

Sub-ns timing for the Pacific Ocean Neutrino Experiment by optical fiber using Gigabit Ethernet

Michael Böhmer,^{a,*} Roman Gernhäuser,^a Lea Ginzkey,^a Rob Halliday,^{bc} Christian Spannfellner^a and Nathan Whitehorn^b for the P-ONE collaboration

^a*Technische Universität München, School of Natural Science,
James-Frank-Strasse, 85748 Garching, Germany*

^b*Michigan State University, Department of Physics and Astronomy, 567 Wilson Rd,
East Lansing, MI 48824, USA*

^c*Elmhurst University, Department of Physics, 190 S Prospect Ave,
Elmhurst, IL 60126, USA*

E-mail: mboehmer@ph.tum.de

Precise timing inside the Pacific Ocean Neutrino Experiment (P-ONE) will be mandatory to allow synchronization of the DAQ systems inside the P-ONE Optical Modules (P-OMs). The DAQ system is designed to digitize 16 channels of photomultiplier tube (PMT) waveforms at high speed, with a large buffer and waveform pre-processing capabilities. This requires a fast and reliable data path between P-OMs and shore, as well as sub-nanosecond synchronization amongst P-OMs, while keeping the cable design as light-weight as possible to reduce mechanical stress. To meet these requirements, an optical network based on standard Gigabit Ethernet (GbE) on Lattice ECP3/ECP5 FPGAs is presented. It allows synchronous operation of the P-OMs by propagation of a central 125 MHz clock in addition to coarse and fine delay measurements on the point-to-point fiber links, and provides Gigabit Ethernet to all endpoints in the network. This allows precise timing and use of commercial off-the-shelf network components inside the P-OMs and the connected mooring junction box (mJB), serving as a network hub. This greatly simplifies the design process and makes well-tested, easily integrated network standards available for data transport.

38th International Cosmic Ray Conference (ICRC2023)
26 July - 3 August, 2023
Nagoya, Japan



*Speaker

1. Concept and requirements

The P-ONE detector [1][2] will consist of clusters of vertical lines, each 1 km long. Optical fiber connections (see figure 1) will have typical lengths between 50 m and 1 km on the line, and up to several kilometer between mJBs and clusters. Precise timing information has to be provided to all P-OMs to allow reconstruction of physics events in the detector volume. As the P-ONE detector will be constructed in steps, a scalable synchronization system with sub ns resolution [3] is required. In contrast to other systems [4], the BlackCat system we developed provides a central clock distribution, broadcast timing pulses, and precise online link delay measurements to provide the same functionality without additional external PLLs¹. BlackCat allows a direct measurement of optical fiber length in situ, as well as constant monitoring of changes even with data taking going on.

Going for a fiber based system has the potential to reduce the amount of necessary copper strands in the backbone cable. This in turn reflects on the weight in air and water of the detector line, and provides more flexibility in the component design of the line. Specifically, the buoy and anchor can be smaller dimensioned and transportation efforts and costs reduced.

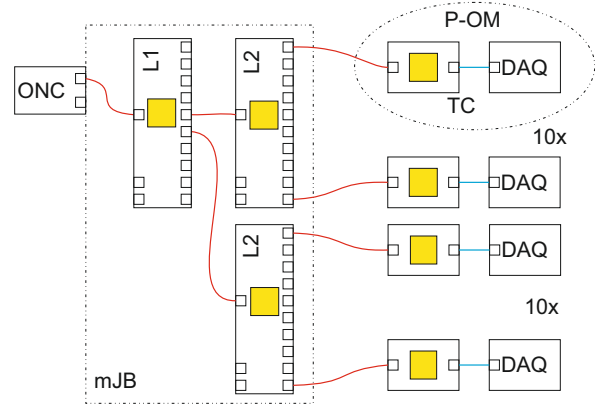


Figure 1: One line of P-ONE, with mJB and optical switch, as well as P-OMs. Each yellow box marks one FPGA containing a BlackCat endpoint.

2. The BlackCat system

BlackCat is a VHDL² code set developed at TUM which implements standard IEEE 802.3 ethernet [5] inside Lattice FPGAs [6], using SerDes³ blocks for communication. Standard cores are used to implement PHY (Physical Layer Device) and MAC (Media Access Controller) functionalities. Generic and vendor specific VHDL packages allow using BlackCat on different FPGA families without the need to change the BlackCat core. For P-ONE implementation the Lattice ECP3 and ECP5 families were selected due to good experiences with reliability and profound knowledge about the internal structures of the devices as well as the implementation tools. All endpoints can be configured as simple endpoint, as bridge (providing timing and standard GbE to a device), and as Fan-In-Fan-Out (FiFo) providing tree like structures for network and timing. BlackCat endpoints implement DHCP⁴, ICMP⁵ and ARP⁶ in the fabric. A new protocol named "Discovery" identifies all endpoints in a network segment and gives access to basic settings, even if DHCP is not available.

¹Phase Locked Loop – a clock tuning circuit

²Very High Speed Integrated Circuit Hardware Description Language – hardware programming language

³Serializer Deserializer – used for high speed serial links

⁴Dynamic Host Configuration Protocol – assigns an IP address to a device

⁵Internet Control Message Protocol – used for ping command

⁶Address Resolution Protocol – provides a MAC address for an IP number

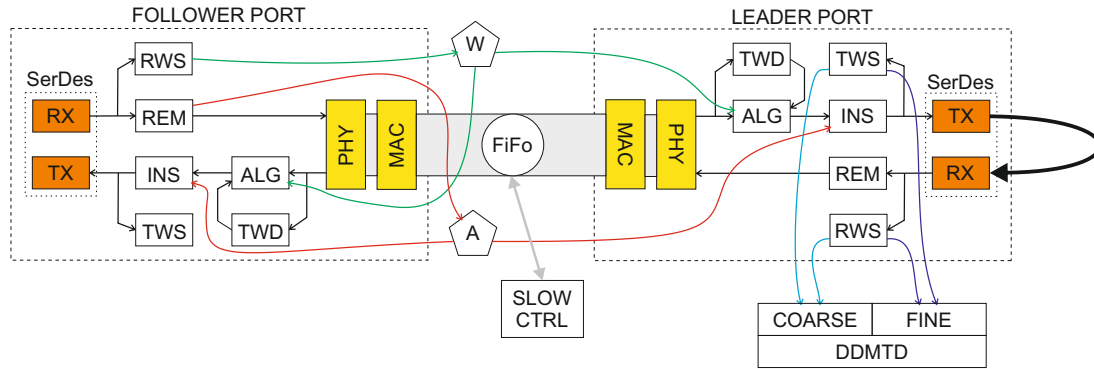


Figure 2: A BlackCat FiFo endpoint (only one leader port shown, loop back enabled) with SlowControl and delay measurement unit (DDMDT), as implemented in one FPGA. Standard IEEE components are marked in yellow, vendor specific ASIC circuits are marked in orange. All others are BlackCat add-ons. Acronyms are being described in detail in the text.

Slow control access is implemented by TRBnet [7] functions, and allows to re-use existing control software without changes. Registers, memories and external SPI⁷ (boot FlashROM) and I²C⁸ (sensors) can be accessed. A forwarder unit allows to send data from the endpoint to a target inside the network. In case of lockups a "ping of death" resets the FPGA logic or reboots the bitstream from external FlashROM.

As requirement for link delay measurements all fiber links in the network need to run with a copy of the central clock, and in addition, need to be aligned to the 16 bit atomic units used by GbE. Clock distribution is handled by BlackCat in a straight-forward way: the reset state machines for SerDes TX (transmit) and RX (receive) channels ensure that the recovered RX clock from a follower port is used as reference clock on all TX SerDes channels (see figure 2). This distributes the central clock to all links on one FPGA.

For 16 bit alignment the 125 MHz data stream on the follower SerDes RX channel is scanned by the RWS (RX Word Sync) unit, which provides a 62.5 MHz word alignment signal (see figure 5). The alignment signal is forwarded to all ALG (Aligner) units which compare to the local word alignment (sensed by the TWS (TX Word Sync) units) and adjust the alignment if necessary.

At this point, all links are 16 bit aligned and run at a syntonous clock with IEEE 802.3 compliant data streams.

3. Test and calibration modes

For calibration, a well defined high precision delay line in the data path is needed. In addition, loop backs in the transmitting and receiving FPGA can be used to measure parts of the signal path independently. An important consideration to choose Lattice ECP3 and ECP5 families is the availability of loop back modes (see figure 3) which can be used to implement a local loop back in the leader port, and a remote loop back in the follower port. Using the loop backs allows to measure

⁷Serial Peripheral Interface – synchronous serial bus

⁸Inter-Integrated Circuit – asynchronous two wire bus for ICs

the delay caused by the SFPs⁹ and fiber without the contribution of the FPGA SerDes and fabric. Both loop backs keep the uplinks synchronized, so the timing distribution stays operational despite ongoing measurements. Only the downlinks need to be re-established. In addition, the RX channel




Figure 3: Local (red), remote (blue) loop back, and normal timing path (green). The FPGA fabric is shown as circle. Loop backs are implemented on analog signal level.

Clock Data Recovery (CDR) can be used as high precision adjustable delay line. It locks during the link establishment on the incoming bit stream, and will notify the fabric about the Word Alignment Position (WAP, see figure 4) on which the recovered RX byte clock is aligned on. Using this feature on a leader port RX line, the round trip delay can be artificially delayed in defined steps of

$$\Delta t_{\text{WAP}} = \frac{1}{125 \text{ Mhz} \cdot n_{\text{WAP}}} = 8 \text{ ns}/10 = 800 \text{ ps} \quad (1)$$

without affecting the clock distribution.




Figure 4: A link scan with all ten WAP values shows the delay steps of 800 ps in the fine time measurement. Peaks are Gaussians with $\sigma \approx 65 \text{ ps}$.

4. Link delay measurements

To measure the round trip delay on a link, a signal needs to be sent from a leader port and returned from the follower port, both in real time (see figure 3). This is not possible on IEEE 802.3 compliant endpoints, as an ongoing ethernet frame transmission must not be interrupted in any case.

⁹Small Form-factor Pluggable – industrial standard optical transceivers




Figure 5: Coarse and fine time measurement, outgoing and incoming data streams shown. /11/ mark ethernet IDLE code groups (16 bit), while DLM marks the link delay measurement code group. Coarse time is measured in units of 8 ns in contrast to 16 ns units of fine time.

We have developed a simple mechanism to overcome this limitation. In the leader port TX line, the INS (Inserter, see figure 2) unit can delay the data stream from the PHY by two clock cycles to provide space to insert a DLM (Link delay Measurement) code group without disturbing the data stream. As soon as IDLE code groups are sensed in the data stream (usually during the Inter Packet Gap (IPG)), the delay circuit is reset and one IDLE code group deleted to arm the inserter unit again.

On the RX path, the REM (Remover) unit adds a latency of two clock cycles to the incoming data stream. Incoming DLM code groups are removed by the unit and the delay circuit is reset. Again, the IPG phase in the data stream is used to restore the IEEE 802.3 data stream by inserting an additional IDLE code group. This restores the original IEEE 802.3 compatible data stream and makes DLM code groups invisible for the receiving PHY.

Received DLM code groups are automatically forwarded from the follower RX channel to all TX channels inside the FPGA, which forwards the code groups to the next level of FPGAs, and echoes it also back to the sending FPGA.

Latencies in SerDes and BlackCat logic allow to split the delay measurement (see figure 5) into a coarse time (counted in whole clock cycles) and a fine time (defined by the delay between the outgoing and incoming 16 bit code groups).

4.1 Coarse time measurement

The coarse time is measured by a simple counter, started by the DLM found in the leader port TWS (TX Word Sync) unit, and stopped by recognition of the returning DLM in the RWS (RX Word Sync) unit. Different latencies in the clock networks inside the fabric cause a shift in the coarse time measurement, which can easily be determined by a local loop back measurement on all ten WAP settings.

The major error source in this measurement is the deviation of the central 125 MHz clock oscillator, which is 25 ppm in our setup. An Oven Controlled Oscillator (OXCO) will be used in the P-ONE setup to improve this number to typically 100 ppb.

4.2 Fine time measurement

For the fine time measurement BlackCat uses the 16 bit alignment signals, which are provided on the leader ports by the TWS and RWS units. They provide a 62.5MHz clock signal, which allows using the Digital Dual Mixer Time Difference (DDMTD) [8] method: a free running sampling clock

ν_{DDMTD} is used to sample the two clock signals, providing an interference pattern of both signals, effectively scaling up the signals in time.

The sampling clock needs to be smaller than the clock to be sampled. In general a large N (see formula 2) is selected to gain high resolution. In BlackCat a 62.4 MHz sampling clock is used as best approximation possible with the FPGA internal PLL (yielding in $N = 624$). One clock cycle in the sampling clock domain corresponds to $\Delta t_{\text{min}} = 25.6 \text{ ps}$ in the original clock domain. This Δt_{min} defines the intrinsic time resolution of the measurement.

$$\nu_{\text{DDMTD}} = \left(\frac{N}{N+1} \right) \cdot \nu_n \quad \Delta t_{\text{min}} = \left(\frac{1}{N \cdot \nu_n} \right) = \left(\frac{1}{624 \cdot 62.5 \text{ MHz}} \right) \approx 25.6 \text{ ps} \quad (2)$$

As BlackCat directly measures the delay between the rising edges of both interference patterns, a simple deglitching circuit with edge detector was implemented to clean the rising edge of the interference pattern. Jitter caused by deglitching affects measurements close to a delay of 0 ns and 16 ns, in which case the WAP setting described in section (3) can be used to shift the measurement into a safe region. Using high statistics (2^{18} data points) on the fine time smooths out effects of jitter and results in a $\sigma \approx 65 \text{ ps}$ typically (see figure 4).

5. Verification and results

To ensure reproducible results over boards and FPGA design revisions, local loop back delay was measured on all accessible SerDes channels on two TRB3sc boards and with three different FPGA designs (see figure 6). Selection of clock networks obviously is mandatory, and placement




Figure 6: Local loop back delay measurement. PCS number denotes different SerDes channels. The first design (violet and green data points) had clocks assigned in wrong fabric resources for SerDes channels A2, A3 and D0. The offset between data sheet value ($\Delta t = 152.701 \text{ ns}$) and measured values (light and dark blue, yellow, orange) can be explained by latencies in local and primary clock networks used for measurement.

constraints for the sampling units relative to the SerDes units is needed to achieve constant results. The local loop back measurement can therefore be used to calibrate the delay measurements in situ.

To verify the BlackCat link delay measurement, one fiber link between two TRB3sc modules was used. TX and RX pulses were provided on LVCMOS25 pins of the source FPGA and the delay measured by a GHz oscilloscope and the FPGA internal BlackCat system. Different fiber lengths between 1 m and 1600 m were measured to cover all P-ONE requirements.




Figure 7: Difference of delay measurements taken by BlackCat and oscilloscope. Data points for loop back measurements are shifted for better view. The time difference originates from routing for scope probe pins inside FPGA fabric.

The delays measured by oscilloscope and BlackCat differ by a constant value for normal and loop back modes, independent from fiber length (see figure 7). Error bars were taken from the scope delay measurement statistics, for BlackCat the fine time peak width was used, and for coarse time the clock frequency deviation of 25 ppm was considered.

To verify results, for some fibers the geometrical length was measured, and a simple approximation used for calculating the fiber length from the delay measurement:

$$l_{\text{fiber}} = \frac{(t_{\text{rlb}} - t_{\text{llb}}) - \Delta t_{\text{SFP}}}{2 \cdot v_{\text{fiber}}} = \frac{\Delta t - \Delta t_{\text{SFP}}}{2 \cdot v_{\text{fiber}}} \quad (3)$$

For length calculation (see table 1) a standard value of $v_{\text{fiber}} = 4.897 \text{ ns/m}$ was used, and the integral latency in SFP TX and RX paths was set to $\Delta t_{\text{SFP}} = 8.8 \text{ ns}$. Both oscilloscope and BlackCat measurements could reproduce the geometrically measured fiber lengths in good agreement.

$l_{\text{nom}} [\text{m}]$	$\Delta t_{\text{osc}} [\text{ns}]$	$\Delta t_{\text{BlackCat}} [\text{ns}]$	$l_{\text{measured}} [\text{m}]$	$l_{\text{calc, osc}} [\text{m}]$	$l_{\text{calc, BlackCat}} [\text{m}]$
1	18.76	18.71	1.01	1.02(2)	1.012(12)
2	28.96	28.87	2.07	2.06(2)	2.049(12)
3	38.84	38.90	3.07	3.07(2)	3.067(12)
10	107.60	107.48	10.10	10.09(2)	10.076(12)
50	502.26	502.19	50.44	50.38(2)	50.376(12)

Table 1: Fiber lengths calculated from oscilloscope and BlackCat measurements.

6. Personal acknowledgements

I would like to thank Grzegorz Korcyl (Uniwersytet Jagielloński) for his work on the Gigabit endpoint for TRBnet. His work provides the base for the ethernet part of the BlackCat system.

I want also to thank Joachim Diener for his ongoing support enabling me to continue working on such an exciting project like P-ONE.

References

- [1] Christian Spannfellner "Design of the Pacific Ocean Neutrino Experiment's First Line". In: *Proceedings of the 38th International Cosmic Ray Conference — PoS(ICRC2023)* (2023)
- [2] *P-ONE: Pacific Ocean Neutrino Experiment Website*.
URL: <https://www.pacific-neutrino.org/>
- [3] J.P. Twagirayezu "Performance of the Pacific Ocean Neutrino Experiment (P-ONE)". In: *Proceedings of the 38th International Cosmic Ray Conference — PoS(ICRC2023)* (2023)
- [4] E. Gousiou et al. J. Serrano M. Cattin. "The White Rabbit Project". In: *Proceedings of IBIC2013, Oxford, UK (2013)*. ISBN 978-3-95450-127-4.
- [5] *IEEE Standard for Ethernet*. URL: <https://standards.ieee.org/ieee/802.3/10422/>
- [6] *Lattice Semiconductor Website*. URL: <https://latticesemi.com/>.
- [7] *TRB Collaboration Website*. URL: <http://trb.gsi.de/>.
- [8] Marc R. Feldman, Dustin Tso, Sarina Kapai. "D-DMTD: Digital Dual Mixer Timer Difference". In: *Sandia Report, SAND2017-10097* (2017).

Full Authors List: P-ONE Collaboration

Matteo Agostini¹¹, Nicolai Bailly¹, A.J. Baron¹, Jeannette Bedard¹, Chiara Bellenghi², Michael Böhmer², Cassandra Bosma¹, Dirk Brussow¹, Ken Clark³, Beatrice Crudele¹¹, Matthias Danninger⁴, Fabio De Leo¹, Nathan Deis¹, Tyce DeYoung⁶, Martin Dinkel², Jeanne Garriz⁶, Andreas Gärtner⁵, Roman Gernhäuser², Dilraj Ghuman⁴, Vincent Gousy-Leblanc², Darren Grant⁶, Christian Haack¹⁴, Robert Halliday⁶, Patrick Hatch³, Felix Henningsen⁴, Kilian Holzapfel², Reyna Jenkins¹, Tobias Kerscher², Shane Kerschtnien¹, Konrad Kopański¹⁵, Claudio Kopper¹⁴, Carsten B. Krauss⁵, Ian Kulin¹, Naoko Kurahashi¹², Paul C. W. Lai¹¹, Tim Lavallee¹, Klaus Leismüller², Sally Leys⁸, Ruohan Li², Paweł Malecki¹⁵, Thomas McElroy⁵, Adam Maunder⁵, Jan Michel⁹, Santiago Miro Trejo⁵, Caleb Miller⁴, Nathan Molberg⁵, Roger Moore⁵, Hans Niederhausen⁶, Wojciech Noga¹⁵, Laszlo Papp², Nahee Park³, Meghan Paulson¹, Benoît Pirenne¹, Tom Qiu¹, Elisa Resconi², Niklas Retza², Sergio Rico Agreda¹, Steven Robertson⁵, Albert Ruskey¹, Lisa Schumacher¹⁴, Stephen Sclafani^{12,α}, Christian Spannfellner², Jakub Stacho⁴, Ignacio Taboada¹³, Andrii Terliuk², Matt Tradewell¹, Michael Traxler¹⁰, Chun Fai Tung¹³, Jean Pierre Twagirayezu⁶, Braeden Veenstra⁵, Seann Wagner¹, Christopher Weaver⁶, Nathan Whitehorn⁶, Kinwah Wu¹¹, Juan Pablo Yañez⁵, Shiqi Yu⁶, Yingsong Zheng¹

¹Ocean Networks Canada, University of Victoria, Victoria, British Columbia, Canada.

²Department of Physics, School of Natural Sciences, Technical University of Munich, Garching, Germany.

³Department of Physics, Engineering Physics and Astronomy, Queen's University, Kingston, Ontario, Canada.

⁴Department of Physics, Simon Fraser University, Burnaby, British Columbia, Canada.

⁵Department of Physics, University of Alberta, Edmonton, Alberta, Canada.

⁶Department of Physics and Astronomy, Michigan State University, East Lansing, MI, USA.

⁸Department of Biological Sciences, University of Alberta, Edmonton, Alberta, Canada.

¹⁰Gesellschaft für Schwerionenforschung, Darmstadt, Germany.

¹¹ Department of Physics and Astronomy and Mullard Space Science Laboratory, University College London, United Kingdom

¹² Department of Physics, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104, USA.

¹³ School of Physics and Center for Relativistic Astrophysics, Georgia Institute of Technology, Atlanta, GA, USA.

¹⁴ Erlangen Centre for Astroparticle Physics, Friedrich-Alexander-Universität Erlangen-Nürnberg, D-91058 Erlangen, Germany.

¹⁵ H. Niewodniczański Institute of Nuclear Physics, Polish Academy of Sciences, Radzikowskiego 152, 31-342 Kraków, Poland.

^α now at Department of Physics, University of Maryland, College Park, MD 20742, USA.

Acknowledgments

We thank Ocean Networks Canada for the very successful operation of the NEPTUNE observatory, as well as the support staff from our institutions without whom P-ONE could not be operated efficiently.

We acknowledge the support of Natural Sciences and Engineering Research Council, Canada Foundation for Innovation, Digital Research Alliance, and the Canada First Research Excellence Fund through the Arthur B. McDonald Canadian Astroparticle Physics Research Institute, Canada; European Research Council (ERC), European Union; Deutsche Forschungsgemeinschaft (DFG), Germany; National Science Centre, Poland; U.S. National Science Foundation-Physics Division, USA; Science and Technology Facilities Council, part of U.K. Research and Innovation, and the UCL Cosmoparticle Initiative.