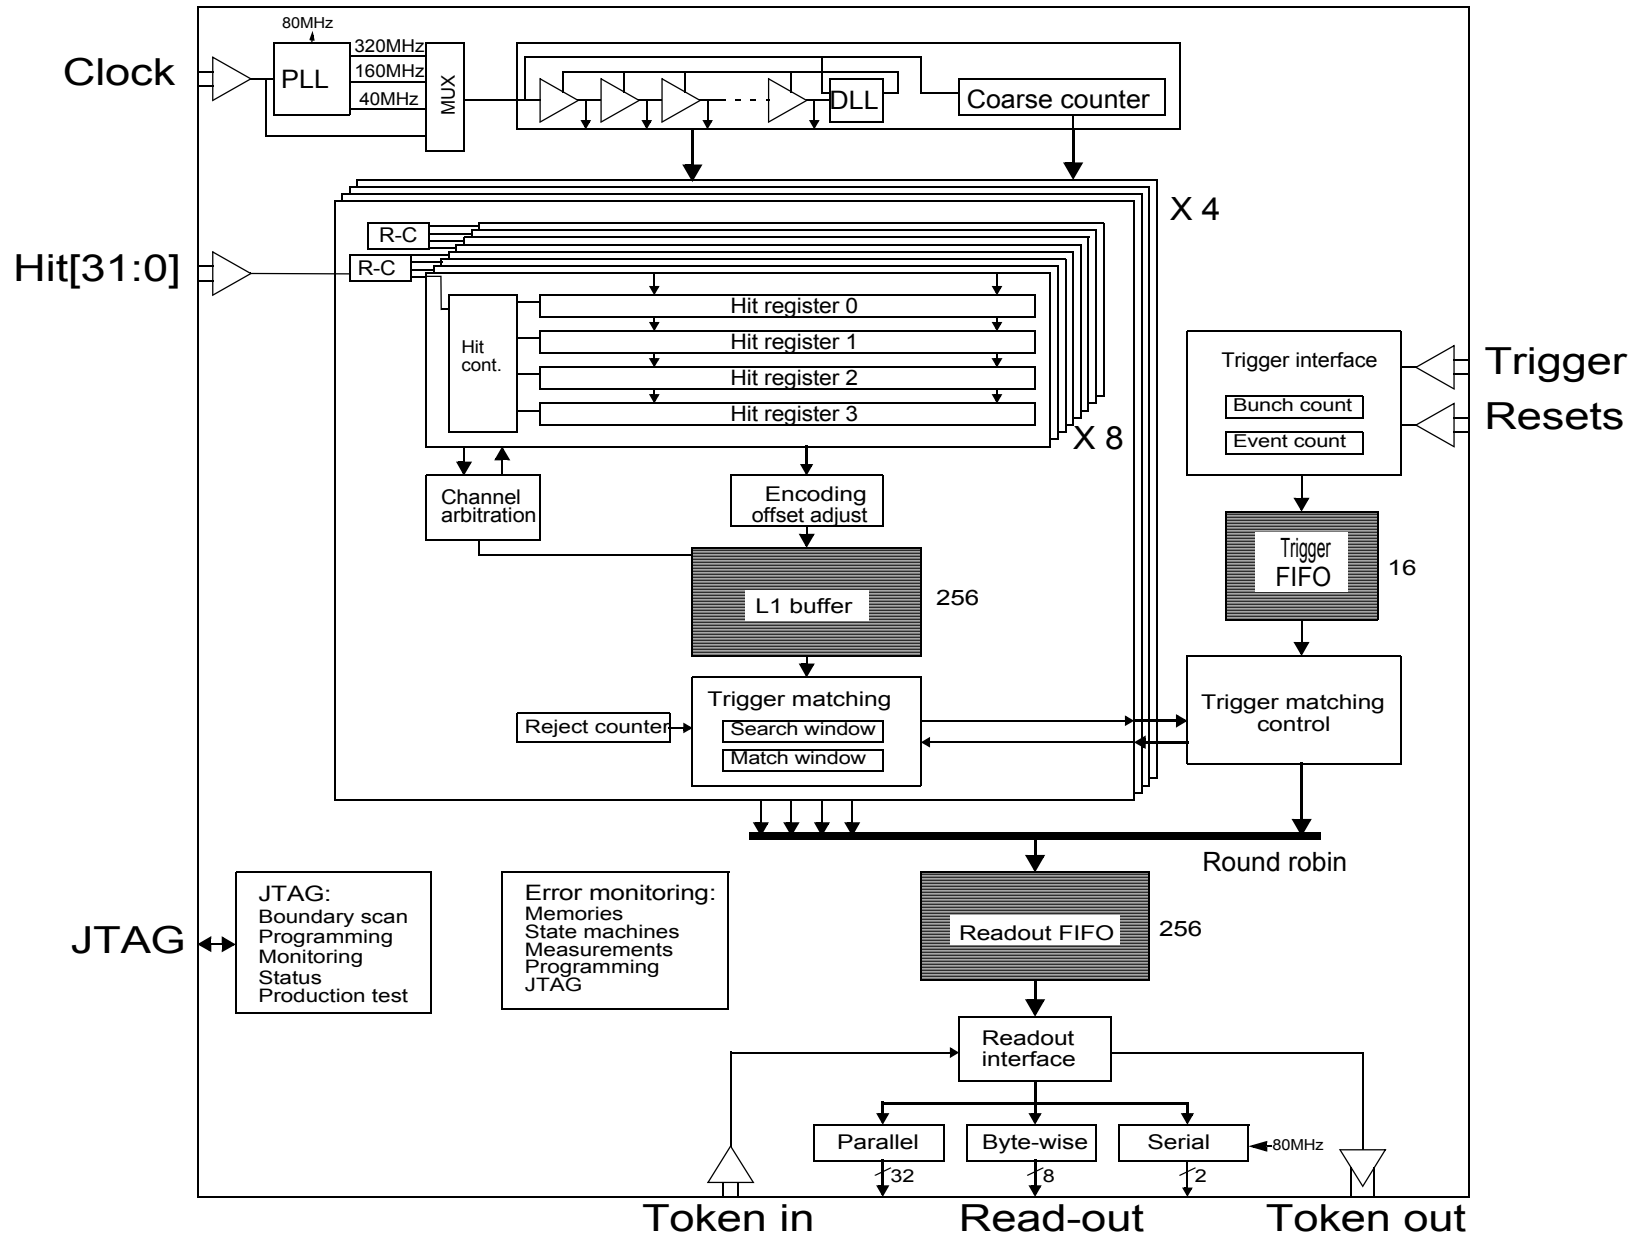


Data driven TDC

- Only stores data when hit detected
- Variable latency over full (1/4) dynamic range
Compromise between hit rate and latency
- Triggered / non triggered mode
- Multiple overlapping triggers
- Channel merging possible via derandomizers
Limits hit rates
- Good double pulse resolution
But complicated dead time analysis
- Buffer occupancies must be seriously analyzed
- Buffer overflows must be handled carefully:
Hit may be lost if marked
Complete events must never be lost
- Wide latency buffer (covers full dynamic range)
- More complicated architecture/implementation
Previous data driven TDC worked well in different applications
Logic complication handled by logic synthesis
Extended verifications at behavioral/register/gate level
- High flexibility

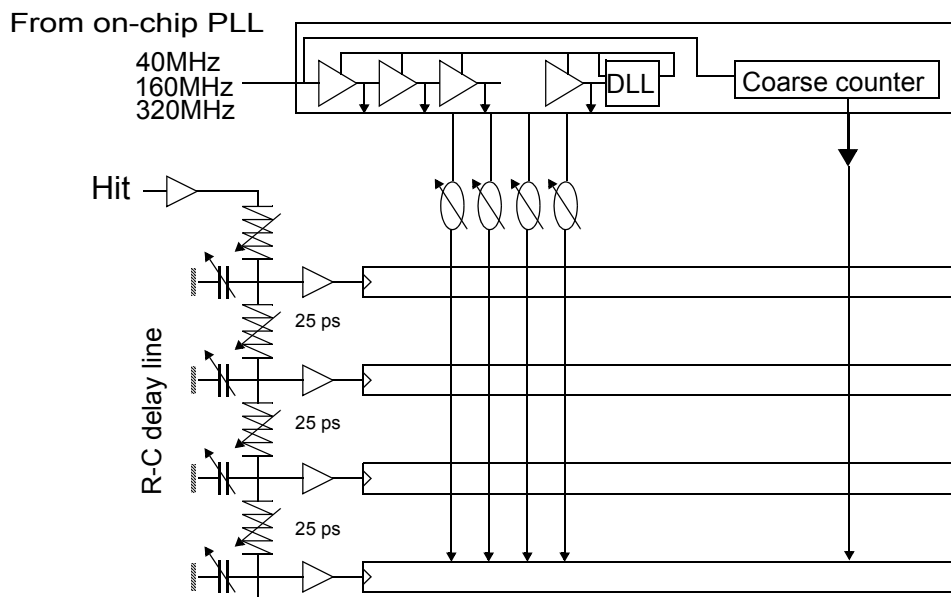


HPTDC architecture



Time measurement

- Coarse time: Count of clock periods
- Fine time: Extracted from Delay Locked Loop with 32 taps



- Low resolution: 40 MHz clock
781ps
- Medium resolution: 160MHz clock
195ps
- High resolution: 320MHz clock
98ps
- Very high resolution: 320 MHz clock and
Channel interpolation
24ps

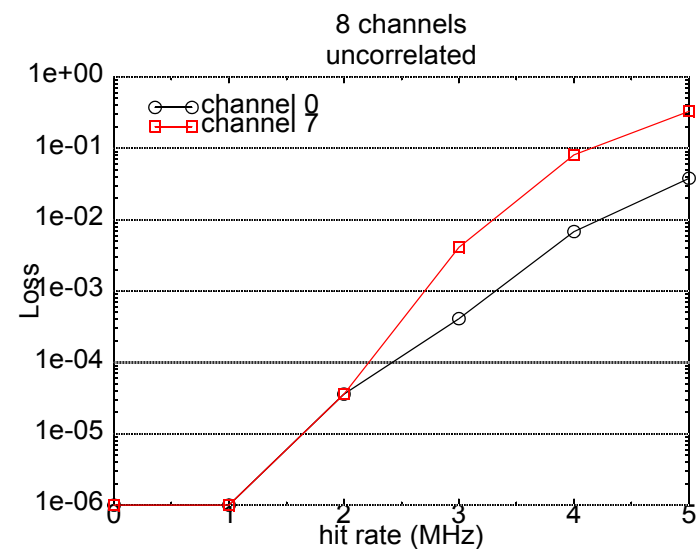

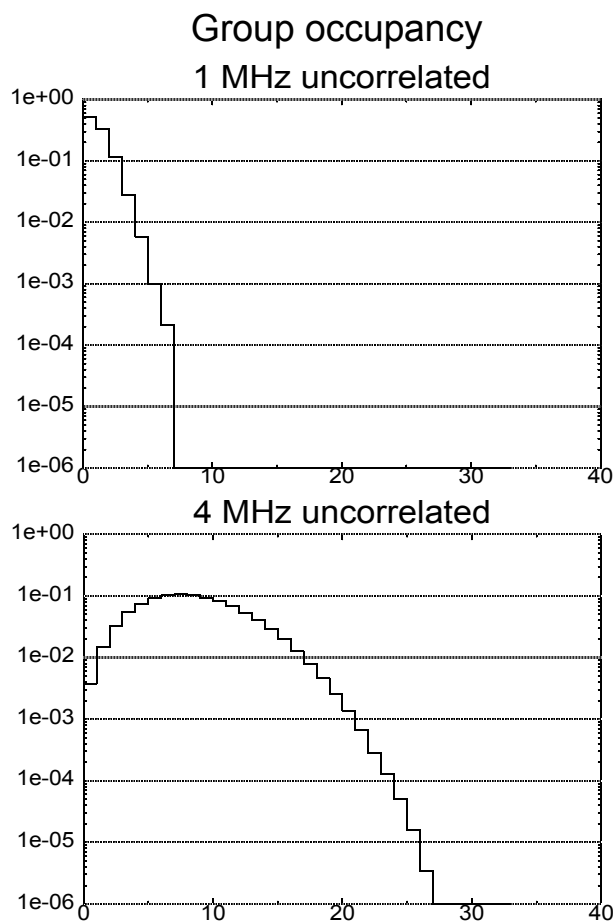
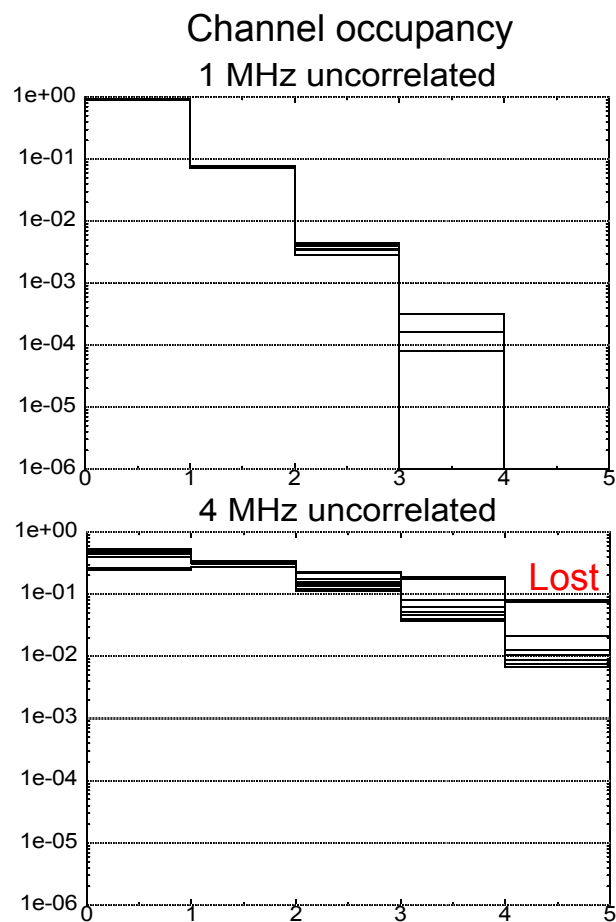
Very high resolution mode:

- R-C delay line dependent on IC processing
- R-C delay line independent of temperature ($\pm 20^{\circ}\text{C}$)
- Infrequent calibration required (once)
- Simple calibration using code density test
- Calibration can be performed with physics data (random over 100 ps range)
- Option of correction of integral errors from DLL

Limitations from using four channels:
8 channels per chip
Hit rates 2-4 times lower
Not possible to have paired leading edge and pulse width

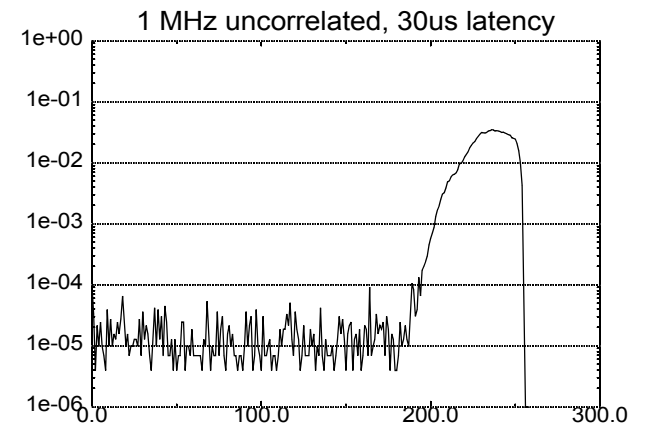

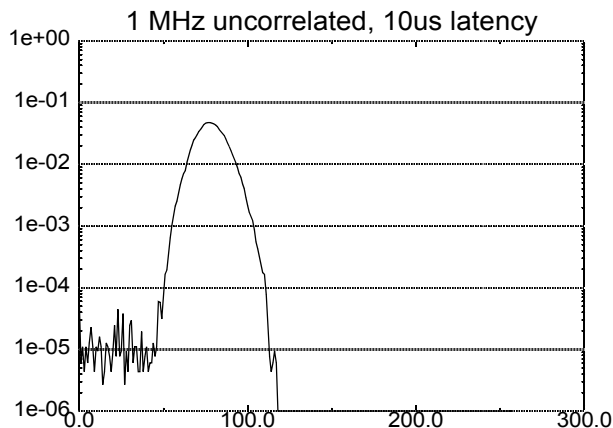
Channel merging

- Hit measurements are derandomized in a 4 deep asynchronous FIFO buffer
- Hits from 8 channels are merged into one L1 buffer
- Arbitration between channels made to be “reasonable” fair




L1 buffer

- Zero suppressed data only, max 256 hits per 8 channels
- Max latency given by dynamic range (1/4)
- High hit rate - > short latency
Low hit rate -> long latency
- Average buffer occupancy:
 $\text{channels}(8) \times \text{hit rate} \times \text{latency} < \sim 1/2 \text{ buffer size } (256)$
- Events with hits lost from L1 overflows marked





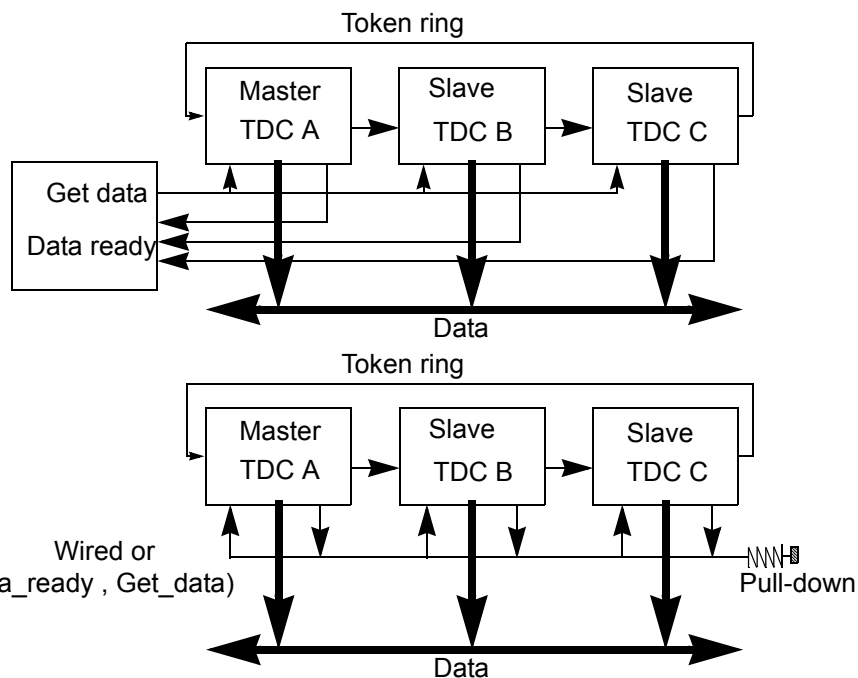
Trigger matching

- Trigger matching based on hit measurements and a trigger time tag
- 16 deep trigger FIFO to receive new triggers while matching function busy
- Trigger matching based on coarse count (25 ns resolution)
- Programmable latency and matching window
- Supports assigning hits to multiple overlapping triggers
- Reject function to remove old hits when no triggers waiting
- Works across counter overflows (3564 for LHC)
- Maximum number of hits per event programmable
- Trigger matching can be disabled



Readout

- 256 deep readout FIFO to de-couple matching and readout
- Readout FIFO can artificially be reduced to prevent data pile-up
- Token based sharing of readout port with bypass option
 - Triggered: Token only passed when all hits in event have been read out
 - Non triggered: Token must be constantly circulating to find TDC's with data
- 32 bit parallel readout for high rate applications
- Byte-wise readout to driver commercial serializers
- Serial readout for low rate applications
- Readout via JTAG possible for debugging



Increasing performance

- **Increasing internal logic clock frequency from 40 to 80 MHz(160):**
 - Clock from internal PLL
 - Higher hit rates can be handled x2, (x4)
 - Trigger matching speed improved (for high trigger rate applications)
 - Occupancy of L1 buffer in principle NOT improved (given by latency)
 - Power consumption increases
 - IO interface kept at 40 MHz
 - Chips currently only production tested at 40 MHz
- **Using fewer channels per channel group (8)**
 - 4 channels (16 per chip) -> double hit rate,
 - L1 buffer occupancy reduction
 - Less than 4 channels per group does not bring significant improvement

SEU Handling

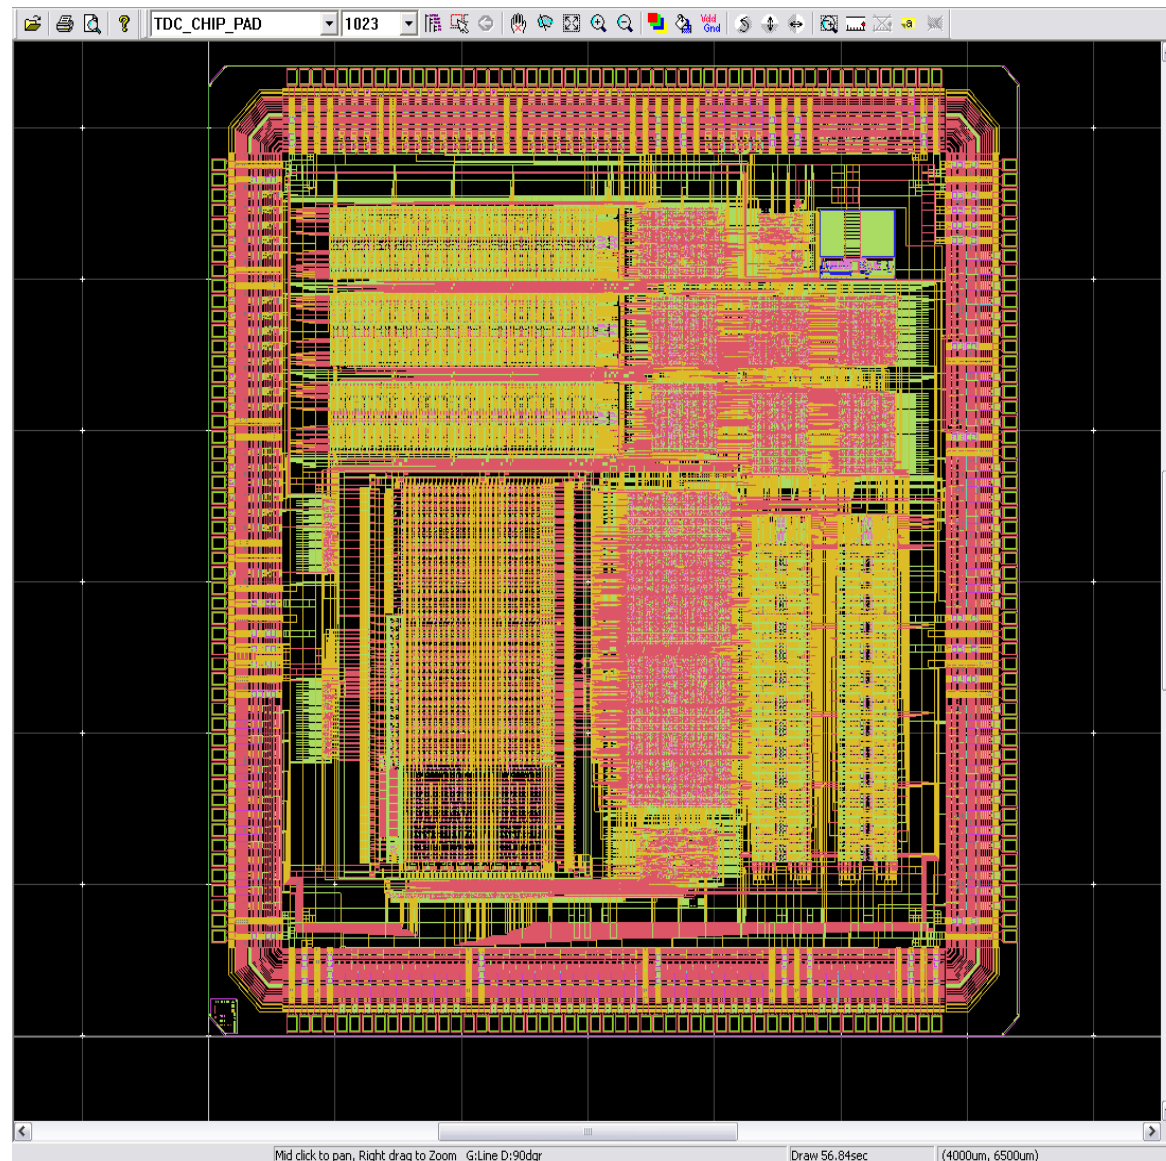
- SEU detection (not SEU immunity)
- Programming data protected with parity check
- All internal memories has parity check
- State machines implemented with one hot encoding and continuous state check.
- Measurements with parity error ignored in matching
- Error status with information about detected parity errors from different functional blocks.
- Programmable global error state which can force TDC into a passive state

JTAG

- Programming: ~700 bits
- Status: Errors, buffer occupancies, etc. = ~60 bits
- Option of readout via JTAG
- Boundary scan
- BIST of on-chip memories
- Scan path to verify time measurements: 750 bits
- Test scan path of all internal flip-flops: 2K bits

Implementation

- Architecture simulated extensively at behavioral level (Verilog simulation environment available)
- Mapped into gates (standard cells) with logic synthesis
- DLL, hit registers, RC delay and PLL implemented as full custom
- 0.25 um CMOS technology
- 6.5 x 6.5 mm
- 1 million transistors
- 225 ball grid array package

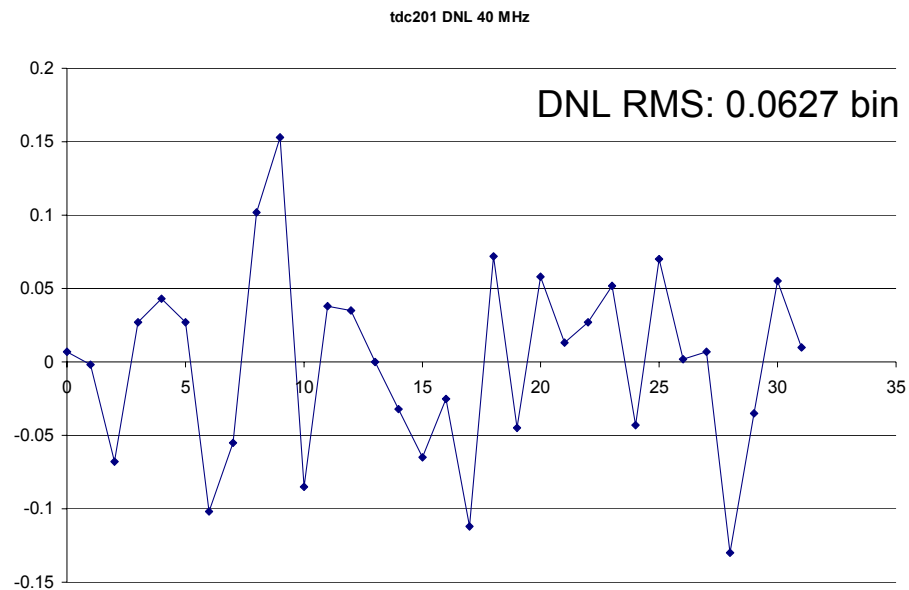


Programmable features


- Resolution
- Integral error correction from DLL
- Channel offsets
- Leading/trailing/pair
- Channel enable/disable
- LVDS or TTL hit inputs
- Channel dead time (5 - 100ns)
- Encoding of triggers and resets
- Trigger matching/no trigger matching
- Trigger latency
- Matching window
- Reject latency
- Roll-over and machine cycle separators
- Limiting number of hits per event
- Readout FIFO size
- Readout of buffer occupancies per event
- Buffer back propagation
- Serial, Byte, 32 bit Parallel or JTAG readout
- Force specific readout pattern
- Serial readout speed (80Mbits/s - 0.3 Mbits/s)
- Use of event headers and trailers
- Token passing scheme
- Generation of global error state
- Low power mode
- DLL and PLL control parameters
- Test modes

This large set of programmable features has made (is making) design verification difficult.
Flexibility does not come for free.

Timing performance: low resolution mode




Effective RMS resolution: 261 ps



High resolution

tdc201 DNL 320 MHz

DNL RMS: 0.0952 bin




Effective RMS resolution: 48ps


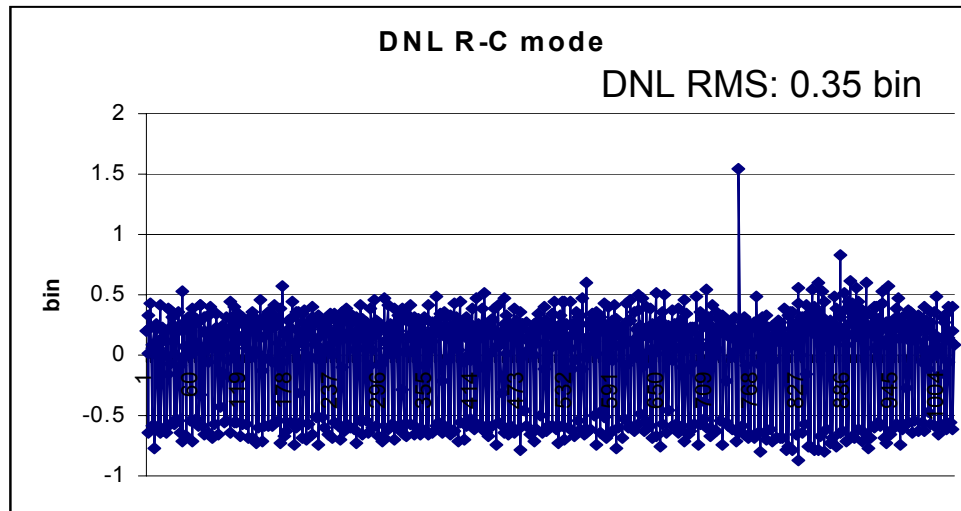
Effective RMS resolution: 34ps with table correction

tdc201 INL 320 MHz

INL RMS: 0.5089 bin



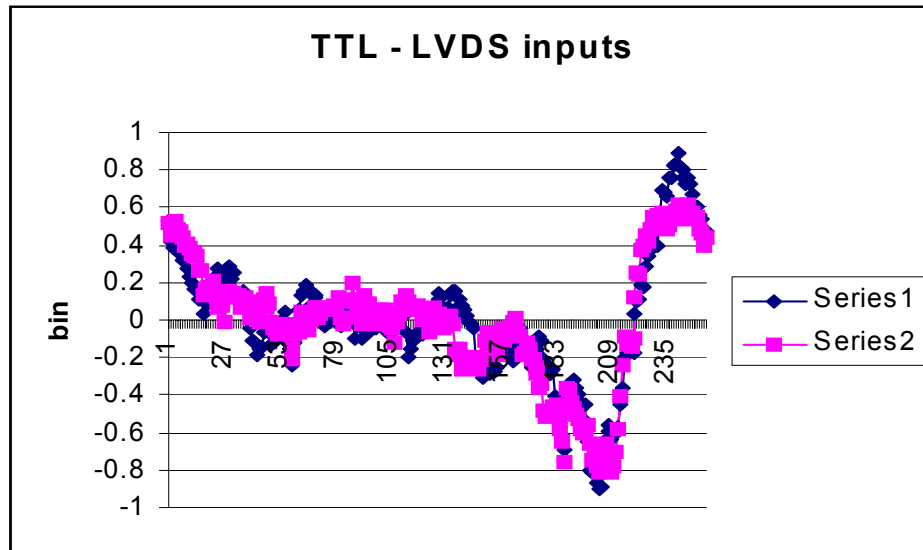
Very high resolution (R-C mode)



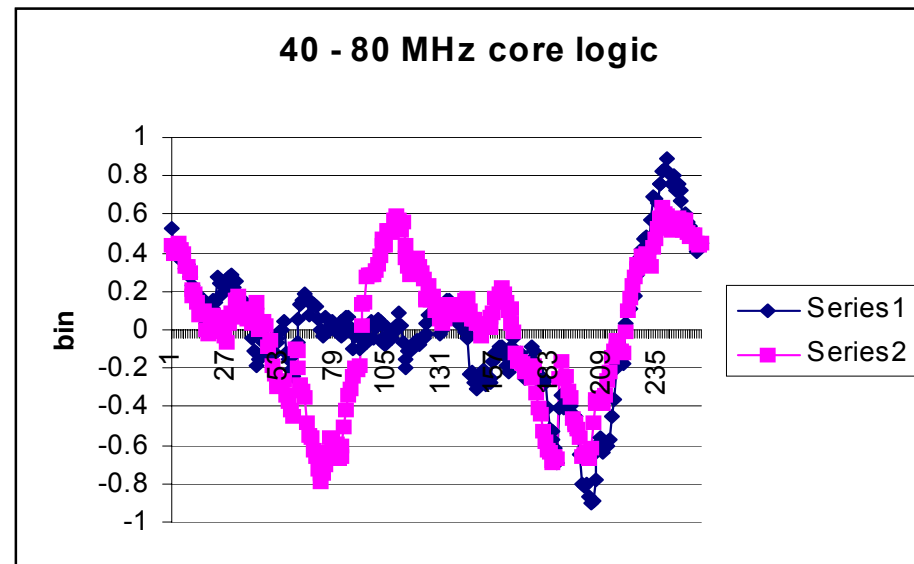
Effective RMS resolution: 40ps without correction

Effective RMS resolution: 17ps with table correction

Cause of INL error



It is clear that the INL comes from on-chip crosstalk from logic part of chip.



As logic clock is the same as the input time reference the INL is a fixed shape which can be compensated for if needed

Status and history

- **Version 1: MPW**

Functionally working

Only 2.5 volt IO levels

Potential problem with high power mode startup for low power applications.

INL problem in high resolution modes

- **Version 2: Engineering run**

3.3 volt IO levels

Ensured to start up in low power mode

INL improved using optimized clock signal distribution and separate power supply for clock drivers. But not yet satisfactory.

New: Hit registers was in some cases found to loose hit information if Vdd decreased below 2.30 - 2.40 volt.

Traced to possible sensitivity to relative N/P MOS parameters

New: DLL lock problem on some chips in low resolution mode at increased Vdd. Traced to possible sensitivity to relative N/P MOS parameters

New: R-C delay line adjustment in some cases problematic

- **Version 3: Engineering run with few modified layers**

Hit register problem resolved by resizing relative N/P transistors

DLL lock problem resolved by resizing relative N/P transistors

R-C delay adjust rescaled to fit observed process parameters.

INL improved by alternative powering scheme. Still not perfect INL in very high resolution mode but satisfactory using simple table lookup correction.

New: L1 buffer parity error at increased Vdd

L1 buffer problem

- The problem

L1 buffer parity error gets detected internally in chip depending on use and logic core power supply voltage.

L1 buffer parity error gets set in JTAG status and measurement is ignored in trigger matching (not read out)

Problem occurs at increased Vdd (highly uncommon failure mechanism)

Seen for first time in latest version (Earlier version do not have this problem up to 3.0 volt)

Nothing has been changed in interfaces to L1 buffers.

Memory macro bought from library supplier recommended by IBM but company do not any more exist

We only have layout and verilog simulation model of memory macro. No schematics that will allow us to try to understand problem by simulations

Low yield also observed in latest version

First production test patterns did not clearly identify chips with this problem.

Current production test still does not seem to be capable to provoke all failures as ~10% of chips delivered to CMS have been found to have problems when mounted on their boards

• Characteristics of the problem

Error does not occur systematically (partly random that makes identification of cause very difficult)

Error depends on Vdd (2.3 - 3.0 volt)

Voltage level to provoke errors is pattern dependent.

Some systematic difference between the four L1 buffers in the chip

Does not seem to depend on activity between channel groups


Some dependency on logic core frequency observed

Does not seem to depend on temperature

No problems with Readout FIFO observed

Latest statistics:

Functional fail:	32%
L1 buffer fail below 2.5 volt:	21%
L1 buffer fail 2.5 - 2.7 volt:	33 %
OK above 2.7 volt:	14%



• Possible reason

1: Unfortunate process parameters that provokes problem (process parameters within defined limits)

2: Alternative power routing used in last version could have had an influence but systematic tests do not seem to fully confirm this

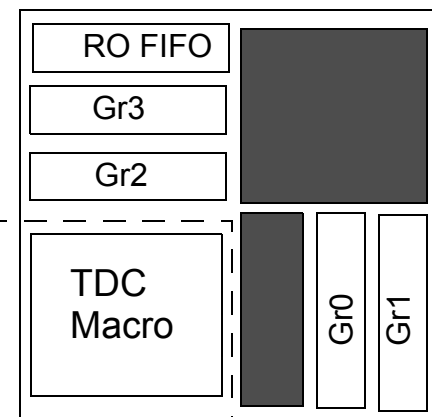
• Remaining tests to be made:

Make test that can provoke failures at same voltage level as seen by CMS

Determine, if possible, if it is a specific bit (difficult as seems random)

Make isolated test of Readout FIFO that uses same memory macro but does not seem to pose problems (isolated test of this is though difficult)


More thorough study of possible clock speed, temperature, etc. effects



Power cut

Failure voltage measured on IC tester

Passed chips



Chips passing production test at 2.6 volt


Group 0: 3.09v

Group 1: 3.12v

Group 2: 2.91v

Group 3: 2.84v

unstable



Chips failing production test at 2.6 volt

Group 0: 3.06v

Group 1: 3.11v

Group 2: 2.81v

Group 3: 2.74v

Returned chips from CMS

		All groups	Group 0	Group 1	Group 2	Group 3
R66 01	ROB 66 - TDC 1	2.46 (2.61)	-	-	2.44 (2.607)	-
R16 03	ROB 16 - TDC 3	2.48 (x)	-	-	-	2.46 (x)
R7 02	ROB 7 - TDC 2	2.44 (2.782)	-	-	2.68 (2.771)	2.42 (2.831)
R145 02	ROB 145 - TDC 2	2.44 (2.628)	2.42 (2.618)	-	2.69 (2.919)	2.63 (2.716)
R13 01	ROB 13 - TDC 1	2.35 (2.694)	-	-	-	2.35 (2.667)
R13 02	ROB 13 - TDC 2	2.28 (2.623)	-	2.63 (2.804)	2.26 (2.612)	2.59 (2.732)
R13 03	ROB 13 - TDC 3	2.39 (2.738)	-	-	2.48 (2.886)	2.38 (2.711)

It is extremely important that I manage to make a test that provokes the problem at the same voltage levels as seen by the users

I need help from users to know at which voltage levels they encounter problems

How to solve this problem

- A: Spend significant amount of time to understand the problem and come up with a design change that may possibly solve the problem.

This will require time and a new design submission with no guaranteed result.

- B: Hope that we are lucky and first production run will not have the problem

- C: ALTERNATIVE “BRUTE FORCE” APPROACH

Run the chips with a lowered core power supply (2.3 volt).

IO will still be 3,3 volt

Does anybody have problems with this ?

Summary

Number of channels:	32 / 8
Clock frequency:	40 MHz external 40MHz / 80MHz / 160 MHz / 320 MHz internal
Resolution:	781 ps (261ps RMS) low resolution mode 195 ps (64 ps RMS) medium resolution mode 98 ps (48 ps RMS) high resolution mode 24 ps (40 ps RMS) very high resolution mode (8 channels) 24 ps (17 ps RMS Corrected)
Dynamic range:	102 us
Double pulse resolution:	5 - 10 ns depending on mode
Hit rate:	Core logic at 40 MHz, Not R-C mode Max. 2 MHz per channel, all 32 channels used Max. 4 MHz per channel, 16 channels used.
Event buffer size:	4 x 256
Read-out buffer size:	256
Trigger buffer size:	16
Power consumption:	300mW - 1500 mW depending on modes.
Hit inputs:	LVDS or LVTTTL

L1 buffer parity error is the major remaining problem

Users and Quantities

User	Contact person	Parts	Delivery
CMS muon	Carlos Willmot	~10000	Samples
			800: Mar. 2003
			1400: Q2 - Q4 2003
			~8000: Q1 2004
ALICE TOF	Pietro Antonioli	24000	Samples
NA48	Sergei Basilev	50	42: Mar. 2003
CAEN	Carlo Tintori	2500	Samples
BES	Jiang Xiaoshan	1400	Samples
RICE	Geary Eppley	?	Samples
Oku	Andrei Siderov	~200	Samples
ATLAS CTP	Georges Schuler	30	Samples
HYTEC	Alan Burley	?	Samples
Orsay	Robert Sellem	?	
Tata institute	Suresh Tonwar	50	
Upsala	Leif Gustavson	?	Samples
LHC machine	Javier Serrano	~1000	Samples
Imago	Dan Lenz	~100	Samples
Ionwerks	Al Schultz	?	Samples
LHCb	Albert Zwart	?	Samples
Frankfurt	Kolja Sulimma	?	Samples
Alice V0	Gwenael Morishau	?	
Struck	Mathias Kirsch	?	

Users from outside HEP community needs special contract with CERN to be capable of buying chips

Users from non EU and US countries may need special export permission from the American department of commerce

Production

- Production lot is 48 wafers with 600 chips per wafer
Yield 80%: $0.8 \times 48 \times 600 = \sim 23\text{k}$ chips
Yield 40%: $0.4 \times 48 \times 600 = \sim 12\text{k}$ chips
- Planning will depend strongly on how to handle L1 buffer problem !
- Possible proposal
First lot mid - end 2003
Second lot begin 2004