

Publication Year	2019
Acceptance in OA@INAF	2021-02-10T14:38:02Z
Title	The XENON1T data acquisition system
Authors	Aprile, E.; Aalbers, J.; Agostini, F.; Alfonsi, M.; Althueser, L.; et al.
DOI	10.1088/1748-0221/14/07/P07016
Handle	http://hdl.handle.net/20.500.12386/30289
Journal	JOURNAL OF INSTRUMENTATION
Number	14

# The XENON1T Data Acquisition System

E. Aprile,<sup>1</sup> J. Aalbers,<sup>2,3</sup> F. Agostini,<sup>4</sup> M. Alfonsi,<sup>5</sup> L. Althueser,<sup>6</sup> F. D. Amaro,<sup>7</sup> V. C. Antochi,<sup>2</sup> F. Arneodo,<sup>8</sup> D. Barge,<sup>2</sup> L. Baudis,<sup>9</sup> B. Bauermeister,<sup>2</sup> L. Bellagamba,<sup>4</sup> M. L. Benabderrahmane,<sup>8</sup> T. Berger,<sup>10</sup> P. A. Breur,<sup>3</sup> A. Brown,<sup>9</sup> E. Brown,<sup>10</sup> S. Bruenner,<sup>11</sup> G. Bruno,<sup>8</sup> R. Budnik,<sup>12</sup> L. Bütikofer,<sup>13</sup> C. Capelli,<sup>9</sup> J. M. R. Cardoso,<sup>7</sup> D. Cichon,<sup>11</sup> D. Coderre,<sup>13</sup> A. P. Colijn, <sup>3,a</sup> J. Conrad,<sup>2</sup> J. P. Cussonneau,<sup>14</sup> M. P. Decowski,<sup>3</sup> P. de Perio,<sup>1</sup> P. Di Gangi,<sup>4</sup> A. Di Giovanni,<sup>8</sup> S. Diglio,<sup>14</sup> A. Elykov,<sup>13</sup> G. Eurin,<sup>11</sup> J. Fei,<sup>15</sup> A. D. Ferella,<sup>2</sup> A. Fieguth,<sup>6</sup> W. Fulgione,<sup>16,17</sup> P. Gaemers,<sup>3</sup> A. Gallo Rosso,<sup>16</sup> M. Galloway,<sup>9</sup> F. Gao,<sup>1</sup> M. Garbini,<sup>4</sup> L. Grandi,<sup>18</sup> Z. Greene,<sup>1</sup> C. Hasterok,<sup>11</sup> E. Hogenbirk,<sup>3</sup> J. Howlett,<sup>1</sup> M. lacovacci,<sup>19</sup> R. Itay,<sup>12</sup> F. Joerg,<sup>11</sup> S. Kazama,<sup>20</sup> A. Kish,<sup>9</sup> M. Kobayashi,<sup>1</sup> G. Koltman,<sup>12</sup> A. Kopec,<sup>21</sup> H. Landsman,<sup>12</sup> R. F. Lang,<sup>21</sup> L. Levinson,<sup>12</sup> Q. Lin,<sup>1</sup> S. Lindemann,<sup>13</sup> M. Lindner,<sup>11</sup> F. Lombardi,<sup>7,15</sup> J. A. M. Lopes,<sup>7,b</sup> E. López Fune,<sup>22</sup> C. Macolino,<sup>23</sup> J. Mahlstedt,<sup>2</sup> A. Manfredini,<sup>9,12</sup> F. Marignetti,<sup>19</sup> T. Marrodán Undagoitia,<sup>11</sup> J. Masbou,<sup>14</sup> D. Masson,<sup>21</sup> S. Mastroianni,<sup>19</sup> M. Messina,<sup>16,8</sup> K. Micheneau,<sup>14</sup> K. Miller,<sup>18</sup> A. Molinario,<sup>16</sup> K. Morå,<sup>2</sup> Y. Mosbacher,<sup>12</sup> M. Murra,<sup>6</sup> J. Naganoma,<sup>16,24</sup> K. Ni,<sup>15</sup> U. Oberlack,<sup>5</sup> K. Odgers,<sup>10</sup> B. Pelssers,<sup>2</sup> R. Peres,<sup>9,7</sup> F. Piastra,<sup>9</sup> J. Pienaar,<sup>18</sup> V. Pizzella,<sup>11</sup> G. Plante,<sup>1</sup> R. Podviianiuk,<sup>16</sup> H. Qiu,<sup>12</sup> D. Ramírez García,<sup>13</sup> S. Reichard,<sup>9</sup> B. Riedel,<sup>18</sup> A. Rocchetti,<sup>13</sup> N. Rupp,<sup>11</sup> J. M. F. dos Santos,<sup>7</sup> G. Sartorelli,<sup>4</sup> N. Šarčević,<sup>13</sup> M. Scheibelhut,<sup>5</sup> S. Schindler,<sup>5</sup> J. Schreiner,<sup>11</sup> D. Schulte,<sup>6</sup> M. Schumann,<sup>13</sup> L. Scotto Lavina,<sup>22</sup> M. Selvi,<sup>4</sup> P. Shagin,<sup>24</sup> E. Shockley,<sup>18</sup> M. Silva,<sup>7</sup> H. Simgen,<sup>11</sup> C. Therreau,<sup>14</sup> D. Thers,<sup>14</sup> F. Toschi,<sup>13</sup> G. Trinchero,<sup>17</sup> C. D. Tunnell,<sup>24</sup> N. Upole,<sup>18</sup> M. Vargas,<sup>6</sup> G. Volta,<sup>9</sup> O. Wack,<sup>11</sup> H. Wang,<sup>25</sup> Y. Wei,<sup>15</sup> C. Weinheimer,<sup>6</sup> D. Wenz,<sup>5</sup> C. Wittweg,<sup>6</sup> J. Wulf,<sup>9</sup> J. Ye,<sup>15</sup> Y. Zhang,<sup>1</sup> T. Zhu,<sup>1</sup> J. P. Zopounidis<sup>22</sup> for the XENON Collaboration

# and M. Pieracci,<sup>26</sup> C. Tintori<sup>26</sup>

<sup>1</sup>Physics Department, Columbia University, New York, NY 10027, USA

 $^2Oskar {\it Klein Centre, Department of Physics, Stockholm University, AlbaNova, Stockholm SE-10691, Sweden}$ 

<sup>3</sup>Nikhef and the University of Amsterdam, Science Park, 1098XG Amsterdam, Netherlands

<sup>4</sup>Department of Physics and Astronomy, University of Bologna and INFN-Bologna, 40126 Bologna, Italy

- <sup>5</sup>Institut für Physik & Exzellenzcluster PRISMA, Johannes Gutenberg-Universität Mainz, 55099 Mainz, Germany
- <sup>6</sup>Institut für Kernphysik, Westfälische Wilhelms-Universität Münster, 48149 Münster, Germany

<sup>7</sup>LIBPhys, Department of Physics, University of Coimbra, 3004-516 Coimbra, Portugal

- <sup>8</sup>New York University Abu Dhabi, Abu Dhabi, United Arab Emirates
- <sup>9</sup>Physik-Institut, University of Zurich, 8057 Zurich, Switzerland
- <sup>10</sup>Department of Physics, Applied Physics and Astronomy, Rensselaer Polytechnic Institute, Troy, NY 12180, USA
- <sup>11</sup>Max-Planck-Institut für Kernphysik, 69117 Heidelberg, Germany

<sup>12</sup>Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 7610001, Israel

<sup>&</sup>lt;sup>a</sup>Also at Institute for Subatomic Physics, Utrecht University, Utrecht, Netherlands

<sup>&</sup>lt;sup>b</sup>Also at Coimbra Polytechnic - ISEC, Coimbra, Portugal

<sup>13</sup>Physikalisches Institut, Universität Freiburg, 79104 Freiburg, Germany

<sup>14</sup>SUBATECH, IMT Atlantique, CNRS/IN2P3, Université de Nantes, Nantes 44307, France

- <sup>15</sup>Department of Physics, University of California, San Diego, CA 92093, USA
- <sup>16</sup>INFN-Laboratori Nazionali del Gran Sasso and Gran Sasso Science Institute, 67100 L'Aquila, Italy
- <sup>17</sup>INFN-Torino and Osservatorio Astrofisico di Torino, 10125 Torino, Italy
- <sup>18</sup>Department of Physics & Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA
- <sup>19</sup>Department of Physics "Ettore Pancini", University of Napoli and INFN-Napoli, 80126 Napoli, Italy
- <sup>20</sup>Kobayashi-Maskawa Institute for the Origin of Particles and the Universe, Nagoya University, Furo-cho, Chikusa-ku, Nagoya, Aichi 464-8602, Japan
- <sup>21</sup>Department of Physics and Astronomy, Purdue University, West Lafayette, IN 47907, USA
- <sup>22</sup>LPNHE, Université Pierre et Marie Curie, Université Paris Diderot, CNRS/IN2P3, Paris 75252, France
- <sup>23</sup>LAL, Université Paris-Sud, CNRS/IN2P3, Université Paris-Saclay, F-91405 Orsay, France
- <sup>24</sup>Department of Physics and Astronomy, Rice University, Houston, TX 77005, USA
- <sup>25</sup>Physics & Astronomy Department, University of California, Los Angeles, CA 90095, USA

<sup>26</sup>CAEN S.p.A. Viareggio, Lucca 55049, Italy

*E-mail:* daniel.coderre@physik.uni-freiburg.de, xenon@lngs.infn.it

ABSTRACT: The XENON1T liquid xenon time projection chamber is the most sensitive detector built to date for the measurement of direct interactions of weakly interacting massive particles with normal matter. The data acquisition system (DAQ) is constructed from commercial, open source, and custom components to digitize signals from the detector and store them for later analysis. The system achieves an extremely low signal threshold by triggering each channel independently, achieving a single photoelectron acceptance of  $(93 \pm 3)$ %, and deferring the global trigger to a later, software stage. The event identification is based on MongoDB database queries and has over 98% efficiency at recognizing interactions at the analysis threshold in the center of the target. A readout bandwidth over 300 MB/s is reached in calibration modes and is further expandable via parallelization. This DAQ system was successfully used during three years of operation of XENON1T.

KEYWORDS: Dark Matter detectors, Data acquisition concepts, Trigger concepts and systems (hardware and software), Control and monitor systems online

ArXiv ePrint: 1906.00819

# Contents

**O**-----

I	Overview		1
2	Hardware and Readout		3
	2.1	Frontend Electronics	3
	2.2	Readout Software	4
	2.3	Busy and High-Energy Veto Systems	5
	2.4	Acquisition Monitor	7
	2.5	Muon Veto Detector Readout	8
3	Software Event Builder		8
	3.1	MongoDB Buffer Databases	8
	3.2	Trigger Logic and Performance	9
4	Ope	eration and Control	11
	4.1	Run Database	11
	4.2	User Interfaces	12
5	Con	clusion	13

#### **1** Overview

The XENON1T experiment is the largest dual-phase liquid xenon time projection chamber (TPC) built to date. It was designed to measure nuclear recoils induced by dark matter in the form of weakly interacting massive particles (WIMPs). The detector has been in continuous operation at the INFN Laboratori Nazionali del Gran Sasso (LNGS) in Italy from its commissioning in 2016 until December 2018 and has produced the strongest exclusion limits to date on the coupling of WIMPs to matter [1–3]. Particles incident on the 2.0 tonne liquid xenon target produce a prompt scintillation signal (S1) and delayed ionization signal, measured via proportional scintillation (S2), separated by a maximum time delay of 670  $\mu$ s. These signals are measured by two arrays of 127 and 121 Hamamatsu R11410-21 3" photomultiplier tubes (PMTs) [5] installed above and below the liquid xenon target. The energy region of interest for the standard spin-independent WIMP search is [4.9 – 40.9] keV for nuclear recoils, but the bulk of background events, as well as regions of interest for other physics searches [4], are on scales up to a few MeV. The detector is installed in a large tank of deionized water instrumented as an active Cherenkov muon veto in order to reduce background and identify cosmic muons [6]. The XENON1T detector is described in detail in [7].

The data acquisition system (DAQ) is responsible for digitizing the signals from each PMT and identifying and recording every interaction with the target. The large variation in signal intensity (keV up to MeV scale), as well as the two very different signal shapes of the S1 and S2 and the

variable time delay between them, create a unique challenge for the DAQ. The DAQ is designed around a triggerless readout paradigm, where data is acquired constantly without a global trigger and physical interactions are identified later in software ("event building"). This allows for the lowest possible energy threshold, but necessitates custom readout logic to provide independent readout of each detector channel while maintaining full time synchronization. The system is centered around a NoSQL MongoDB database [8], which provides fast storage, sorting, and retrieval of data. A software event builder queries this database asynchronously to identify interactions and write them to files. During normal operation the event builder runs within minutes of acquisition, but can be delayed for hours or days in case of disruptions further down the pipeline. The files are transferred to external sites and processed into a format useful for analysis by the XENON1T computing system [9, 10]. A diagram of the DAQ is shown in Figure 1 and the subsystems are detailed in the following sections.



**Figure 1**. A diagram of the XENON1T data acquisition system showing the main components. Signals from the TPC are digitized and reformatted for insertion into the MongoDB buffer. The event builder scans the buffer for events and writes them to storage. The system is controlled via the run database backend, which is exposed to the user via a web interface. Details on each subsystem are explained in the text.

In dark matter search mode, the event rate identified by the event builder is about 5 Hz and the readout runs with <1% deadtime, which is caused by long, periodic electrical noise emitted intermittently by individual channels. Nuclear recoil calibrations, which measure the detector response to a WIMP-like signal, are performed with an  $^{241}$ AmBe source or a deuterium-deuterium neutron generator [11] placed in the water tank outside the cryostat and produce event rates of 10 Hz and 25 Hz, respectively. Other calibrations utilize radioactive sources mixed into the xenon target itself. <sup>83m</sup>Kr is injected once every few weeks for detector stability monitoring [12]. <sup>220</sup>Rn is used to measure the response to background-like low energy electronic recoil calibrations have maximum event rates of 200 Hz and 100 Hz, respectively. Because XENON1T is a rare event search, the full digitized waveform of every event must be recorded for later scrutiny, necessitating a high throughput capability in calibration modes. In addition to these calibrations using radioactive sources, the PMT gains are calibrated with pulsed LEDs configured to provide signals at the single photoelectron level [5]. Table 1 shows an overview of the various operating modes as well as their maximum event rates and average event sizes.

**Table 1.** Event rates and sizes for various operating modes in XENON1T. For the internal sources  $^{220}$ Rn and  $^{83m}$ Kr the event rate is highest right after the source is injected and decreases to the background rate over the course of several hours, at which point either another injection is performed or the calibration campaign is ended. The numbers with an asterisk indicate the event size after application of the hardware high energy veto, described in Section 2.1.

Operation Mode	Maximum Event Rate (Hz)	Average Event Size (MB)
Dark matter search mode	6	1.68
<sup>220</sup> Rn	100	1.66 (0.77*)
<sup>83m</sup> Kr	200	0.33
<sup>241</sup> AmBe	11	2.19
Neutron Generator	30	2.84 (1.02*)
LED	450	0.30

## 2 Hardware and Readout

The readout portion of the DAQ encompasses all components up to the MongoDB buffer in Figure 1. This stage is responsible for detecting signals from the TPC and putting them into a digital format suitable for software triggering.

#### 2.1 Frontend Electronics

The electronics of XENON1T are housed in a counting room within a service building adjacent to the water tank containing the detector [7]. PMT high voltage is controlled via the XENON1T slow control system [14] and supplied by CAEN A1535N modules connected via custom-made passive low-pass filter boxes with a 2.8 MHz cutoff frequency to reduce electronic noise. The filter boxes reduce the RMS of the baseline from a mean of  $(0.91 \pm 0.19)$  mV to  $(0.31 \pm 0.06)$  mV, with the remaining dominant 25 kHz noise frequency originating from the high voltage power supplies themselves. Photons incident on the PMT photocathode are detected with an average quantum efficiency of 34.5% and amplified by a channel-dependent gain of  $(1.0 - 5.0) \times 10^6$  [5]. The signals are routed by 12 m PTFE-coated RG-196 coaxial cables [5] to the counting room, passing through potted vacuum feedthroughs (RHSeals), and connected to Phillips model 776 amplifiers [15]. The amplifiers provide a fixed DC-coupled, non-inverting voltage gain of 10. One of the two outputs is connected to to one of eight input channels on a CAEN model V1724 digitizer module in order to record signals from each PMT [16]. 31 of these modules are operated synchronously to record data from the TPC and one additional module records diagnostic data from the DAQ itself. One V1724 provides a 50 MHz clock synchronization signal that is propagated to all of the others such that the clock drift is constant across all digitizers. Acquisition start is synchronized with a NIM logic signal sent from a CAEN V2718 crate controller in order to start the clock counters on all digitizers at the same instant. The V1724s digitize the input signals at a rate of 100 MHz with 2.25 V input range, 40 MHz input bandwidth, and 14 bit resolution. Data is read out from the digitizers by server-grade PCs via 7 optical links connected to CAEN A3818 PCIe cards with either 4 or 8 daisy-chained digitizers per link. Each optical link supports a readout bandwidth of up to 90 MB/s. By parallelizing readout over multiple optical links and PCs, a maximum readout rate of 300 MB/s is achieved for the complete system, which is sufficient to record the 100 Hz target rate for the <sup>220</sup>Rn calibration mode. The maximum readout bandwidth is defined by the saturation point of the first optical link to reach its peak bandwidth of 90 MB/s, since beyond that point the DAQ will be capable of recording only partial events. Above about 100 Hz in <sup>220</sup>Rn data, events begin to overlap due to long periods of activity following large S2 signals, during which many single electrons are extracted into the gas region [17].

Except during LED calibration measurements, which utilize an external periodic NIM trigger provided by a CAEN V2718 crate controller module, digitizers are operated using a custom selftriggering firmware developed together with CAEN as a modification of their DPP-DAW firmware. It implements an independent threshold trigger on each channel ("self-trigger"), which stores and reads out the sampled PMT waveform of a single channel any time a signal on that channel exceeds a configurable threshold. The empty baseline data outside of the self-trigger window is not recorded. The firmware allows user configuration of the self-trigger threshold and the pre- and post-trigger acquisition windows. These windows are each set to 50 samples, corresponding to 500 ns, in order to provide enough empty baseline around each pulse for accurate fitting and subtraction. In case signals exceed the configured acquisition length, the window is automatically extended until the signal goes under threshold again for a user-configurable number of samples, allowing deadtimeless acquisition of both short, small S1 pulses and much larger and longer S2 pulses. This behavior is demonstrated in Figure 2, which shows two self-triggered pulses. The pulse on the left axis is a single photoelectron (PE) and is the shortest length pulse possible (1  $\mu$ s), while the pulse on the right is a larger S2 that is recorded in its entirety by dynamically extending the acquisition window, in this case to more than 8  $\mu$ s. The maximum self-trigger length is defined by the size of the digitizers' internal memory and is over 3 ms. This is orders of magnitude longer than the longest pulses found in normal readout conditions during a XENON1T science run. The self-trigger threshold was set to 15 ADC counts, or 2.06 mV for the majority of channels, but is raised to a maximum of 30 ADC counts for 30 channels exhibiting abnormally high self-trigger rates due to periodic electronic noise from the high voltage power supplies or abnormal PMT behavior. The channel-dependent single PE acceptance averages to  $(93 \pm 3)\%$  [1].

#### 2.2 Readout Software

Readout PCs run the kodiaq software package [18], which reformats data from the binary CAEN format to MongoDB documents. Data is transferred from the digitizers to the readout PCs via block transfers mediated by the CAENVMElib C++ library. The readout software constantly polls the boards in round-robin fashion to read out any available new data. Upon completion of a block transfer, the transferred data is parsed into individual pulses. A *pulse* refers to a single self-trigger from a single channel, which, together with its identifying meta-data, constitutes a single entry in the buffer database. During acquisition one software thread is dedicated to constantly polling the boards for new data in order to maximize readout bandwidth, while several other threads work to reformat the data as it arrives. MongoDB is a NoSQL database which organizes data into collections, analogous to tables in an SQL database. Single database entries in MongoDB are schema-less and are called documents (similar to SQL rows). Each document written by kodiaq contains the timestamp since run start in units of nanoseconds with 10 ns resolution, digitizer module number, channel number, compressed waveform samples as a binary field, and a MongoDB ObjectID field



**Figure 2**. Example self-trigger pulses from one PMT connected to a V1724 ADC. The left pulse is a single photoelectron (PE) and the right pulse is a long S2 candidate. The configured baseline (red) and self-trigger threshold (blue) are shown as dotted lines. The y-axis is zoomed to show detail, but the dynamic range of the digitizer extends to zero on this scale. The ADC scale is inverted, with the baseline at 16000 and the maximum (minimum) sample amplitude at zero (16383), where this range corresponds to the 2.25 V dynamic range of the digitizer.

as a unique identifier. Kodiaq clients on different physical hosts (readout PCs) are controlled by a single *dispatcher* software module, which interfaces with the user interface described later. The dispatcher monitors the state of the readout nodes, collects diagnostic information, sends commands, and creates and indexes collections in the buffer databases, which store data that is waiting to be triggered. In case one or more readout processes suffer a malfunction, the dispatcher stops the entire DAQ and propagates the error message to the user interface as well as to the XENON1T slow control system, where an alarm may be raised.

The dispatcher is also responsible for interfacing with the runs database (see figure 1) and creating metadata records for each run (see section 4.1). The dispatcher/slave architecture facilitates parallelization across several physical readout servers, enabling the system to reach a peak readout bandwidth of 300 MB/s. Communication between the dispatcher and slaves is performed via a custom protocol running over network sockets within the DAQ network. There is no hard limit on the level of parallelization possible and the system could in principle easily be extended by adding more hardware, with the eventual limit coming at the event building stage where all the independent readout pipelines must be combined by one process (see Section 3).

#### 2.3 Busy and High-Energy Veto Systems

The CAEN V1724 digitizers have 8 MB onboard dual-ported memory per channel and will populate this buffer until the data are read by the readout software, at which point the memory is made available again. If the incoming data fills the buffer faster than it can be cleared, this memory can fill completely and the digitizer stops storing new events. In this case a *busy* signal is raised by the digitizer on its front panel LVDS interface. The busy signals from all digitizers are propagated to a CAEN V1495 general purpose module where a custom FPGA firmware is configured to send a NIM logic signal in case any of the 31 V1724 data-transferring boards report a memory full condition.

The NIM signal is fanned to an input on each board and the presence of this signal inhibits the board from recording new data. Therefore, if one board goes into the busy state, the data stream for all boards is effectively blocked until this state is lifted. This busy state is monitored via an additional digitizer, the *acquisition monitor*, described in Section 2.4. When the DAQ is operating using a hardware trigger, for example during LED calibrations and noise monitoring, the trigger NIM logic signal is passed through the V1495 module and inhibited in case any single board reports a busy state. This scheme ensures that for every recorded event, all data from all active channels is available. Partial events during which a busy condition was raised are removed from the later analysis using data from the acquisition monitor.

During neutron generator and <sup>220</sup>Rn calibration modes, the size of the data is greatly increased by high energy events outside the region of interest for low energy calibrations. In the case of <sup>220</sup>Rn these are high energy  $\alpha$ -decay events and in the case of neutron generator data these are typically  $\gamma$ -decays following neutron capture processes. These high energy events are followed by long tails of delayed single electrons leading to a very high storage and bandwidth footprint and can cause unacceptably high dead-times, despite not being useful for most analyses. The (compressed) storage size in bytes of a XENON1T event is S2 dependent and is determined by data to be approximately

Event Size = 
$$10 \text{ kB} + 4 \frac{\text{B}}{\text{PE}} \times \text{S2}$$
 (2.1)

Here S2 represents the size of the S2 in PE. The constant term can be interpreted as the storage size required for digitizing an S1 and S2 plus the storage overhead, while the S2-dependent term represents the contribution of the post-S2 tail. There are multiple orders of magnitude in size difference between a typical low energy event (S2 ~  $10^2 - 10^3$  PE) and a typical  $\alpha$ -decay event in <sup>220</sup>Rn calibration data (S2 >  $10^5$  PE). In order to suppress these high energy events and preserve readout bandwidth for pertinent low-energy data, a veto system was developed based on a Skutek digital pulse processor DDC-10. It uses a BlackVME S6 motherboard hosting a Blackfin BF561 processor, programmable Spartan-6 LX150 FPGA, and a 10 channel 14-bit ADC card sampling at 100 MHz [19]. The input to this unit is the analog sum of the bottom array PMT channels, which is created using the second output of the Phillips 776 amplifiers and a series of Phillips 740 linear fan-in modules, the last two of which are modified with 30 dB attenuators to adjust the signal to be within the 2 V dynamic range of the DDC-10 digitizer input. A custom FPGA firmware on the high energy veto module constantly scans the input signal and searches for pulses with shapes compatible with S2s and integrated areas above a user-configurable threshold. This is set to approximately  $1.5 \times 10^5$  PE, corresponding to several hundred keV. When a signal meeting these requirements is found, a NIM logic signal with a fixed length of 10 ms is generated and propagated to the digitizers using the same fan array used by the busy veto in order to inhibit acquisition of both the high-energy S2 and the most intense component of the subsequent single-electron tail. The storage size for one hour of <sup>220</sup>Rn calibration data is about 150 GB without or 35 GB with the high energy veto active with no impact on the low energy calibration signals. Figure 3 illustrates the functionality of the high energy veto in an example event.

Both the high-energy and busy veto systems take advantage of another feature of the XENON custom CAEN V1724 digitizer firmware where data is not processed by the internal logic immediately but delayed at the input stage in a FIFO by a user configurable time of up to  $10.22 \ \mu s$ 

(XENON1T uses this maximum value). Thus, a high energy S2 inducing a high-energy veto signal will be removed from the data stream by the veto signal that it prompted. This can be seen in Figure 3, in which the veto signal is received in the data stream (red line) before it is actually issued in real time (purple line).



**Figure 3**. An example event from <sup>220</sup>Rn calibration removed by the high energy veto. The top panel shows the attenuated analog sum waveform from the bottom PMT array as recorded by the acquisition monitor and also used as the input to the DDC-10 high energy veto module. The large S2 fulfills the veto condition and a veto signal is issued after processing this peak (purple line). Due to the internal delay within the V1724, the data is vetoed already before the peak happens (red line), effectively removing it. The bottom panel is the digital sum waveform computed from the full-resolution recorded data, in which the S2 is no longer present. Note that the y-axis spans five orders of magnitude.

#### 2.4 Acquisition Monitor

In addition to the 31 V1724 modules used to digitize the signals from the TPC PMTs, one additional V1724 module, the acquisition monitor, is used to digitize information about the DAQ system itself. This module is not connected to the busy module and is fed with a low enough rate of signals to ensure that it is never busy. In order to allow each channel to function independently, this digitizer is also equipped with the XENON custom firmware and operated in self-triggering mode. One channel records the attenuated analog sum of the bottom-array PMTs that is used for the high-energy veto input. If the DAQ does go busy due to high rate or a high-energy veto signal there is still a reduced level of information recorded from the detector during this time, as can be seen when comparing the top (sum signal recorded by the acquisition monitor) versus bottom (full resolution data) panels of Figure 3. The other channels record the busy on/off, the high energy veto on/off, and the trigger from the muon veto system, which is used as an additional synchronization for the two systems (see section 2.5). The acquisition monitor data is stored in the main data stream and is available to analysts in high level data. The first level of selection conditions in the XENON1T analysis, which remove events where the busy, high-energy, or muon veto systems were active, relies on this data [20].

#### 2.5 Muon Veto Detector Readout

The water-Cherenkov muon veto detector surrounding the TPC consists of 84 Hamamatsu R5912-HQE PMTs, read out with a second DAQ system that uses a coincidence trigger module to induce acquisition in 11 CAEN V1724 digitizers via a hardware trigger [6]. Besides the different triggering paradigm, this system is nearly identical to the TPC system described in detail here and both systems share the same readout software, data flow framework, and analysis software. The data from the two detectors is combined at the analysis stage in order to remove events originating from a muon shower, especially muon-induced nuclear recoils.

The 50 MHz clock synchronization signal used by the TPC digitizers is also propagated to the muon veto modules to keep the clocks for both readout systems synchronized. In addition, a GPS-based absolute timing module is embedded into the muon veto system to provide time synchronization with potential external phenomena, such as galactic supernovae [21]. A receiver module located on the laboratory surface computes the current time according to signals from GPS satellites and is connected by a fiber optic cable to a second, underground module located in the XENON counting room. Each time the muon veto is triggered, this module records and stores the current GPS time. In addition, a NIM logic signal is emitted every 15 seconds, which triggers the GPS module and is simultaneously recorded by the acquisition monitor. To compute the exact, global time for a given event, the DAQ retrieves the GPS time of the most recent muon veto or synchronization trigger and adds the time difference according to the digitizer clock. There is at most one bin (10 ns) of jitter over these few seconds, which is negligible for the purposes of this measurement. The DAQ for the upcoming XENONnT experiment will continue to use this method for time synchronization.

### **3** Software Event Builder

The continuous stream of data produced by the readout must be segmented into individual events for use in XENON1T's data processor PAX [9]. The data in between events, consisting primarily of PMT dark counts and single electrons trailing high-energy events, is discarded to save storage space, except in special run modes for studying low-threshold, high-background processes.

The software trigger and event builder [22] perform this segmenting and pruning of the data stream. They read from the MongoDB buffer, search for events compatible with interactions in the detector, and produce archive files with events that are ready for further processing.

#### 3.1 MongoDB Buffer Databases

The readout software, kodiaq, inserts data into three standalone MongoDB databases configured on three event builder computers. MongoDB's "sharding" feature would have allowed operation of a single multi-host database, but it proved too slow (in various configurations) to maintain real-time insertion and query rates in calibration modes, likely due to overhead in maintenance of the combined index. Instead, the readout is configured such that a roughly equal amount of data is inserted into each of the three standalone databases by assigning one third of the digitizer modules to write to each database. In order to work around some aspects of the MongoDB internal memory management, most notably that the deletion of a document does not automatically free up the corresponding disk space, a new MongoDB collection is created every 21 seconds on each machine and the data load is transfered to the newest collection. When the data in the collection is no longer needed, the entire collection is dropped, which deletes all data and indexes and frees up the corresponding storage space. The 21-second time window corresponds to the wraparound of the 31-bit, 100 MHz V1724 digitizer clock counter. All collections are indexed at initialization by timestamp, module identifier, and channel number, allowing for fast queries at the event builder stage. The trigger logic is able to detect pulses that span two collections by postponing processing of pulses appearing near the end of one collection until the next collection is read.

The database on each machine can use up to 128 GB RAM, 1 TB solid state storage, and several TB hard disk space, allowing the DAQ to record data for days even if the connection to the file storage system is interrupted. During calibration data-taking, data is cleared by dropping each 21-second collection immediately after it is processed by the trigger, which typically happens less than two minutes after acquisition. During dark matter data-taking, pre-trigger data is stored in the buffer databases for six hours before being cleared. An external program monitors the Supernova Early Warning System (SNEWS) mail list and only flags data for purging if there is no supernova email warning [23]. This would enable the study of low-energy signals from supernova neutrinos, which may otherwise be partially below the trigger threshold [24].

# 3.2 Trigger Logic and Performance

The XENON1T software trigger logic is a single thread process included as part of the PAX data processor [9]. It uses a rolling time window to query the MongoDB database and search for time-coincident pulses. Unpacking the full data payload of each pulse and combining this information would require a system with a prohibitively large amount of computing power, especially during high rate calibrations. Therefore, the payload (digitized waveform) of each self-trigger document is never examined, only the timestamp and channel number. These fields are indexed by the database, allowing the trigger to function in real time. The trigger logic is able to effectively identify S1 and S2 signals by recognizing the tight time coincidence (within 100 ns) of S1 signals compared to the broader distribution of S2 signals, which can be several  $\mu$ s long. After classification as S1 or S2, a trigger decision is made based on the number of pulses that contribute to the signal, with the coincidence level set to 50 pulses for S1 and 60 pulses for S2.

The low energy threshold of the XENON1T spin-(in)dependent WIMP search analyses [1–3] are defined by a 3-fold coincidence requirement on the S1 signal. Due to the amplification of the S2 signal, interactions with S1s at this energy feature S2s that are orders of magnitude higher. Thus, it is the S2 trigger channel that is relevant for the low-energy performance of the DAQ. The 60 pulse threshold eliminates triggers from one- or two-electron S2s, which are common following high-energy events, while achieving a high efficiency at the target analysis S2 threshold of 200 PE. Lower-energy signals can still be studied in special run modes where the trigger is disabled. Due to larger instrumental backgrounds at these low energies, only a small amount of such data is needed to get background-limited studies.

Figure 4 shows the trigger efficiency as a function of S2 size, as determined from simulated waveforms (curves) and measured using <sup>220</sup>Rn calibration data events near the cathode (points). Due to large charge losses at the cathode (caused by the field inversion), events originating there can have S2s around the trigger threshold even if their S1s are large enough to induce a trigger.



**Figure 4.** Efficiency of the XENON1T S2 trigger versus observed S2 size in photoelectrons (PE, bottom x-axis) or equivalent number of extracted electrons (top x-axis). The colored lines show the efficiency determined from simulated waveforms, at different depths in the TPC. The dots show the S2 trigger efficiency, measured on S1-triggered events at the cathode in  $^{220}$ Rn calibration data, with statistical error bars. The lower panel shows this same efficiency measurement with the simulation-derived efficiency at the cathode subtracted. The trigger has a rapid activation between 100 and 150 PE and is >98% efficient in the middle of the TPC above the analysis threshold of 200 PE (dashed line).

The S2 trigger efficiency is measured by the fraction of S1-triggered events for which the S2 also triggered. The dependency of the trigger efficiency on the depth of the interaction in the TPC is illustrated in Figure 4 for the top, middle, and bottom of the TPC. At the analysis threshold of 200 PE, the trigger accepts over 98% of valid S2s in the center of the TPC.

A time window of 1 ms is stored around each triggering S1 or S2 signal found above the trigger threshold, which was chosen to exceed the maximum drift time (670  $\mu$ s) by several hundred  $\mu$ s in order to always record some time before and after each event. Overlapping event windows are combined up to a maximum event size of 10 ms, at which point they are artificially truncated. This truncation causes a short amount of dead time, but since it only takes place for events that are over 15 drift lengths long, it is irrelevant for data acquired under normal conditions. The operating principle of the software trigger is presented in Figure 5.

The output of the event builder is a series of start and end times corresponding to event ranges. These are sent to a file-building system, which queries all pulses corresponding to the event time ranges, reformats them to the storage format, and writes out each event. This system consists of  $\sim 20$  worker processes on each of the three PCs that also host the MongoDB buffer databases, controlled via two *RabbitMQ* task queues [25]. The event builder puts event time ranges on a first queue, from which the workers read. The workers put the finished event objects in a second queue.



**Figure 5**. Sum waveform from 100 ms (about 14 full drift times) of XENON1T data at the input stage of the software trigger. S1 (blue) and S2 (red) peaks above the trigger threshold of 50 and 60 coincident signals, respectively, are marked and correspond to the right y-axis. Peaks detected below these coincidence thresholds (gray X markers) do not induce a trigger. The event windows are drawn as orange boxes and would be the only portion of this data written to disk during normal data-taking. The interesting portion of the middle event is shown in the upper panel.

Finally, a single process reads the event objects, then streams them to file in the order in which the events were acquired. At the XENON1T background event rate of  $\sim$ 5 Hz, the output data stream is approximately 30 GB/hour. During calibration, output rates up to a 250 GB/hour are possible for data collected without the high energy veto active.

# 4 Operation and Control

Besides the main processes reading and triggering data, various support systems are used to combine these processes into a working, controllable system and to keep track of the data collected. This section describes the runs database and web interface, shown as yellow boxes in Figure 1.

# 4.1 Run Database

The communication backend of the DAQ, and also of the data processing system not described here (see [10]), is another MongoDB database called the runs database. This is a three-member replica set with the primary at LNGS and two secondaries located at the University of Chicago and the University of Stockholm, which are the main XENON1T computing sites. Locating the replicas at the remote computing sites allows for local read access from computer nodes and seamless operation of many components of the processing and analysis systems independent of the LNGS network.

Data recorded by the DAQ is organized into data sets called *runs*, where each run corresponds to data acquired for a given consecutive range of time (usually one hour) under identical settings. The runs database holds a record of all DAQ, trigger, and data processor settings, important operational conditions, operator comments, and an index of all locations where any type of data is stored for each run of the DAO. Metadata for muon veto runs are also stored in this database in the same way. This index is used within the DAQ system for subsystem communication. The digitizer readout code, kodiaq, makes an entry at run start that it wrote untriggered data to the MongoDB buffer. The event builder registers this and begins triggering the run, while also making an entry that it began writing a data file to the local storage array at LNGS. The transfer system registers this and transfers the written raw data out to the storage and processing sites, and so on. The computing system uses the runs database as an index to track the location of all data associated with a particular run. For example, one run may have raw data written to tape archive, one to two additional copies of raw data in cloud storage, and one or more copies of processed data at analysis sites, all of which is tracked in the run document. Access to this database is also embedded into the XENON1T analysis software and analysts query the run database to select data with certain properties, for example data acquired with a specific calibration source or acquired with certain operating conditions.

## 4.2 User Interfaces

The DAQ system is accessed by users via a web interface with two entry points: the graphical interface, which is written in HTML and JavaScript for human interaction, and the API, which is for programmatic access. The backend is powered by the Django framework [26]. The human user interface allows configuration of the DAQ by experts into pre-defined readout modes that define the configuration for all parts of the system for calibration, dark matter search, and other data-taking. Both the TPC and muon veto subdetectors are controlled over this interface either independently or in synchronous operation. Operators can monitor the data rates, block transfer rates, and memory and CPU usage of each readout node. The software trigger also provides diagnostic information including its status, database buffer storage capacity, current run, and current event rate. A monitoring app allows inspection of the lone and coincident self-trigger rate per channel for any run and a live oscilloscope displays the most recent pulses recorded in any channel at 10 ns sampling rate. Both of these have proven instrumental for immediate remote diagnosis of detector anomalies.

A graphical interface to the runs database allows operators to apply single-word tags to mark exceptional runs (e.g., "test" for test runs not intended to be used in the dark matter exposure), add extended explanatory comments associated with single runs, view the processing status and data locations of all data types associated with a run, and browse and search runs by their properties. The interface also provides general logging functionality, user control, and an organizational platform for operation shifts. Figure 6 shows a screen shot from the control page of the web interface.

The API is implemented with the Django tastypie framework [27] and provides programmatic access to the runs database and DAQ backend. This API is world-accessible and is used by the computing and data transfer system to read and record processing statuses and storage locations for individual runs. The slow control system uses this API to query the DAQ status and monitor the system. Key parameters, including the data transfer rate, are also stored in the slow control databases. Disruptions within the DAQ are also propagated to the slow control system over the

Readout     Option::     Data use     ●	18 Neuros 1 30 Sec. Sec. 1 19 Martin de la constante de la co	Rest 2000	Add Run DAQ: 0N (a Stop mon,mode_syne_triggersentPC won,mode_syne_triggersentPC won,mode_syne_triggersentPC @eso Common 2000 fueuron_generator_hower @eso Common 2000 @eso Common 200	sveren(_shifter)
Trigger Status: Healthy	Triggering run: 181010_0732(22864)	Checked in: 2018.10.10 at 09:52 40% 20%	TPC running 5.67 MB/s	
012 0 50 750 750 1000 Event rate: 5.00 events/5 D9 Monitor: 2013.10.10 at 09.52 Deadtime last: 0.00% total: 2.73%	1250 1500 1700 2000 2250 2500 eb0 eb1 eb2	21%0 3000 3250 3500 0% 4.7.4 GB 7.07 GB 7.25 GB	Funct 151100 0722     Mode: nextRon       Started by: starting_user     Time: 00.20       reader2 (2)     CPU.16%       4 digitizer(s)     RAX 2%       reader3 (2)     CPU.15%       4 digitizer(s)     RAX 3%       reader3 (2)     CPU.15%       4 digitizer(s)     RAX 3%	generator_heveto
Log: Make a new log entry 2018-10-10707:42 alarm Log alarm from	n slow control system	Submit	4 digitizeri(5) RAM: 3% reader6 (2) CPU: 5% 8 digitizer(5) RAM: 1% reader7 (2) CPU: 5% 8 digitizer(5) RAM: 3%	01:1:1334 bib/s Oth::1:2914 MP/s 01:1-0955 bib/s Oth::1:1628 MP/s 01:1:4158 bib/s
2018-10-1010/125 alarm Log alarm Trom 2018-10-101707:14 dispatcher Log message fr 2018-10-101707:14 dispatcher Log message fr 2018-10-101707:14 dispatcher Log message fr 2018-10-101707:14 dispatcher Log message fr	i sow contor system rom dispatcher rom dispatcher rom dispatcher rom dispatcher		Muon     Veto running 0.00 kB/s       Run: 181010_0732     Mode: muon_mode.       Started by: joing     Time: 00:20       reader5 (2)     CPU.0%       11 digitzer(s)     RAM: 10%	sync_triggersentTPC Data: 0 MB/s IUT: 0 Bi/s

**Figure 6**. One of the views on the DAQ user interface allowing for control of the system. Clockwise from top left: (1) transfer rate history for readout nodes, (2) run control panel allowing data-taking to be started and stopped with various settings, current readout status for (3) TPC and (4) muon veto processes, (5) DAQ log messages, (6) software trigger status showing event rate, dead time, and pre-trigger buffer occupancy.

API, and alarms are raised in the case of DAQ errors, loss of connectivity to the API, or if the data rate has been zero for an extended time. The slow control alarm system then alerts operators and experts via email and telephone [14].

# 5 Conclusion

The XENON1T DAQ system takes advantage of the speed and data-structure of NoSQL databases to implement a flexible software trigger. At low data rates such as in dark matter search mode, this allows the trigger decision to be deferred, possibly indefinitely, or performed remotely. The flexible software trigger is made possible by the use of a custom digitizer firmware to stream all data from the detector to software without any reduction. Every photon measured by the PMTs is inserted as a document in the MongoDB buffer database and used to identify events.

Additionally, the use of a modern web application to control the experiment allows a greater degree of operational flexibility and has allowed operators to forego a conventional control center. Touch-screen tablets allow full control of the system and inform operators of disruption, day or night. This system is the precursor to another iteration, currently being designed and commissioned at the XENONnT experiment, which will further expand upon the strengths of the XENON1T DAQ by forgoing the trigger stage completely in favor of fast, online data processing [28].

# References

- [1] E. Aprile et al. (XENON Collaboration), Phys. Rev. Lett. 121 (2018) 111302.
- [2] E. Aprile et al. (XENON Collaboration), Phys. Rev. Lett. 122 (2019) 141301.
- [3] E. Aprile et al. (XENON Collaboration), Phys. Rev. Lett. 122 (2019) 071301.

- [4] E. Aprile et al. (XENON Collaboration), Nature 568 (2019) 532âĂŞ535.
- [5] P. Barrow et al., 2017 JINST 12 P01024.
- [6] E. Aprile et al. (XENON Collaboration), 2014 JINST 9 P11006.
- [7] E. Aprile et al. (XENON Collaboration), Eur. Phys. J. C 77 (2017) 881.
- [8] MongoDB. https://www.mongodb.org/
- [9] XENON Collaboration, (2018, March 12). The pax data processor v6.8.0 (Version 6.8.0). Zenodo. https://doi.org/10.5281/zenodo.1195785
- [10] B. Riedel et al., PEARC '18 Article No. 9 (2018).
- [11] R. F. Lang et al., Nucl. Instrum. Meth. A 879 (2017) 31.
- [12] A. Manalaysay et al., Rev. Sci. Inst. 81 (2010) 073303.
- [13] E. Aprile et al. (XENON Collaboration), Phys. Rev. D 95 (2017) 072008.
- [14] J. Cardoso et al., RT 2016. DOI: 10.1109/RTC.2016.7543178 (2016).
- [15] Phillips PS776 Amplifier. http://www.phillipsscientific.com/pdf/776ds.pdf
- [16] CAEN. Digital Pulse Processing in Nuclear Physics. White Paper WP2081. https://www.caen.it (2011)
- [17] P. Sorensen, K. Kamdin, 2018 JINST 13 02032.
- [18] D. Coderre, C. Grignon, C. Tunnell. XENON1T/kodiaq: Production release (Version 1.1). Zenodo. https://doi.org/10.5281/zenodo.165486 (2016).
- [19] Skutek DDC-10. http://www.skutek.com/ddc10.htm
- [20] E. Aprile et al. (XENON Collaboration), arXiv:1906.04717 (2019).
- [21] M. De Deo et al., 2019 JINST 14 P04001.
- [22] J. Aalbers, Doctoral Thesis, University of Amsterdam (2018).
- [23] A. Habig et al., Nucl. Phys. B 143 (2005) 543.
- [24] R. F. Lang et al., Phys. Rev. D 94 (2016) 103009.
- [25] RabbitMQ, Version 3.6.5 (2016) Retrieved from https://www.rabbitmq.com
- [26] Django, Version 1.8.1 (2015) Retrieved from https://www.djangoproject.org
- [27] Tastypie, Version 0.14.2 (2019) Retrieved from http://www.tastypieapi.org
- [28] E. Aprile et al. (XENON Collaboration), JCAP 04 (2016) 027.