

# The ROD Crate DAQ software framework of the ATLAS Data Acquisition System

S. Gameiro, G. Crone, R. Ferrari, D. Francis, B. Gorini, M. Gruwe, M. Joos, G. Lehmann, L. Mapelli, A. Misiejuk, E. Pasqualucci, J. Petersen, R. Spiwojs, L. Tremblet, G. Unel, W. Vandelli, and Y. Yasu

**Abstract**—In the ATLAS experiment at the LHC, the ROD Crate DAQ provides a complete software framework to implement data acquisition functionality at the boundary between the detector specific electronics and the common part of the data acquisition system. Based on a plugin mechanism, it allows selecting and using common services (like data output and data monitoring channels) and developing software to control and acquire data from detector specific modules providing the infrastructure for control, monitoring and calibration. Including also event building functionality, the ROD Crate DAQ is intended to be the main data acquisition tool for the first phase of detector commissioning. This paper presents the design, functionality and performance of the ROD Crate DAQ and its usage in the ATLAS data acquisition system and during detector tests.

**Index Terms**—Data acquisition, detector commissioning, software design, VMEbus.

## I. INTRODUCTION

ATLAS [1] is one of the five experiments at the Large Hadron Collider accelerator at CERN. The data rate from the whole detector after Level 1 trigger rejection is about 150 Gbyte/s which are distributed across 1600 readout links each with a maximum possible throughput of 160 Mbyte/s and a fragment rate of 100 kHz. The ATLAS Trigger and Data Acquisition system (TDAQ) consists of a High Level Trigger (HLT) which performs event selection reducing the event rate by a factor of 300 and a dataflow system that transports data from the detector readout to the HLT system and selected events to mass storage.

Within TDAQ, the ReadOut Driver (ROD) is a sub-detector specific front-end element [2]. It is located, in the event data flow, after the first level of online event selection, between the front-end electronics and the ReadOut System (ROS), as

Manuscript received June 17, 2005; revised December 2, 2005.

S. Gameiro, D. Francis, B. Gorini, M. Gruwe, M. Joos, G. Lehmann, L. Mapelli, J. Petersen, R. Spiwojs and L. Tremblet are with CERN, CH 1211 Geneva 23, Switzerland.

G. Crone is with the Department of Physics & Astronomy, University College London, London WC1E 6BT, U.K.

R. Ferrari and W. Vandelli are with the Università di Pavia and I.N.F.N., Pavia IT-27100, Italy.

A. Misiejuk is with Royal Holloway, Department of Physics, University of London, Surrey TW20 0EX, U.K.

E. Pasqualucci is with the Università di Roma I “La Sapienza” and I.N.F.N., Rome IT-00185, Italy.

G. Unel is with the Department of Physics & Astronomy University of California, Irvine, CA 92697-4575 USA.

Y. Yasu is with the High Energy Accelerator Research Organization, Ibaraki 305-0801, Japan.

Digital Object Identifier 10.1109/TNS.2006.873001

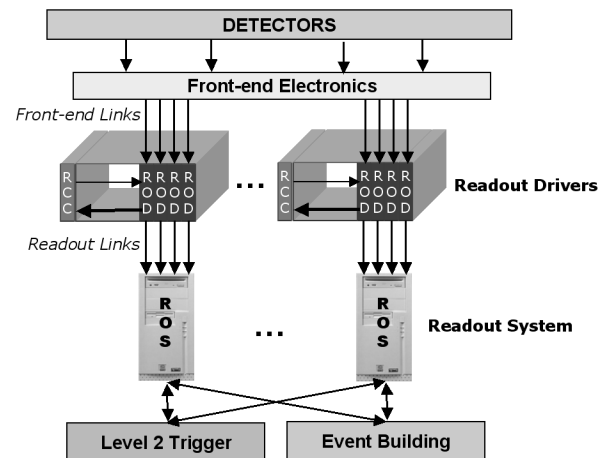


Fig. 1. Simplified diagram of the TDAQ chain.

illustrated in Fig. 1. The ROD receives data from one or more front-end links and sends data over the readout links to the ROS.

The ROD Crate DAQ (RCD) was developed due to the need of subdetectors for common data acquisition functionality at the level of the ROD crate, for single or multiple ROD crates, in laboratory setups, at the assembly of detectors, at test beams, and at the ATLAS experiment during commissioning and production.

The ROD system covers all RODs and other functional elements at the same hierarchical level in the event data flow between the front-end electronics and the ROS; those elements are grouped in crates. The crates contain ROD crate modules which can be: RODs, other modules, e.g., for controlling the front-end electronics or the trigger system, and one or more ROD crate processors. In ATLAS, ROD crates are custom 6U or 9U VME64x [3] crates, each one housing a single board computer running under the Linux operating system (OS), acting as the ROD Crate Controller (RCC).

RCD comprises all the software to operate one or more ROD crates. It provides the functionality for configuration and control, data readout, ROD emulation, monitoring, calibration and event building across multiple ROD crates.

RCD is based on existing ROS software originally developed to meet the requirements of the main dataflow at the level of the ROS: to buffer fragments which are input via standard ATLAS readout links and provide fragments on request from the Level2 trigger and Event Building systems. An analysis of the ROS framework indicated that this could be generalised to the detector specific level of the ROD crates thereby extending the




Fig. 2. ROD emulation use case.

domain of ATLAS DAQ for which common software support is provided.

## II. RCD USE CASES

This chapter describes the RCD use cases [4].

### A. Module Control

The RCD is used to control one or more modules in a ROD crate. Controlling a module consists in communicating with it, passing relevant information at each of the TDAQ system state transitions (see Section III-C). An example of modules that use RCD only for controlling are trigger modules (see Section III-D). All of the other use cases presented in this chapter include the RCD control functionality.

### B. ROD Emulation

In this case, the RCD is used to read non formatted data from one or several modules, and to build ROD fragments, which can be later sent to the TDAQ system. This is typically used when a detector ROD is not available, and allows testing a detector with simpler electronics. Fig. 2 illustrates this use case.

### C. ROD Data Readout

This is the main use case for the first phase of detector commissioning, when the RCD is the DAQ tool available for detector readout. The RCD reads ROD fragments from one or several ROD modules via VMEbus, and builds ROS fragments which can be either stored locally or sent out via a TCP connection to a data driven event builder, as described in Section II-E.

### D. ROD Monitoring

This is the typical use case for the RCD during the ATLAS production phase. The ROD modules are fully controlled by the RCC, but the modules themselves manage the primary data flow, sending events to a ROS through a readout link, as shown in Fig. 3. In this case, the RCD acquires sampled fragments from RODs for monitoring purposes, sending them to the monitoring system or the histogramming service (see Section III-D). The RCD should also be able to detect errors signalled by the ROD modules and perform operational monitoring.

### E. Data Driven Event Building

When acting as data driven event builder (REB), the RCD receives ROS fragments coming from VMEbus crates or ROSs

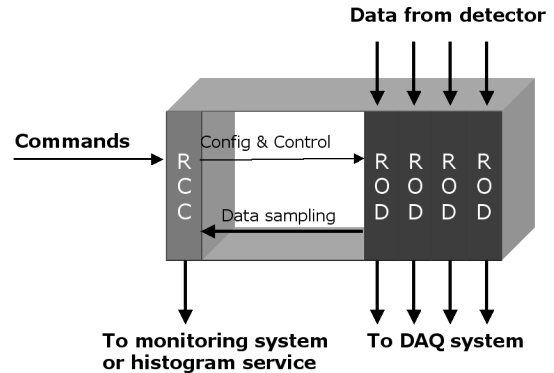


Fig. 3. ROD monitoring use case.

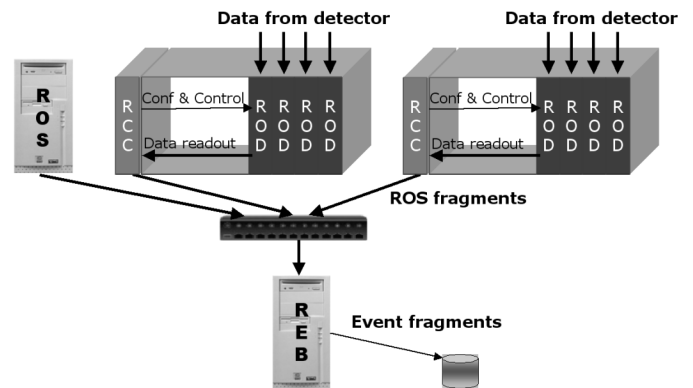


Fig. 4. Data driven event builder use case.

via Ethernet TCP connections, and builds full events (Fig. 4). This is mostly needed during the detector commissioning, when the full TDAQ system and in particular the High Level Trigger is not yet available.

## III. RCD ARCHITECTURE

The RCD software is based on a C++ common application framework originally developed for the ROS and referred to as IOManager (IOM). The framework loads specific plugins to customize its behaviour for the different RCD deployment scenarios.

This chapter describes the RCD architecture, focusing on the thread and plugin structure.

### A. Threads structure

The RCD has the structure of a single, multi-threaded process based on Linux POSIX threads as illustrated in Fig. 5. The trigger thread is activated on the occurrence of an event which may be external or internal to the RCD. It builds a request which describes actions to be performed by the RCD in response to the event, e.g., reading out a number of data sources, and then writes this request to a queue. A request handler thread processes the requests: reads an element from the request queue and executes it. It typically consists in extracting fragments from the data sources, and assembling these into a larger fragment, which is written to an output device. Several request handlers may work in parallel thereby achieving a better CPU utilization.

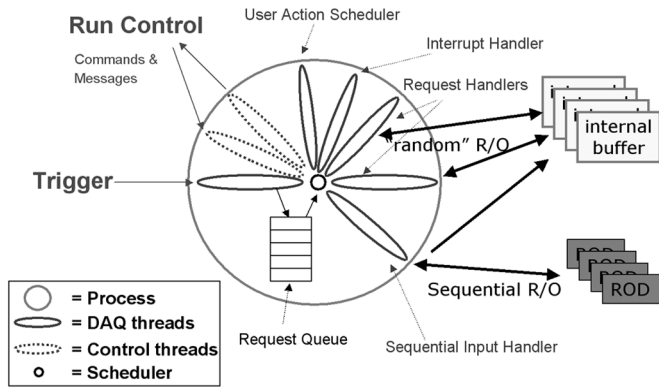


Fig. 5. Thread diagram of the RCD application.

The framework provides a number of optional threads. The *Sequential Input Handler* thread supervises the input of data fragments from several ROD modules, based either on polling of registers in the RODs or VMEbus interrupts. The data fragments are read into a set of internal memory buffers, one per ROD, from where they can be accessed (randomly) via the request handlers. The *User Action Scheduler* thread allows activation of user written handlers when a defined time slice (per handler) has expired. Similarly, when a VMEbus interrupt occurs, the *Interrupt Handler* thread activates a user written handler identified by a VMEbus interrupt vector.

In addition to the threads related to the data flow, there are control threads which communicate with external controller applications for configuration, error handling and statistics.

The strategy for thread execution is based on ‘poll and yield’: a thread polls on a resource, executes until the resource is exhausted and then yields, i.e., gives the CPU back to the OS. This approach minimises the number of thread context switches and the time for one context switch. The RCD is not a hard real time application in the sense that it does not have to react to external events (such as VMEbus interrupts) in a guaranteed time. Rather the design aims at maximising the data throughput.

### B. Plugin Mechanism

The application framework implements the common core functionality of the system: interface with the run control commands, error handling and recovery, activity scheduling, configuration management, and event selection and sampling, among others. The application relies on specific plugins for the implementation of the different I/O protocols.

All the plugins are loaded by the application at run time (dynamic loading) and do not need to be linked to the main application. This allows users to add new plugins without having to modify the application binary.

Fig. 6 shows a diagram with the four different types of plugins: configuration, trigger, module and output. The configuration plugin is the first one to be loaded, and it informs the application about the other plugins. For each of the four types of plugins, an abstract base class defines an API that needs to be implemented for every specific instance. A more detailed description of each plugin is now presented.

1) *Configuration*: The standard way of configuring the RCD is via a configuration database. The database is based on a

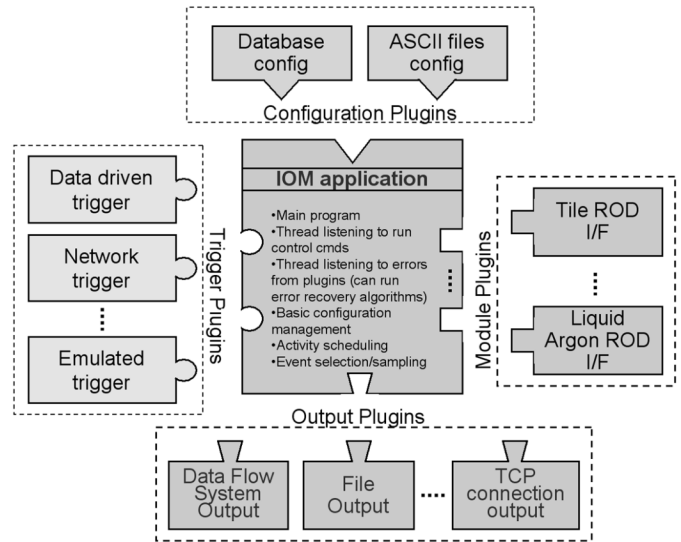


Fig. 6. Plugin diagram of the IOM application.

core schema which includes the definitions of all the classes of the system. A modular structure allows creation of new sub-schemas with user defined classes, which may use other classes from the core schema, and can be included in the user’s database. An alternative way of configuring the RCD is via simple ASCII files, more suitable for development/debugging phases.

2) *Trigger*: The trigger plugins implements the functionality of the trigger thread, see Section III-A. The RCD software provides several trigger plugins characterised by the type of trigger source which may be external or internal to the application: arrival of a message on a network; input of data fragments in the internal buffer via sequential data channels, see Fig. 5 (data driven trigger); internal generation of the trigger (emulation). In addition the user may develop trigger plugins driven by detector specific trigger hardware, e.g., VMEbus modules.

3) *Module*: A module describes a hardware or software component that is controlled by the framework. It can perform any appropriate action associated with a state transition; the base class associated with a module is called *ReadoutModule*, and it includes virtual methods for each of the state transitions. Typical actions include reading the configuration database, initializing the hardware modules or publishing statistics.

A module may be associated with the readout of data. In that case, one or more data channels are defined within the module, with virtual methods for requesting, retrieving and deleting fragments. A specialised implementation of a data channel was provided for RCD deployments, which include virtual methods for polling on a data source and retrieving data (ROD fragments), sequentially.

4) *Output*: As explained in Section III-A, a request handler writes the event fragments to an output device. This functionality is implemented via an output plugin. The RCD software provides several plugins associated with the type of output device, e.g., network (DataFlow System Output) or local disk. In addition the user may develop specific output plugins although this should be an exception.




Fig. 7. Simplified diagram of the controller's State Transitions.

### C. Supervision System and Graphical User Interface

The TDAQ Supervision System is in charge of performing the initialization and shutdown of TDAQ firmware and software, distributing commands to TDAQ elements, synchronizing operations between them, and performing error handling [5]. The building block of the Supervision is the *Controller*. The supervision will generally contain a number of controllers organized in a hierarchical tree.

In order to regulate the control activity, a state machine model was introduced in the controller's core: a controller always has a state, reflecting the possible states of the TDAQ system. The states and transitions of all controllers in the tree are synchronized. Each state has a defined set of authorized transitions that bring the system into a new state, as shown in Fig. 7.

An Interactive Graphical User Interface (IGUI) [6] allows the user to control the operation of the RCD and to obtain status and error information. The IGUI interfaces to the supervision system, the Information Service [7] and the Message Reporting System [8] which are standard components of the TDAQ Online System. The same IGUI is used throughout ATLAS from small test systems to the final production system.

### D. Monitoring, histogramming and calibration.

The RCD allows to transport monitoring data, in the form of event fragments or histograms, from modules in the ROD crate to either the Event Monitoring Service [9], a framework for developing detector specific event monitoring software or to the Histogramming Service [9] which provides facilities for receiving and displaying histograms.

Some detectors perform calibration at the level of the ROD crate when access to a full event is not required. In this context the triggering of the readout of the calibration data is important. In ATLAS, standard trigger modules have been developed [10]–[12] with software support in RCD. For calibration, the trigger modules are programmed to run in local, detector specific mode allowing to take data independently of other detectors.

## IV. RCD DEPLOYMENT

This chapter deals with aspects of recent deployments and developments of RCD: its usage during the 2004 Combined Test

Beam, enhancements of the functionality and its role in the detector commissioning phase.

### A. 2004 Combined Test Beam

In the 2004 Combined Test Beam, a slice of the ATLAS detector was assembled with the objective of integrating and testing all the components of the experiment under real conditions. Elements from nine different types of detectors, including the inner detector, calorimeter and external muon chambers, were installed in a beam line and exposed to different kinds of particle beams.

The detector readout electronics was connected to a scaled-down version of the final ATLAS DAQ system, which was operated in realistic data taking conditions (24 hours a day 7 days a week) for about six months.

The readout electronics for the various detectors was mostly available in early prototyping versions, sometimes with limited functionality; hence the RCD software had to be used for different tasks.

- All the detectors but one *controlled* and *configured* their modules via the RCD.
- Some detectors used the *ROD emulation* functionality to read non-formatted data from different types of modules, combine it and build formatted fragments, as the final ROD modules were not available at the time.
- More than half of the detectors had real ROD modules and some performed *data readout* with the RCD.
- All the detectors but one used the RCD *monitoring* functionality.

RCD was also used to control and configure the VMEbus modules (see Section III-D) used to interface the readout electronics with the first-level trigger, to handle incoming timing and trigger signals and to provide feedback on busy conditions.

### B. RCD enhancements and performance measurements.

After the 2004 Combined Test Beam, the usage of the RCD was analysed in detail, and a number of changes were performed in order to simplify the API and add some functionalities [13].

To test the new software, and while the final RODs of the detectors were not available, a test bed was set up to emulate the behaviour of a ROD. It consists of a VMEbus memory module plus a VMEbus interrupt module [14], which reproduce the three main ROD functionalities: memory, registers and interrupt capability. The memory is filled with ROD fragments to emulate the ROD data source. The trigger to read the ROD fragments is provided by the VMEbus interrupt module, which can be used in either polled or interrupt driven mode.

When doing polled data read out, a register in the VMEbus interrupt module is read for the presence of a "data available" bit. This bit is set via a NIM pulse which runs at a given frequency. When the bit is set, a fragment is read from the memory. When using interrupts, every time a fragment is available (given again by a NIM pulse), the VMEbus interrupt module issues an interrupt, and a fragment is retrieved from the memory.

Fig. 8 shows the performance measurements done for such a test system running with two readout modules, each one of them with two data channels. The measurements were performed for




Fig. 8. Performance measurements of a ROD crate test system running four data channels.

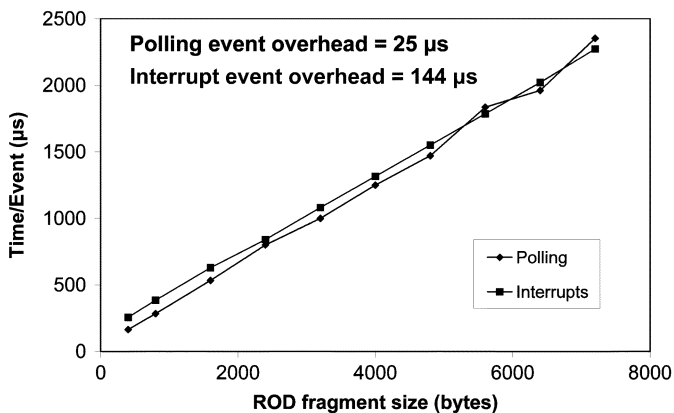


Fig. 9. Performance measurements with a ROD emulator connected to an RCD event building system.

both polling and interrupt modes, to compare the event overhead and in general the ROD fragment transfer time.

The total transfer rate measured in this system, using single cycles to read from the VMEbus memory, is 3.6 Mbyte/s, which is close to the maximum rate that can be achieved in this system.

In polled mode the event overhead is about 25  $\mu$ s and in interrupt mode 100  $\mu$ s. In both cases the event overhead is small compared to the typical transfer of a 1 kbyte fragment across VMEbus which is the performance requirement of RCD.

One of the enhancements to the RCD framework was the capability of acting as an event builder (see Section II-E). To validate the new event building software, a PC running an RCD event builder was added to the test system described above. Full events are built out of the ROS fragments received from the VMEbus crate. The network link between the ROD crate and the RCD event builder is a Fast Ethernet connection (100 Mbit/s). The performance measurements can be viewed in Fig. 9.

The performance is comparable to that shown in Fig. 8, which shows that adding an RCD event builder has a small impact on the performance of the system.

### C. Detector Commissioning

The commissioning of detectors in ATLAS has already started and will last until 2007. Within this context, RCD will be the main tool concerning the data acquisition functionality.

The usage of the RCD will be similar to that of last year's test beam, with the ROD emulation modules gradually being replaced by the final detector RODs. Whenever a larger part of a detector or several detectors will have to be read out, requiring more than one ROD crate, multi-crate event building will be used to acquire combined data.

## V. CONCLUSION

The ROS software framework, originally developed to meet the dataflow requirements at the level of the common readout system in the ATLAS DAQ (see Fig. 1), has been extended to the detector specific domain, largely based on VMEbus, between the front-end electronics and the readout system. Due to the modular design of the ROS software, and in particular the technique of using plugins, new functionalities could be added so as to meet the additional requirements of RCD.

A first version of RCD was tested in the ATLAS Combined Test Beam of 2004 and the requirements were met both in terms of functionality and performance. Based on the feedback from the test beam, RCD was further developed mainly to add the event building capability and to improve the handling of VMEbus modules. The RCD software will be an important data acquisition tool for use in test beds, for detector commissioning and as a component of the final ATLAS DAQ for detector control, monitoring and calibration.

## ACKNOWLEDGMENT

The authors would like to thank the ATLAS TDAQ community for their contributions to the work presented in this paper.

## REFERENCES

- [1] *ATLAS experiment*, [Online]. Available: [http://en.wikipedia.org/wiki/ATLAS\\_experiment](http://en.wikipedia.org/wiki/ATLAS_experiment).
- [2] *ROD Crate DAQ Task Force*, [Online]. Available: <https://edms.cern.ch/document/344713/1>, Data Acquisition for the ATLAS Read-out Driver Crate (ROD Crate DAQ) - Definition.
- [3] *VMEbus International Trade Association*, [Online]. Available: [www.vita.com](http://www.vita.com).
- [4] G. Crone, S. Gameiro, B. Gorini, M. Joos, E. Pasqualucci, and J. Petersen, *ROD Crate DAQ User's Guide* [Online]. Available: <http://edms.cern.ch/document/577958/1>.
- [5] A. Kazarov, G. Lehmann, D. Liko, S. Wheeler, and H. Zobernig, *ATLAS TDAQ: Controller Requirements* [Online]. Available: <http://edms.cern.ch/document/431663/2.1>.
- [6] S. Camara, *Design of the Integrated Graphical User Interface of Atlas DAQ/EF Prototype-1* [Online]. Available: <http://atddoc.cern.ch/Atlas/Notes/144/Note144-1.html>.
- [7] *Information Service*, [Online]. Available: <http://atddoc.cern.ch/Atlas/DaqSoft/components/is/Welcome.html>.
- [8] *Message Reporting Service*, [Online]. Available: <http://atddoc.cern.ch/Atlas/DaqSoft/components/mrs/Welcome.html>.
- [9] M. Barczyk, *Online Monitoring software framework in the ATLAS experiment* [Online]. Available: <http://doc.cern.ch/archive/electronic/cern/others/atlnot/Note/daq/daq-2003-005.pdf>.
- [10] P. Farthouat and P. Gallno, *TTC-VMEbus Interface* [Online]. Available: <http://ttc.web.cern.ch/TTC/TTCviSoec.pdf>.
- [11] P. Farthouat, P. Gallno, G. Schuler, and R. Spiwoks, *Local Trigger processor* [Online]. Available: [http://ific.uv.es/tical/rod/docs\\_vme\\_daq/ltp-spec-v6.pdf](http://ific.uv.es/tical/rod/docs_vme_daq/ltp-spec-v6.pdf).
- [12] Per Gallno, *ATLAS ROD Busy Module MkII* [Online]. Available: [https://edms.cern.ch/file/456631/4/ROD\\_BUSY\\_MkII\\_manual\\_27.pdf](https://edms.cern.ch/file/456631/4/ROD_BUSY_MkII_manual_27.pdf).
- [13] G. Crone, B. Gorini, M. Joos, E. Pasqualucci, J. Petersen, and W. Vandeli, *Extensions of the IOManager architecture for ROD Crate DAQ* [Online]. Available: <http://edms.cern.ch/document/554806/1>.
- [14] "Creative Electronic Systems S.A.," *RCB 8047 CORBO VME Read-Out Control Board, User's Manual*, Switzerland, Aug. 1993.