
Muon New Small Wheel Readout with RCE/ATCA

Version: V4-3

Updated: 11 Nov 2013

Authors: Rainer Bartoldus, Ric Claus, Nicoletta Garelli, Michael Huffer, Su Dong, Charlie Young (SLAC)
Tony Liss, Todd Moore, Ruo-yu Shang (UIUC)
Andy Lankford, Raul Murillo Garcia, Michael Schernau (UC Irvine)
Peter Onyisi (UT Austin)

The ATLAS muon New Small Wheel (NSW) upgrade requires a more performant readout system which is capable of handling the higher data rate with a more granular detector for the high luminosity at phase-1 and beyond. The generic high bandwidth DAQ R&D with the Reconfigurable Cluster Element (RCE) concept on ATCA platform provides a natural candidate solution for this upgrade. This note describes some initial thoughts on the RCE/ATCA implementation for the NSW readout. Some of the key advantages of this readout example:

- A generic DAQ concept with wide range of applications from simple lab testing to full DAQ systems; and from ATLAS upgrade of different subsystems to other experiments. The broader involvements provide synergy and extra resources benefiting each participating project. Knowledge gained through this project has likely reuse values for other upgrade projects and beyond.
- An integrated hardware and software approach with extensive utilities and well defined interface to empower a wider range of users for fast progress with low entrance threshold. The processor-centric RCE concept has enabled most user calibration/tests to be programmed in software to allow fruitful contributions from a broader community through a systematically supported software infrastructure.
- It provides a full set of resources for system dependent user applications with direct access to the frontend links for error corrections and convenient calibration and testing.
- A versatile concept that gives flexibility besides the obvious dramatic bandwidth capacity increase to deliver compact upgrade systems and allow smooth by stage deployment with forward and backward compatibility. Interesting new architectures can be explored with same hardware in situ. In particular, with the full TTC state information available, advanced architectural evolution such as dynamic event routing can be fully supported.
- Mature prototype hardware and software utilities already in place, and well advanced integration with ATLAS TDAQ software infrastructure. Extensive applications in pixel IBL project, and target technology for the urgent muon CSC readout replacement with close resemblance to the NSW application. The initial configuration is similar to the present ROD with direct front end control access still at hands of subsystem users backed by extensive user application freedom to ease detector calibration and tests.
- Flexible and compact test stand solutions with same programming infrastructure.

The RCE Concept and ATCA

By the time of restart in 2014 after phase-0 shutdown, the readout electronics in ATLAS will be already ~10 years old. Rebuilding the RODs are already becoming impractical now just from the inaccessibility to obsolete components. Even at the time of designing the present ATLAS RODs, it was already known that the old technology of VME backplane was only useful to support timing/trigger distributions and slow configuration/monitoring data, while the real high rate DAQ already had to use other means of data flow such as S-Link so that the VME dependence is only peripheral. With rapid growth of the telecommunication industry in the last two decades, much more powerful technologies are now available to serve the future ATLAS readout needs.

To support future experimental activities, an R&D program on generic high bandwidth DAQ started at SLAC in 2007, to capture the commonality of DAQ systems in a set of generic building blocks, as well as an industry standard packaging solution, aimed to flexibly serve many different applications. This concept has the RCE (*Reconfigurable Cluster Element*) and CI (*Cluster Interconnect*) as its fundamental building blocks. These two components are physically packaged together using ATCA (*Advanced Telecommunication Computing Architecture*). The initial round of generation-I hardware are already widely deployed in Photon Science experiments at the LINAC Coherent Light Source (LCLS) at SLAC, as well applied to ATLAS upgrade needs such as the pixel IBL module tests, stave loading tests at Univ Geneva and stave test readout at SR1. Planning of their use for future experiments, such as LSST, LBNE and ATLAS upgrade are well underway. Initial discussion of possible use of the RCE/ATCA R&D for ATLAS upgrade started with presentations at the ATLAS Upgrade Week in Feb/2009 [2] and ACES Mar/2009 [3]. RCE training workshop was held in 2009 [4] in conjunction with establishing an RCE test stand at CERN [5] for common R&D, and e-Groups for RCE R&D communications. Many subsequent reports can be found in IBL general meetings and ATLAS Upgrade Week tracking readout sessions. The relative ease in applying this concept to the designs of these very different projects has provided significant validation as to the correctness of this approach. The urgent need to upgrade the muon CSC readout upgrade [6] for 2014 has also called upon this solution as the targeting technology. Given the similarity of the application to the CSC readout, it is natural to consider the RCE concept also for the NSW upgrade as a smaller next step to ensure the NSW will have a usable DAQ system to serve the testing and commissioning early enough to make sure the efforts during the limited time window of LS2 shutdown can concentrate on debugging the new detectors and not the DAQ system.

The Reconfigurable-Cluster-Element (RCE) Concept

This generic DAQ concept consists of two types of building blocks:

- A generic computational element. This element must be capable of supporting different models of computation, including arbitrary parallel computing implemented through combinatoric logic or DSP style elements, as well as traditional procedural-based software operating on a CPU. In addition, the element must provide efficient, protocol-agnostic mechanisms to transfer information into and out of the element.
- A mechanism to allow these elements to communicate with each other both hierarchically and peer-to-peer. This communication must be realized through industry standard, commodity protocols. The connectivity between elements must allow low latency, high-bandwidth communication.

In turn, this research has satisfied these needs with the *Reconfigurable Cluster Element* or RCE and the *Cluster Interconnect* or CI.

The RCE is a generic computational building block based on programmable *System-On-Chip* (SOC) technology. The current 3rd generation (“Gen-III”) hardware implemented the RCE on mezzanine cards with the core element of each processing RCE being a XILINX ZYNQ XC7Z045T SOC in conjunction of an associated memory system of 1 GB of DDR-3 housed in commodity SO-DIMM. The emerging new technology trend through the ZYNQ series with the natively built-in cross bar connecting processing elements to memory and I/O is now well aligned with the RCE concept. Each RCE has up to 12 lanes (each lane has 2 differential pairs) of high speed user I/O, up to 10Gb/s for each lane. There is also a 10GE Ethernet port on each RCE which is expected to be upgraded to 40GE in near future. The RCE SOC core hosts three different types of computational resources:

- Dual-core *Cortex A9* ARM processors @900Mhz. Standard GNU tools are available for cross-development. The processor is running under the control of a Real-Time kernel called *RTEMS*. RTEMS is an *Open Source* product, which contains, along with the kernel, POSIX standard interfaces as well as a full TCP/IP protocol stack. Processor can also run LINUX for development and testing.
- 900 Multiple-And-Accumulate (MAC) units. Each MAC is capable of one cycle, 18 x 18 fixed-point multiplication summed into a 48-bit accumulator. MACs may be operated either independently or cascaded together.
- Generic combinatoric logic and high-speed block RAM. The 437K CLB FF is almost 10x of the Virtex-5 FX-70 FPGA targeted for the deprecated Gen-II RCE.

The Cluster Interconnect (CI) consists principally of a 10-Gigabit *Ethernet* switch. The number of ports per switch may vary from eight (8) to twenty-four (24), and the physical interface for each port is *XAUI*, *KFR*, or *KR4*, or *SM-II* (100, 1000 Base-T). Physically, the switch is in a very compact form of an ASIC based on the *Fulcrum Microsystems* (now part of Intel) *FM22xx* or *FM6xxx* chip families. To manage the switch, the CI section also has a controlling RCE.

ATCA

While the RCE is actually a platform independent concept, its implementation is clearly more effective when taking advantage of a modern platform with latest technology. ATCA is a communication and packaging standard developed and maintained by the PCI Industrial Computer Manufacturers Group (PICMG). This specification grew out of the needs of the telecommunication industry for a new generation of “carrier grade” communication equipment. As such, this standard has many features attractive to the HEP community, where “lights-out”, large-scale systems composed of multiple crates and racks are the norm. This specification includes provision for the latest trends in high speed interconnect technologies, as well as a strong emphasis on improved system Reliability, Availability and Serviceability (RAS) to achieve lower

cost of operation. While a detailed discussion of ATCA is well beyond the scope of this note (See [111] for details), these are its most pertinent features:

- A generous board form factor (8U x 280 mm with a 30.38 mm pitch). The form factor also includes a mezzanine standard (AMC or the *Advanced Mezzanine Card*) allowing construction of substrate boards.
- A chassis-packaging standard, which allows for as few as two boards and as many as sixteen boards.
- The inclusion of hot-swap capability.
- Provision for *Rear-Transition-Modules* (RTM). This allows for all external cabling to be confined to the backside of the rack and consequently enables the removal of any board without interruption of the existing cable plant.
- Integrated “shelf” support. Each board can be individually monitored and controlled by a central shelf manager. The shelf manager interacts with external systems using industry standard protocols (for example RMCP, HTTP or SNMP) operating through its Gigabit *Ethernet* interface.
- By default, external power input is specified as low-voltage (48 V) DC. This allows for rack aggregation of power, which helps lowering cost of power distribution for a large-scale system.
- It defines a high speed, protocol-agnostic, *serial* backplane. The backplane does *not* employ a data-bus; rather it provides point-to-point connections between boards. A variety of connection topologies are supported, including dual-star, dual-dual star as well as mesh and replicated mesh.

The ATLAS VME replacement taskforce is now recommending ATCA as the base infrastructure for the many trigger and DAQ systems.

Cluster on Board (COB)

For ATCA board-level integration, a new architecture is now adopted with each front-board in a shelf managing its own, individual cluster. These clusters are themselves clustered together using a full-mesh backplane over an ATCA shelf. Thus, one 14-slot shelf is an integrated set of up to fourteen (14) clusters making an integrated *Ethernet* of up to 112 RCEs. This board is called the *Cluster-On-Board* or COB. Purpose-build software and firmware of the RCEs contained on each board allows the board to be configured differently to serve the various applications. The block diagram of the COB and its RTM is shown in Figure 1. For the discussions to follow, we will conservatively quote the bandwidth for the Gen-III version as used for the nCSC readout for 2014 which should be already adequate for NSW.

The COB would service data from up to 96 serial channels from a detector’s Front-End-Electronics. These data would be processed by up to eight RCEs with each RCE servicing up to 12 bi-directional serial channels of data operating at up to 10 Gigabits/s per channel. Reduction of input data would be implemented by subsystem specific code running on the RCEs. The physical interface for the input data would be provided by the RTM, which can be allowing each tailored to subsystem specific needs while still preserving the commonality of the COB. The back-end of the RTM goes to a common Zone-3 connector on the COB carrying its serial data to the COB and subsequently to the appropriate RCEs. For output, one *Ethernet* port (up to 10 Gigabits/s) from each RCE goes to one port of the Fulcrum switch on the COB motherboard, which is connected to the full mesh backplane fabric (P2) provided by the ATCA shelf. This full mesh backplane is organized such that each COB has a data communication bandwidth of 10 Gigabits/s with any

other COB of its shelf (crate) simultaneously. For a smaller systems requiring only a 6-slot shelf (NSW MM or sTGC alone, or each Endcap can fit in a 6-slot shelf), the backplane mesh is naturally replicated to 2x10 Gigabits/s between any pair of slots.

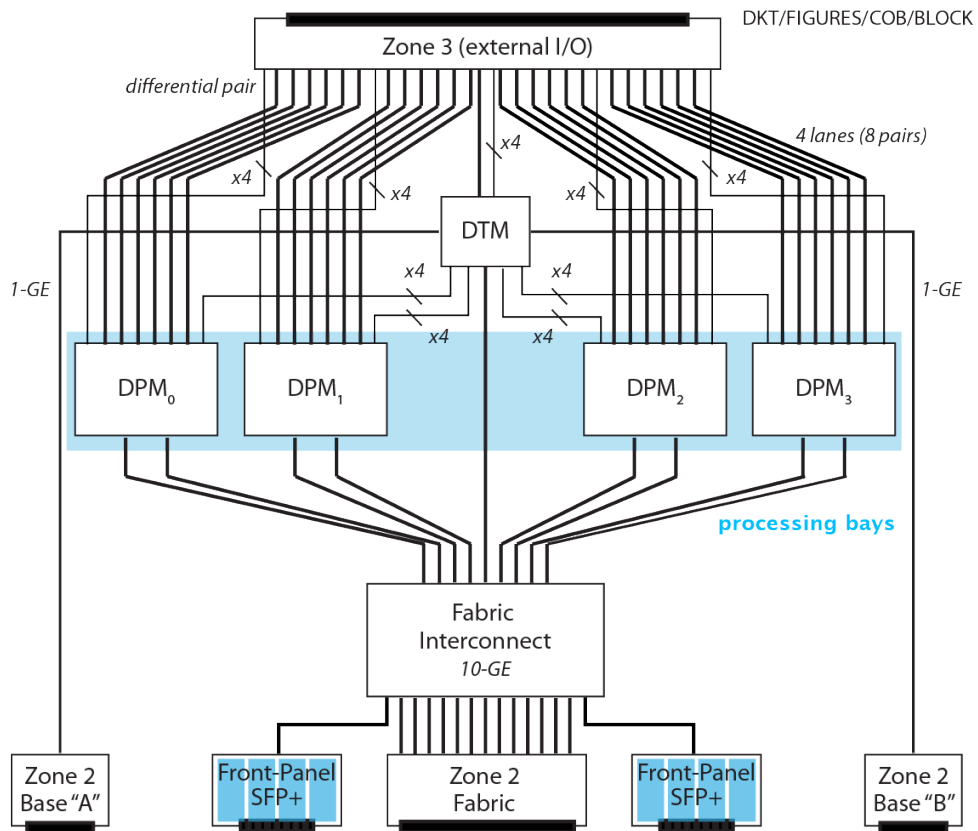


Figure 1: Block diagram of the Cluster-On-Board (COB). Each Data Processing Module (DPM) contains two RCEs. The Data Transportation Module (DTM) contains a single RCE for central control and hosts the IPM controller (IPMC) that manages and monitors the board’s configuration. The Fabric Interconnect is a Fulcrum switch directly mounted on the COB motherboard. The SFP+ front panel ports host 2x10GE independent Ethernet connection for each COB.

As the COB has encompassed most of the resources for common ROD functionalities, the RTMs are typically simpler boards that will carry the detector-dependent interfaces. The simpler demands on the RTMs in turn gives the flexibility to design different versions of the RTMs to serve different applications with less concern of complexity and compatibility as they are easier to replace for future evolution. The SFP+ Ethernet ports are now on the COB front panel to simplify further the RTM to be just 8x12 MPO/SNAP12 user I/O ports. One important real estate saving is that the S-Links will be implemented as a standard protocol plug-in in the RCEs so that they can just use the MPO/SNAP12 I/O ports in a much more compact form than the bulky HOLA cards. This implementation is a simple downgrade from the standard 3.1Gb/s data communication already widely used on the Gen-I RCEs. For forward compatibility with other modern protocols which TDAQ may wish to use for receiving data, building new RTMs for any other new form of interface for a relatively small DAQ system such as NSW is a minor upgrade.

The DTM would also process the L1 timing and control signals sourced from the TTC interface on the RTM or from backplane and fan those signals out to the board’s RCEs. The TTCrx ASIC, widely used in LHC experiments for timing and trigger, is mounted on the TTC interface mezzanine card FTM on the RTM as the external interface to ATLAS TTC system. This will allow each RCE to receive the original standard TTC distribution from the LTP/CTP system, eliminating the need for intermediate interpretation of TTC information, as in, for example, the TIM modules

used by the current VME-based muon and silicon subsystems. The COBs also implemented interconnects to Base fabric as an alternative TTC distribution network so that a COB sitting in a hub slot interfacing to TTC via FTM externally can also distribute/gather TTC to/from the rest of the shelf via the Base fabric. In addition, the DTM RCE with the 1GB memory can also act as an essentially unlimited (in time), standalone trigger and timing emulator.




Figure 2: Picture of the COB V6 prototype loaded with Gen-III pre-Prod DPMs, a prototype DTM.

The NSW project with targeted detector installation in 2017 will demand the readout system to be ready no later than that date to make the best use of the limited LS2 time window to fully commission the new detector. A mature concept of RCE/ATCA has a good chance to deliver a readout system in time for new detector tests during LS2.

NSW Readout Implementation

NSW Detector and Frontend Readout

Each of the two NSW endcaps has 8+8 (large+small) sectors. The longitudinal depth is segmented into 4 quadruplets of 4-layer detector segments. The inner-most and outer-most quadruplets are finely segmented Thin Gap Chambers (sTGC) with fast timing for primarily triggering. The two middle quadruplets are MicroMegas (MM) chambers primarily for precision tracking but also providing trigger information. The DAQ data to be transported upon L1 trigger are envisaged to be carried by GBT links operating 5Gb/s with 3.2Gb/s usable payload bandwidth to allow encoding with e.g. 8b/10b. Separate GBT trigger links are envisioned to serve the data from each bunch crossing for L1 trigger logic. We will not discuss the separate trigger path here but only concentrating on the regular DAQ path after L1 trigger. Both sTGC and MM will use the VMM ASIC for the frontend readout chains which will perform zero suppression to only send sparsified data for DAQ. The configuration/TTC data and DCS data expected to be implemented on the same bidirectional GBT link as the DAQ data. For service modularity concerns, the DAQ data from 4 layers in the same MM multiplet are combined but the sectors are divided into 4 radial stations with separate data fibers, following the chamber station boundaries. The envisaged detector frontend readout mapping to GBT links is illustrated in Figure 4. For each of the MM layers within a station, there are 2048 readout strips of 0.5mm pitch read out by 32 VMM chips with each covering 64 channels. 8 VMM chips are housed by one MMFE board to collectively feed one of the 20 GBT E-link ports at 160Mbps [13]. The two MM modules within each station will have one GBT link for each module. Note that only 16 of the 20 E-links ports are used for the MMFE DAQ data while 4 E-link ports @160Mbps per GBT chip can be used for DCS and TTC data. There are in total 256 GBT links for the MM (2 endcaps; 16 sectors; 4 stations; 2 modules). The sTGC DAQ data are grouped together for two layers in the same quadruplet over a whole sector so that the total NSW DAQ data for sTGC is carried over 128 GBT links (2 endcaps; 8+8 sectors; 4 double layers) [8].

The same local frontend channel maps for all stations at different radii for the MM implies very different real data occupancies on the GBT links between the inner and outer stations, given the known rapid falling hit rate vs radius. A detailed hit rate table [13] was used to examine the actual data rates on the links for different regions of the detector. Given the requirement that the NSW need to survive through phase-2 era, the hit rate used for the bandwidth estimate is extrapolated from 1×10^{34} to 5×10^{34} and another $\times 1.4$ safety factor, with an overall scaling factor on the same hit rate vs radius dependence which should have rather little variation vs. luminosity. The data rates are summarized in Table 1 for the phase-2 L1 trigger rate of 200Khz. It can be seen that the actual data rate is significantly below the 3.2Gbps data payload rate for a GBT link everywhere. The highest rate E-link input to a GBT is also safely below the E-link speed of 160 Mbps. The

Muon NSW Readout Replacement with RCE/ATCA

total data volume for the entire MM system is 96 Gbps. The additional DCS and TTC data returning from the detectors should be <10% of the DAQ data.

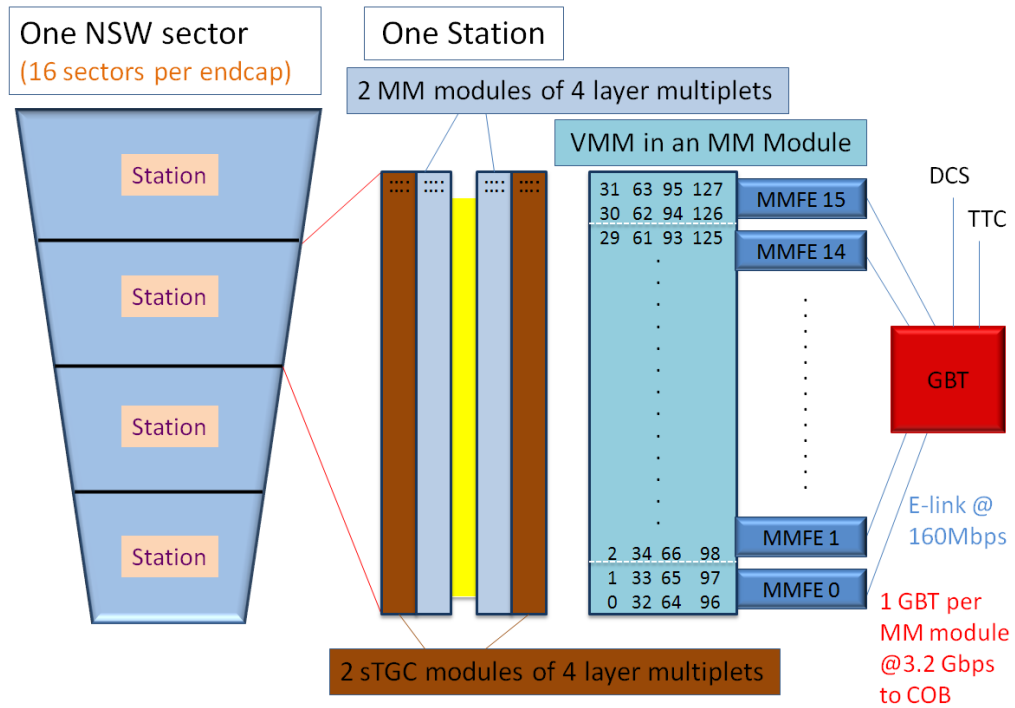


Figure 3: NSW sector layout and MM readout channel mapping.

Station	Inner radius (cm)	Inner edge hit rate (particle/cm ² /s)	Inner edge hit rate (Mhz/chip)	Inner edge VMM2 data rate (Mbps/chip/BCID)	Inner edge strip occupancy	Inner edge E-link rate (Mbps)	GBT data rate (Mbps)
0	100	14100	2.02	1.62	0.40%	65	692
1	202	3140	0.92	0.73	0.18%	29	373
2	305	1320	0.58	0.46	0.11%	18	252
3	407	711	0.42	0.33	0.08%	13	189

Table 1: MM hit and DAQ data rate for 5×10^{34} with x1.4 safety factor and 200 Khz L1 rate for phase-2. The strip occupancy assumed an average of 5 strips per hit particle and each hit only affects 1 BCID. Each hit strip has 32 bits of data. The last two columns of GBT E-link and full GBT data rates assumed reading out 5 BCIDs per event.

NSW Readout Backend Model

The envisaged NSW readout model with RCE/ATCA COBs is illustrated in Figure 4 for the MM system that is still output data to ROSEs using S-links. This readout model is configured in a way that the 4 COBs on the top perform the traditional ROD functionality but they do not generate the S-link output directly. Instead, they use the high-bandwidth ATCA full mesh backplane to transmit the output data to another set of 2 COBs at the lower half which format the data into S-link output to send out over the 96 S-link channels on each COB. The system fits within a 6-slot ATCA crate. Note that each COB performing ROD functionality has ~204Gb/s maximum input for the 64 GBT fibers (out of 96 possible input channels) but only <31 Gb/s real event data expected and the same volume of data needs to exit the RCE to the backplane since the data is already sparsified. Each slot-slot bandwidth is 2x10Gb/s (6-slot shelf has replicated mesh) so that the data out of each "ROD" needs to be divided to go to at least 2 different COBs performing the S-link formatting. sTGC readout channels is half the size of the MM system so that they can also fit easily into one crate.

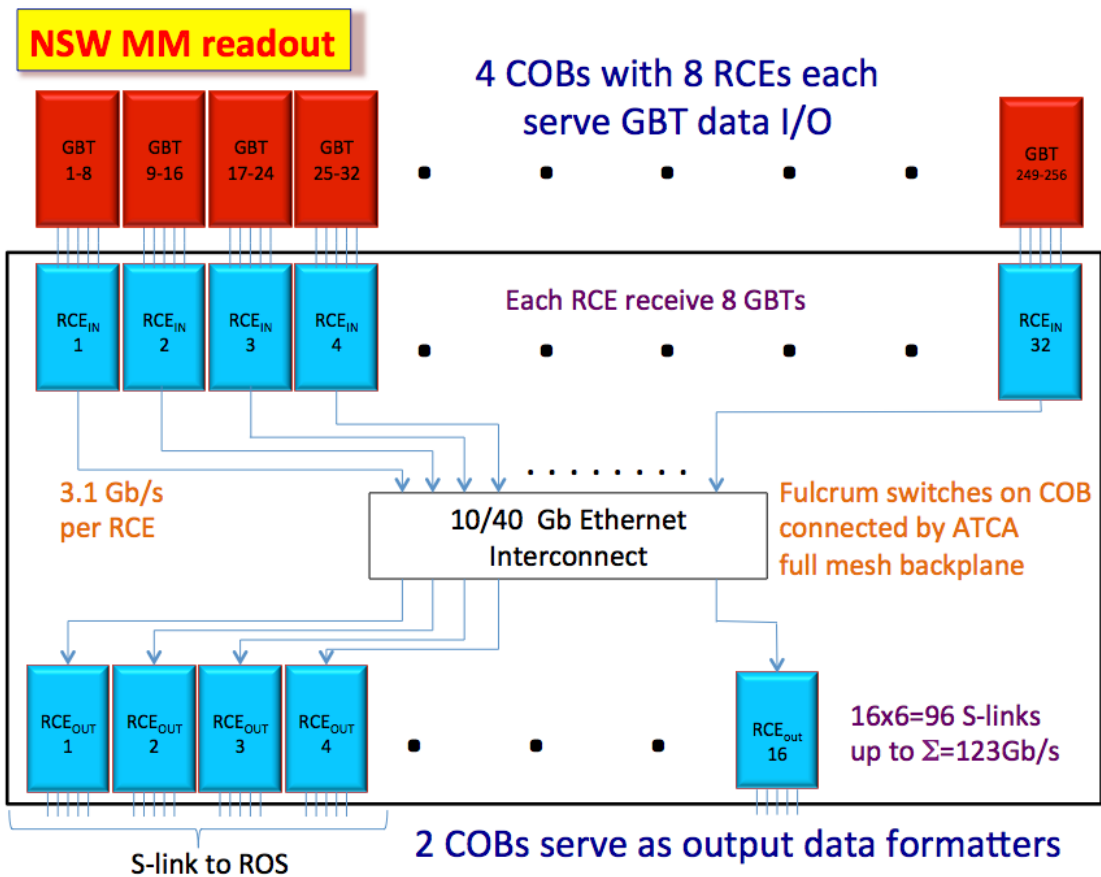


Figure 4: NSW MM readout model with RCE/ATCA COBs and S-link output for ROSEs. The event data rate within each GBT fiber is up to 0.7 Gb/s.

The DCS data arriving together with the DAQ data on GBT link will be separated out from the data stream by the GBT protocol-plugin firmware in the RCE which is continuously running without software configuration as long as there is power in the readout crate. They are gathered by the DTM RCE and put onto Ethernet from the COB front panel ports independent of the rest of the DAQ system. If the whole crate is powered, they can also be aggregated through the backplane to be made available be on a single RCE of a formatter COB dedicated to serve DCS.

Forward/Backward compatibility

Even though there are now many branches of upgrade R&D in ATLAS moved on to ATCA, it is still necessary to assure ourselves that the upgrade readout with ATCA can coexist within the current TDAQ architecture. We will discuss this evolution through the illustration in Figure 5. The intermediate phase-1 scheme corresponds to the model described in Figure 4 for NSW. Comparing the Phase-1 intermediate scheme with the current readout scheme with the ROL-ROS path preserved, one can see that the RCE/ATCA shelf is exactly equivalent to the original VME ROD crate architecturally with identical interfaces to the rest of the experiment. The crate internal communication technology being VME or ATCA is a local protocol mechanics that actually does not bear on the global readout architecture. The only difference for the outside world is the slight difference in configuration data content communicated to the crate via *Ethernet* to the SBC or the COBs individually (each COB has its own *Ethernet* port). The software residing in the SBC of the existing systems using various standard TDAQ inter-process communication utilities was already demonstrated to run on the RCEs in the RTEMS operating system in the IBL calibration test examples since 2010. The minor code extensions for RCE platform compatibility are already in the official TDAQ releases.

Looking into the future, it is clearly beneficial to a common readout design whenever possible. The new GBT links entering the scene at phase-1 is an obvious new application with some complexity that is desirable to have a common interface solution. The RCEs are self-contained versatile GBT transceivers. The variety of resources of ARM processors, FPGA fabric in the SoC and the associated memory are combined to flexibly serve high bandwidth dataflow as well as user applications for DAQ error corrections, feature extraction and calibrations, still with direct access to the frontend link for convenient user application programming. Owing to the generic design of the RCE, it does present a common hardware solution and many common firmware and software utilities so that the adaption for subsystem specific tasks would need only limited effort. A COB can host all the ROD functionalities in the current TDAQ architecture with more resources to support them, especially the software infrastructure with the RCEs.

At the output end, in case the TDAQ interface to new ROSes is already moving to a modern I/O scheme beyond the S-link at phase-1, such new protocol can be hosted within the RCE as another protocol-plugin, just like the PGP (3.2Gb/s) already in operation. In a more significant evolution step, the COBs performing S-link formatting can actually take over the ROS functionalities because the generic design of the RCEs embedded sufficient bandwidth capacity and memory to effectively perform the ROS functionalities. Even for this upgrade scheme with the ROS relocated to be inside a COB to give much simpler physical setup eliminating the bulky ROL S-links, the interface of the whole readout entity of RCE-based “ROD+ROS” is still the same as before with the output appearance just like a ROS. These evolution steps are therefore backward compatible to the current readout architecture at the crate level, to give a completely flexible upgrade path independent of the upgrade timing elsewhere in the experiment. Note also that in the upgrade scheme, the “ROD” (COB) and “ROS” COBs (COB_R) are distinct so that detector systems doing calibrations and TDAQ tests with the “ROS” COBs can go on in parallel without interference.

There is much to be explored towards new directions, such as dynamic event building at the “ROD”-“ROS” transition in this above scheme, afforded by the powerful interconnectivity and high bandwidth of the full mesh ATCA backplane. The dynamic event routing is a recommended requirement for future readout architectures which is entirely natural to the RCE concept. The COBs performing ROD functions can direct the event to any given RCE on the formatting/ROS COBs based on BCID so that a ROS RCE can collect all data channels in the same shelf for the same event together (i.e. event building). For the NSW MM readout case with just one shelf, each event building RCE is collecting the entire NSW MM system. While frontend COBs in the ROD mode has to work for 200kHz L1 rate for Phase-2, such dynamic routing will give each of the 2x8

Muon NSW Readout Replacement with RCE/ATCA

event building RCEs $200\text{KHz}/16=13\text{KHz}$ events to ease the building task. Another recommended requirement to allow dynamically routing separate L1 trigger types can also naturally be met as individual GBT link data for different events are kept separate so that routing by trigger type is not fundamentally different.

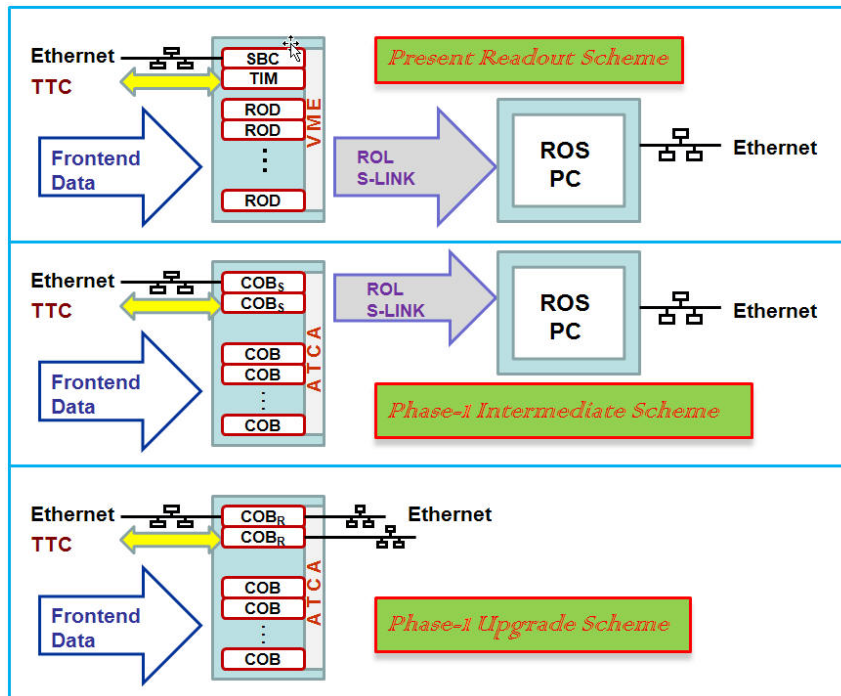


Figure 5: Illustration of a possible path of smooth evolution from the present VME readout to ATCA based schemes with increasing versatility and decreasing complexity.

While the forward/backward compatibility has paved a promising path for time evolution, the wide range of applications afforded by the generic nature of the RCE concept also serves a constant reminder the other dimension of advantage. Applications in at least many tracking and muon upgrade projects are already establishing a pool of resources for development, production and even future operation and maintenance.

In the initial recommendation of the ATLAS readout working group, a future readout architecture scenario is promoted with a dedicated common GBT receiver named FELIX followed by a large network switch while system specific ROD functionalities reside on the other side of the switch with not yet defined hardware platform. There are significant concerns with regarding this architecture. The FELIX concept itself will need some time to define the actual design and demonstrate viability through prototyping. The RCE COBs have flexible enough design to operate as the ROD carrying the system dependent applications on the other side of the switch and doing DAQ data formatting for S-links (or other modern protocols) to ROS, but details of such model will require significant conceptual advance on FELIX to further develop.

Current Activities (as of Nov/2013)

The intense core RCE development for the Gen-III version has converged to a stable set of hardware with both the COB motherboard and mezzanine boards DPM/DTM close to production quality and expected distribution to users by Spring 2014. Most infrastructure and connectivity

Muon NSW Readout Replacement with RCE/ATCA

tests already performed. The core firmware and software tasks that are also well advanced and expected to be completed by Spring 2014 to serve the needs of various projects at Phase-0:

- Core firmware for Gen-III RCE implemented on ZYNQ's native cross-bar.
- S-link protocol plug-in.
- TTC interface to ATLAS.
- DAQ dataflow integrated with ATLAS TDAQ.

Some subsystem dependent applications aimed at 2014 Phase-0 that will bring relevant experience for NSW:

- CSC G-link data and front-end communication protocol plug-in.
- CSC DAQ feature extraction firmware/software.
- CSC calibration implementations.
- Pixel/IBL calibration and DAQ

The NSW data volume is not drastically more than the nCSC case so that the Gen-III design and infrastructure is already sufficient even into the phase-2 era so that the main core development remaining for NSW is the GBT protocol-plugin and associated RTM. This makes the expected 2016 production and ready to run commissioning DAQ by 2017 a realistic goal.

The NSW specific tasks:

- Front-end communication protocol plug-in / TTC integration through GBT.
- DCS data path implementation.
- MM/sTGC DAQ feature extraction.
- MM/sTGC calibration.
- NSW readout chain integration with detector and test beam.

On the application front for NSW in the near term, design is underway at UIUC for a test interface board to house the VMM readout chip and an optical link to communicate with RCEs to allow the development of the RCE protocol plugin for the frontend communication and contributing to the VMM testing. This can already be done with existing Gen-I RCE hardware and the 3.1 Gb/s PGP data link protocol. This will evolve to incorporating the GBT link through a GLIB test set already in hand. The GBT link tests will benefit from a shared effort with the RCE applications to pixel/tracker readout. Once the VMM-RCE link is established, we are hoping to instrument prototype MM chambers for a complete readout chain test against cosmics and test beam. This can again benefit from the existing pixel telescope at SLAC test beam with RCE readout hopefully reach this milestone of the complete readout chain. Some proposed near term milestones:

- Apr/1/2014: Pre-Production COB with Gen-III mezzanine boards and core RCE firmware ready.
- Apr/1/2014: Single VMM test board and RTM link setup ready.
- Nov/1/2014: Basic communications between VMM and RCE/COB through GBT links demonstrated.
- Mar/1/2015: Test beam with prototype MM chambers and full chain of readout electronics prototypes, including the GBT link.

We believe we are on track to deliver a full readout backend system before the 2017 NSW installation to allow the 2018 commissioning to concentrate on the new detector and trigger with an established DAQ.

References

1. *PICMG 3.0* “AdvancedTCA Short From Specification”, January 2003
2. *Presentations at the Feb/2009 ATLAS upgrade week*:
<http://indico.cern.ch/materialDisplay.py?contribId=23&sessionId=20&materialId=slides&confId=45460>;
<http://indico.cern.ch/materialDisplay.py?contribId=9&materialId=slides&confId=52930>
3. *Presentation at ACES (CERN, Mar/2009)*: “High bandwidth generic DAQ research and application for LHC detector upgrade” by Mike Huffer
<http://indico.cern.ch/materialDisplay.py?contribId=51&sessionId=25&materialId=slides&confId=47853>
4. *RCE training workshop*:
<http://indico.cern.ch/conferenceOtherViews.py?view=standard&confId=57836>
5. *RCE Development lab Twiki*: <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/RCEDevelopmentLab>
6. *Muon CSC readout replacement review*:
<https://indico.cern.ch/conferenceDisplay.py?confId=208610>
7. *NSW MM L1 trigger (Venetios Polychronakos), TDAQ Mini-workshop Jul/9/2012*:
<https://indico.cern.ch/contributionDisplay.py?contribId=25&sessionId=1&confId=190420>
8. *NSW sTGC L1 trigger (Lorne Levinson), TDAQ Mini-workshop Jul/9/2012*:
<https://indico.cern.ch/contributionDisplay.py?contribId=24&sessionId=1&confId=190420>
9. *Ken Johns, “MM frontend readout”, NSW readout electronics workshop, Les Houches, Dec/2012*:
<https://indico.cern.ch/contributionDisplay.py?contribId=2&confId=216557>
10. *John Oliver, NSW MM data rate estimate xls for NSW readout electronics workshop, Les Houches, Dec/2012*