Texas active target (TexAT) detector - part 3: Acquisition and analysis infrastructure

E. Uberseder, G.V. Rogachev, E. Koshchiy, and E. Pollacco[1]
[1]*IRFU, CEA Saclay, Gif-sur-Yvette, France*

The TexAT (Texas Active Target) detector under development at the Cyclotron Institute will contain a Time Projection Chamber (TPC) with a high level of segmentation in the readout plane. Additionally, the tracking volume will be surrounded by 58 quad-segmented Si detectors backed with an additional CsI detector. All together, the TexAT detector consists of 1024 MicroMegas, 232 Silicon, and 58 CsI readouts for a total 1314 channels. Such a large number of channels, if analyzed using traditional electronics, would lead to an extreme cost for the acquisition system. With the increasing use of TPCs in nuclear physics, the problem of large channel numbers has been addressed by the community. A collaboration of CENBG, GANIL, IRFU, NSCL and RIKEN has developed a low-cost high-fidelity data acquisition system, known as GET, for use with highly segmented TPCs [1].

The GET system is built upon individual "AGET" ASIC chips of 64 channels each. These chips provide, on a per channel basis, customizable pre-amplification and shaping of each signal. Additionally, the analog signal is stored over 512 sequential time bins via a switched capacitor array (SCA) with a selectable clock frequency. Four of these chips are loaded onto a single "AsAd" card (see Fig. 1), which samples and

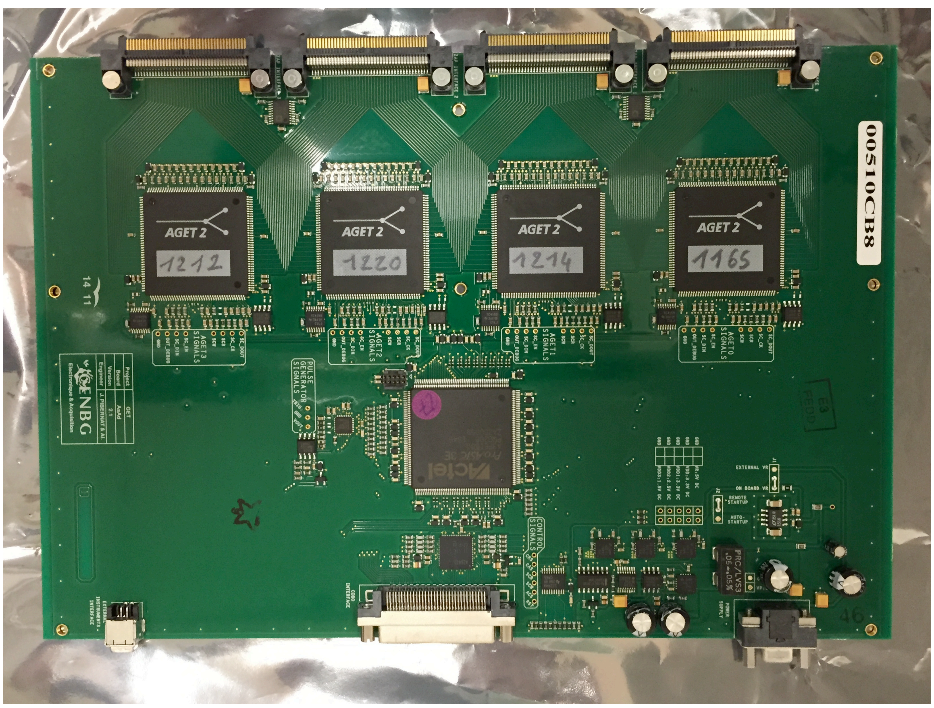

FIG. 1. Image of the AGET chips mounted on an AsAd board.

digitizes all 512 time bins for each of the 256 channels. These "AsAd" boards are read out by a MicroTCA module referred to as "CoBo" (see Fig. 2), which has the ability to package and send the data to a PC from four AsAd boards — a total of 1024 independent channels. The CoBo board has a theoretical throughput of 1 Gbps, and the MicroTCA architecture can transmit 10 Gbps through a fiber uplink, allowing for the operation of 10 CoBo boards, or 10240 channels, within a single crate. An additional MicroTCA board, called

"Mutant", distributes the clocks to the CoBo boards and handles complex custom trigger decisions. The GET system has a specialized set of software tools for control, monitoring, and data offload, and the system has been integrated into the GANIL Narwal-based data acquisition package.

FIG. 2. Image of MicroTCA crate including the two CoBo cards

To date, we have procured five of the six necessary AsAd boards for the TexAT detector, as well as the two needed CoBo boards. The production run for the Mutant boards has not yet taken place, though one has been ordered. The MicroTCA architecture (crate, MCH, power supply, etc.) has also been delivered. The associated software has been preliminarily installed on a staging computer, and data generated via an onboard pulser on the AsAd cards has successfully been collected using the system. To attach the GET system to a detector, a customized "ZAP" board must be manufactured which allows for biasing and reverse-protection of the ASIC chips. These boards are currently under design. As the MicroTCA crate can offload up to 10 Gbps, the associated computing system must be capable of writing such a throughput to disk. The GANIL data acquisition system can operate across several distributed nodes, though also this requires synchronization of the experiment-specific configuration files across each node. Another solution is currently under consideration for the TexAT detector that would incorporate a single multi-processor multi-drive node as the backbone of the acquisition system.

To be performed in a timely manner, the analysis of the large data sets associated with TPC type detectors requires parallel data processing. We have conceived and tested a cluster design unique in nuclear physics, adopted from the data science industry, to fulfill the parallel processing needs for the TexAT detector. The cluster is based on the open-source Hadoop Distributed File System (HDFS) [2], where each compute node contains several terabytes of data storage, and the files are distributed, with replication, across the nodes. This is in contrast to a traditional cluster, where most of the storage is centralized on a single file server and must be pulled to the computed node over the network. The open-source cluster engine used in the present case, Spark [3], is aware of the location of the date file on the cluster, and computations on that file occur on locally on the particular node. As such, very little data needs to travel over network, and only a reduced set of histograms traverse the network to the master node, where they are reduced. The Python interface to the Spark API allows for very seamless integration of physics analyses through the Python hooks into the ROOT [4] analysis package (PyROOT). While the ROOT developers also have a parallel analysis system with data-local capabilities, PROOF, a large advantage of using Spark is found in generality of the cluster engine. The same cluster engine used for data-local analysis can be used to perform parallel batch computations, such as Geant4 calculations, with only a small python wrapper comparable in length to a Condor submit script.

To test the cluster design, we have purchased four data/compute nodes, and a single master node. Each compute node consists of six cores and four 2 Tb hard drives. The total configured space of the distributed file system is 32 Tb, with 24 cores available for computation. The nodes are connected through a 1 Gbps network backbone. The entire configuration has been built inside of a Docker [5] image, such that deployment on a new node takes only minutes after operating system installation. We have benchmarked the system against a data set of 30 GB. A serial analysis consisting of a simple *TSelector* processing a *TChain* on the master node, with no data-locality, took 870 seconds to run. The same analysis performed on 4 nodes with 24 cores and data locality preferred took only 40 seconds, which represents a 90% parallelization efficiency. Additionally, this represents a 6 Gbps data processing rate, a value that demonstrates the power of data locality, as it would ordinarily exceed the maximum possible network transfer rate for only four compute nodes on a 1 Gbps network (4 Gbps).

In summary, the TexAT detector will utilize an entirely new data acquisition system, GET, designed specifically for TPCs. This system allows for a high data throughput at a low cost per channel. We have obtained and tested nearly all of the hardware and software components of the system. To process the large amount of data that will be generated by the detector, a locality-aware parallel cluster design, utilizing software adopted from the data science industry, has been constructed and benchmarked. The parallel data processing efficiency is extremely good, with the data rate exceeding the network bandwidth thanks to the data-local feature of the design.

[1] E. Pollacco *et al.*, Physics Procedia **37**, 1799 (2012).
[2] https://hadoop.apache.org/
[3] https://spark.apache.org/
[4] R. Brun and F. Rademakers, Nucl. Instrum. Methods Phys. Res. **A389**, 81 (1997).
[5] http://www.docker.com/