Overview and future developments of the FPGA-based DAQ of COMPASS

Overview and future developments of the FPGA-based DAQ of COMPASS

Y. Bai,[a] M. Bodlak,[b] V. Frolov,[c,d] V. Jary,[e] S. Huber,[a] I. Konorov,[a] D. Levit,[a] J. Novy,[d,e] D. Steffen[a,d,1] and M. Virius[e]

[a] *Physik-Department E18, Technische Universität München,*
 James-Franck-Straße 1, 85748 Garching, Germany

[b] *Department of Low-Temperature Physics, Charles University,*
 V Holesovickach 2, 180 00 Praha 8, Czech Republic

[c] *Joint Institute for Nuclear Research,*
 Joliot-Curie 6, 141980 Dubna, Moscow region, Russia

[d] *European Organization for Nuclear Research — CERN,*
 CH-1211, Geneve 23, Switzerland

[e] *Faculty of Nuclear Sciences and Physical Engineering, Czech Technical University,*
 Brehova 7, 115 19 Prague 1, Czech Republic

 E-mail: dominik.steffen@cern.ch

ABSTRACT: COMPASS is a fixed-target experiment at the SPS at CERN dedicated to the study of hadron structure and spectroscopy. Since 2014, a hardware event builder consisting of nine custom designed FPGA-cards replaced the previous online computers increasing compactness and scalability of the DAQ. By buffering data, the system exploits the spill structure of the SPS and averages the maximum on-spill data rate over the whole SPS cycle. From 2016, a crosspoint switch connecting all involved high-speed links shall provide a fully programmable system topology and thus simplifies the compensation for hardware failure and improves load balancing.

KEYWORDS: Data acquisition concepts; Data acquisition circuits

---

[1]Corresponding author.

Contents

1 Introduction

COMPASS (COmmon Muon and Proton Apparatus for Structure and Spectroscopy) [1] is a fixed-target experiment at the M2 beam line of the SPS accelerator at CERN. This unique beam line provides muon as well as hadron beams of high intensity and therefore attracted a strong community for Quantum Chromodynamics (QCD) studies in the not well-understood non-perturbative region. Merging the efforts of different groups finally led to a highly flexible and versatile spectrometer setup, which can be adapted to various measurements and physics programs. Data taking of the first physics run started in 2002. Since then, the number of channels increased from 190,000 to approximately 300,000 as well as the trigger rate from 5 kHz to 30 kHz. After one decade of data taking, the physics program of COMPASS was extended in 2012. The COMPASS-II program [2] started with a physics run for the study of the polarized Drell-Yan (DY) process in the years 2014 and 2015 followed by a run dedicated to Deeply Virtual Compton Scattering (DVCS).

After twelve years of successful data taking, it became difficult to maintain the Data AcQuisition system (DAQ) based on PCI technology. In order to provide a reliable and modern system for the next physics program, a new DAQ was designed and commissioned during the technical shutdown of CERN accelerators in 2013-2014. Whereas the old DAQ employed traditional "software" event building based on distributed computers interconnected via a high performance network, the new DAQ exploits the superior properties of Field Programmable Gate Arrays (FPGAs) and executes the event building completely in hardware. In the following, a detailed overview of the new FPGA-based DAQ (FDAQ) of COMPASS will be given.

2 Trigger Control System, front-end electronics, and data concentrator modules

Before going into details of the new FDAQ, this section will describe the parts of the electronics in the readout chain which are not affected by the upgrade and stay untouched. These parts define the framework requirements for the FDAQ since it has to interface with them and preserve their functionality.

In COMPASS, data taking is initiated on a trigger decision. The triggers are distributed by the Trigger Control System (TCS) [3] together with a reference clock and event labels to all data concentrator and event builder modules via a passive optical fiber network. The event labels consist of three numbers: the spill number in the current run, the event number in the spill, and the event type. The first two numbers together with the run number uniquely identify each event. The data concentrator modules (CATCH [4], HGeSiCA [5], or GANDALF [6]) forward the trigger and the common clock to the connected front-end electronics. The data, which they receive upon each trigger from the front-ends, are merged with the event identification of the TCS.

The detectors in the COMPASS spectrometer use different types of front-end electronics according to the requirements of the particular detector. They all have a pipeline architecture to cope with the high data rate (~1.5/ Gb/s) and use either Time to Digital Converters (TDC) or Analog to Digital Converters (ADC) for digitization.

From the data concentrator modules the data are transmitted to the event builder (EB) via optical links using the S-Link protocol [7] developed at CERN. Because the data rates are different for different detectors, an optional intermediate multiplexing stage has been developed. It allows to merge the data streams of up to four concentrator modules of detectors with low data rate. This lowers the number of required optical connections to the EB by exploiting the S-Link bandwidth more efficiently (see figure 1). Behind this stage, there are between 64 and 120 optical links entering the EB depending on the current physics program and spectrometer setup.

3 Design of the hardware event builder

Due to aging of components and related loss of modules, the DAQ's event building architecture needed a revision [8]. Using a traditional event building architecture, it had based on distributed online computers interconnected via an Ethernet Gigabit network. However, the communication pattern in an event building network is different from that in a general network, thus handled inefficiently by switches, and hence requires external traffic shaping [9]. Concurrent developments in FPGA technology, such as increased I/O bandwidth (> 3 Gbps) and support for high performance SDRAM even on low-cost chips, made also FPGAs suitable for event building purposes. The expected higher reliability due to increased compactness of the system led to the decision to change to a fully FPGA-based DAQ.

The EB of the COMPASS DAQ has to fulfill two main requirements: first, it has to be capable of processing the data rate of 1.5 GB/s in average during the on-spill period of the SPS of 5 to 10 seconds, with a maximum peak data rate of 8 GB/s, followed by an off-spill period between 10 and 20 seconds. For the second requirement, it has to be versatile and scalable in order to address the different spectrometer setups for the various physics programs of the COMPASS experiment.
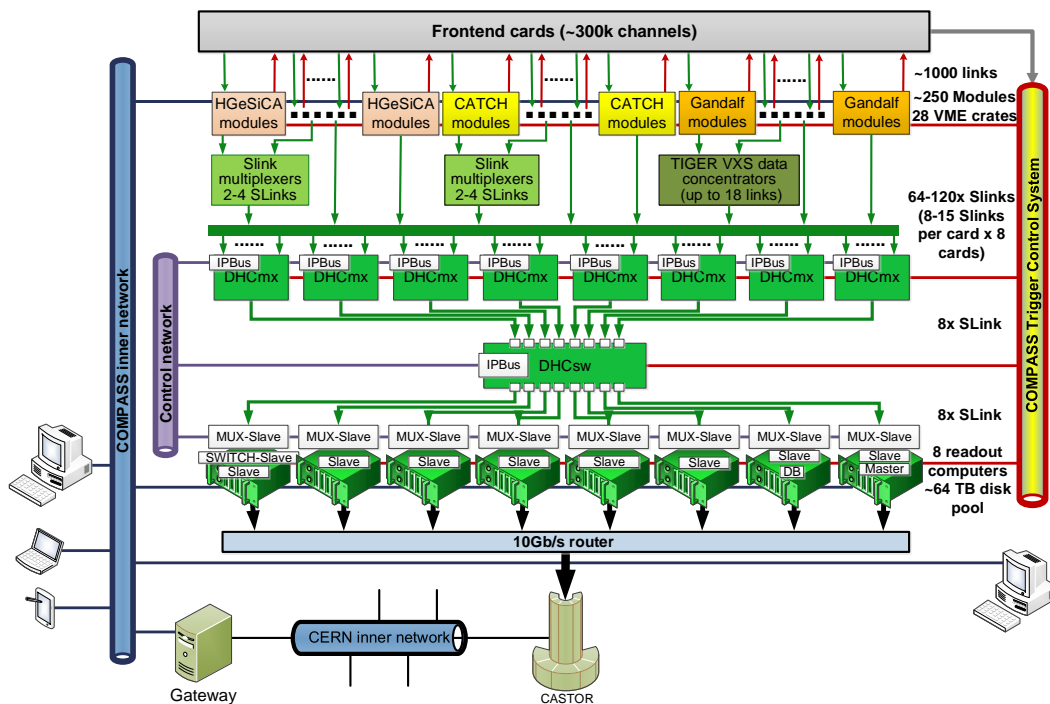
Figure 1. DAQ topology with hardware event builder.

The hardware EB of the new FDAQ consists of nine custom designed DAQ units with identical PCB layout and eight-off-shelf FPGA cards, called spillbuffer cards, which are plugged into eight readout computers. Despite their identical layout, the DAQ units are used differently in the EB, as can be seen in figure 1: eight of them are programmed with a firmware that allows to use them as 15x1-multiplexer (DHCmx), and one is programmed as 8x8-switch (DHCsw), distributing the data sent by the multiplexers to the spillbuffers in a round robin manner on event basis. The spillbuffer cards receive full events from the switch and forward them via a PCI-E interface to the buffer memory in the readout computers.

The system architecture makes use of the spill structure of the SPS beam by buffering data on different levels of the EB and averaging the maximum on-spill data rate over the whole SPS cycle. As a result, the online readout computers work independently from the initial detector rate and deal only with a lower averaged sustained rate which depends on the SPS supercycle. In the moment, the sustained rate is limited by the PCI-E performance to 125 MB/s per readout computer, thus 1 GB/s for eight readout computers in total.

The system uses three independent network interfaces for synchronization, event building, and configuration or monitoring. For synchronization and distribution of event information the already mentioned TCS is used. The data flow and thus the event building follows the S-Link specifications, such that already existing hard- and firmware could be used. For configuration and data flow control of the hardware EB, the IPbus [10] protocol is used which runs over Ethernet network. This functional division simplifies software architecture, optimally synchronizes processes, and offers an efficient usage of the network bandwidth.
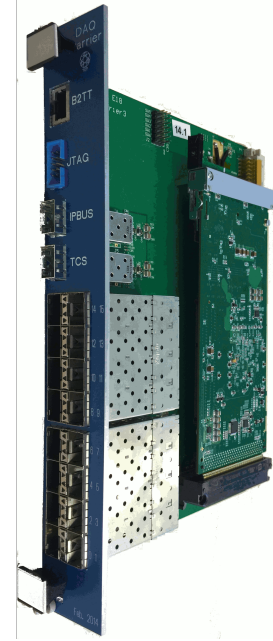
Figure 2. Five DAQ units in a VME crate (left) and more detailed view of one unit consisting of DAQ carrier card and DHCxx module (right).

The new FDAQ was deployed and comissioned for the two months of DY pilot run in 2014. The DY spectrometer setup included only a sub set of COMPASS spectrometer which leads to an average event size of ~ 25 kB and an on-spill data rate of ~ 600 MB/s. Due to the lower data rate, the DAQ setup was scaled down to five DHCmx and four readout engines.

3.1 Hardware layout

One DAQ unit consists of a carrier card and one FPGA board called Data Handling Card (DHC), as shown on the right in figure 2. Depending on which operation mode (multiplexer or switch) is loaded into the FPGA, the DHC is called DHCmx or DHCsw respectively. The DAQ carrier design complies to the 6U VEM standard, although only power and cooling of the VEM crate is used. All custom designed modules fit into a single VME crate (left side in figure 2) and provide the DHCs with all necessary interfaces over the front panels: TCS receiver, Gb Ethernet for the IPbus slow control network, a JTAG interface for backup FPGA programming, and 16 serial data links for data transmission. The DHC complies to the Advanced Mezzanine Card (AMC) specification and is therefore compatibel to the Advanced Telecommunications Computing Architecture (ATCA). The heart of the DHC is a Xilinx XC6VLX130T FPGA from the Virtex6 family. Moreover, it is equipped with 4 GB DDR3 SDRAM.

3.2 DHC firmware architecture

An important feature of the FDAQ is the data consistency check implemented in the firmware on each stage of the event building. Together with the information received by the TCS, such as event label, trigger signal, and Start Of Burst (SOB) signal, the consistency check is able to recognize following errors in the data stream:
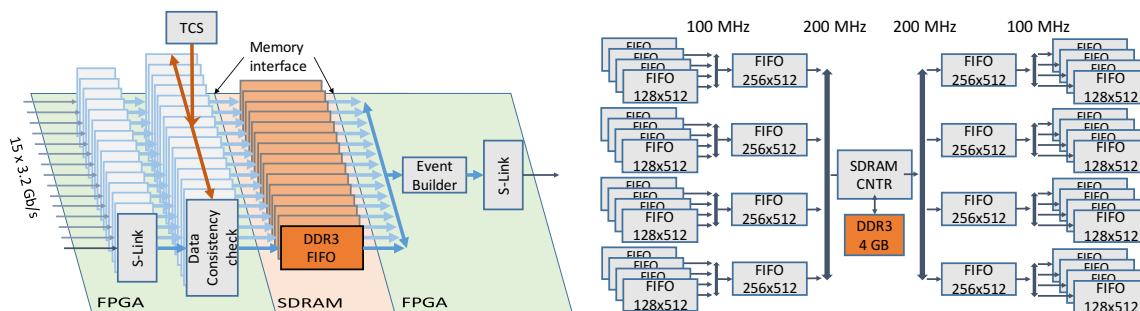
Figure 3. Block level diagram of the firmware for the DHCmx (left) and more detailed view of its memory interface (right).

- transmission errors detected by the S-Link

- truncation errors which may occur in the front-end electronics due to limited buffer sizes; in this case, there is a mismatch between the declared and the real data block size.

- inconsistency of the event label, a mismatch between the actual and the expected event label

- missing data, i.e. data are not received within a programmable time interval after the trigger

Data blocks with errors are dropped unless the error is explicitly masked and empty blocks with correct event labels are generated. This guarantees data consistency and a decodable data structure and enables the readout computers to recognize and localize origin of occured errors in the data stream. Additionally, the information about these errors is collected in local registers which are accessible by the run control software via the IPbus network. Thus, the FDAQ provides two different ways of error diagnostics and keeps running regardless of occurring errors.

A block level diagram of the firmware implemented in the DHCmx can be seen on the left side of figure 3. After the obligatory data check of each incoming data stream, the data are transferred into the SDRAM, which is organized as multiple, sizeable, FIFO-like buffers. For efficient use of the DDR memory volume, the size of the preallocated buffers is programmable and allows for sharing the memory proportionally to the incoming data rate. In order to use the maximum bandwidth of 6.4 GB/s of the DDR3 RAM, the memory interface between FPGA and SDRAM (figure 3) has two stages of multiplexing and demultiplexing. The double layer architecture provides a higher access rate to the SDRAM than the incoming S-Links can deliver. Furthermore, the logic in the SDRAM controller grants priority to the incoming stream to avoid data losses. After the buffer, the subevent building logic block sequentially collects fragments of the same event into one subevent and transmits it to the DHCsw.

The current version of the firmware for the DHCsw is an extension of the firmware for the DHCmx, as can be seen on the left of figure 4. It features the same data consistency check and memory interface, but it scales the number of attached memory buffers per source which are filled in a round robin based manner. The buffers are read by the same amount of event builder, each effectively filling two S-Links. Due to limited logic resources on the FPGA, the number of buffers per source could only be scaled to two at six incoming links. Hence, the current firmware supports up to six incoming S-Links from the DHCmx modules and four outgoing to the readout computers
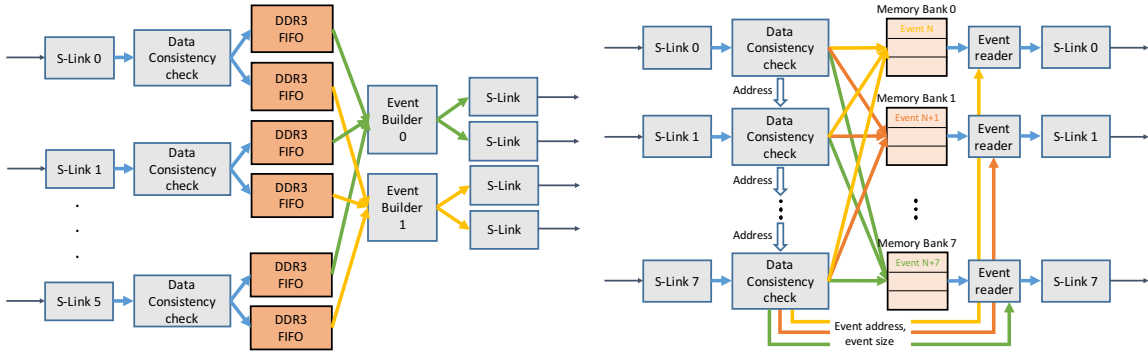
Figure 4. Block level diagram of the current (left) and future (right) firmware for the DHCsw.

as this is sufficient for the DY setup in the years 2014/2015. For next years's DVCS-run however, the firmware will be updated, as shown on the right side of figure 4. The event fragments will already be sorted by the data check logic. The buffer memory will therefore be preallocated into equal-sized banks of 0.5 MB, which corresponds to the maximum possible event size in COMPASS. As soon as the last data check logic has filled its event fragment into the memory bank, it passes the memory address and size of the event to an event reader, which reads the whole event from the DDR memory and sends it via S-Link to its corresponding spillbuffer. This firmware setup requires a revision of the memory interface and provides less efficient use of the DDR3 SDRAM, but it has full throughput bandwidth of 1.6 GB/s.

3.3 Software architecture

The Run Control, Configuration, And Readout Software (RCCARS) [11] for the new hardware EB is a multilayer system centered around the master process. The communication between processes of the software uses the DIM library [12].

The RCCARS incorporates different kind of processes. The master process exchanges status information and control commands between the slave processes and the Graphical User Interfaces (GUIs). It also contains the main application logic, e.g. the state machine, which controls the whole system and is responsible for the start and end of a run. The Run Control GUIs can be started on any computer inside the COMPASS network for monitoring purposes, but only one of them can interactively control the master process. Moreover, there are slave processes which control and monitore the DAQ hardware units. These slaves access the memory registers on the FPGAs via the Ethernet network. Finally, the readout slaves are the only realtime software processes of the system. They are responsible for data transfer from the spillbuffers to hard disks. By doing so, they transform the data format given by the structure of the hardware event builder into the old DATE-format to ensure compatibility of the new FDAQ with already existing online monitoring and analysis tools. Apart from the readout slaves, the software is not involved in data processing, but only supervises the hardware EB.
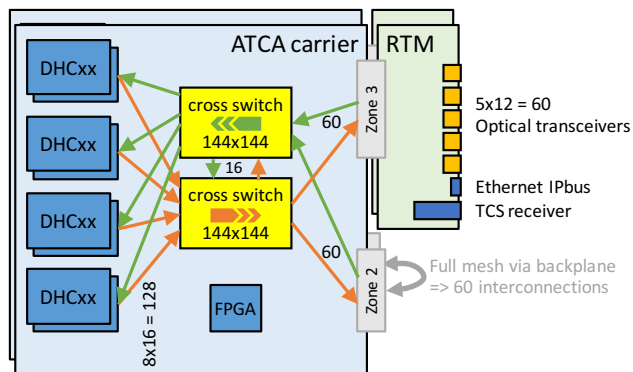
Figure 5. Diagram of the funtionality of the two ATCA carrier cards implementing the crosspoint switch.

4 Future developments

For the DVCS-run in 2016, a fully programmable crosspoint switch will be introduced, which will wire all point-to-point high speed links. The switch will therefore provide a fully customizable DAQ network topology between data concentrator modules, the event building hardware, and the online computers. The adaptability of the system topology allows to compensate for hardware failure in the EB by activating spare resources replacing broken or malfunctioning ones.

In the first step of the implemention, changing the network topology will only be possible between two runs. This avoids synchronous and time critical reconfiguration of the crosspoint switch in interaction with the RCCARS. Later on, the feature of the TCS to distribute any timing information to all hardware nodes of the EB simultaneously will be used for reconfiguration the event building network topology on-the-fly. Algorithms shall, in case of hardware failure or load imbalance, synchronously reconfigure the FDAQ topology to recover system performance. In this way, the fully programmable crosspoint switch will improve system reliability and minimize data loss.

The implementation of the crosspoint switch will come along with a change from the VME to ATCA standard. One ATCA carrier card will be equipped with up to eight DHC Mezzanine Cards. For the connections with the data concentrator modules and with the eight readout computers, 120 high-speed optical links are needed. This entails the use of two Rear Transition Modules (RTMs). Each RTM will be equipped with five CXP 12-channel parallel optical transceivers providing 60 optical links, an Ethernet interface for the IPbus network, and a TCS receiver. The data links from the optical transceivers are routed to the AMC daughterboards via two 144x144 programmable crosspoint switches. As crosspoint switches, the VSC3144-02 [13] providing 6.5 Gb/s fully asynchronous data throughput was chosen. They will be controlled by a Xilinx XC7A200T FPGA from the Artix7 family. Moreover, the FPGA will act as a hub, which distributes the IPbus and TCS commands to the DHC daughterboards using the Universal Communication Framework (UCF) which is currently developed at Technische Universität München. The UCF is designed for star-like network topologies and supports TCS, IPbus, and JTAG transmission. A functional diagram of the ATCA carrier cards is shown in figure 5.

ATCA's switch fabric supports a full mesh interconnect via the backplane. The shelf, which will be used, is designed for two carrier cards and will therefore provide the needed 60 high-speed

links for interconnection of the two ATCA modules with the possibility to house up to 16 DHC modules, enough to integrate backup modules.

5 Results and conclusion

The FDAQ was successfully deployed and commissioned in 2014 and allowed to successfully take data for nominal Drell-Yan conditions during the run 2015. The hardware EB consisting of five DHCmx modules and four spillbuffer cards is capable of dealing with the on-spill data rate of 600 MB/s. The realtime readout processes of the RCCARS, likewise running on four computers, withstood the averaged sustained data rate of 200–250 MB/s. However, the run also revealed minor problems. Some data format errors are not recognized by the data consistency check and therefore may cause problems in the event building process. In combination with rare crashes of the RCCARS, this leads to problems and interruption of data taking about once a day. The errors on software as well as on firmware side have been identified and will be removed after a phase of tests for the next run. All in all, the initial demands in terms of reliability are already fulfilled and the development of the system still continues. Updates for 2016 not only comprise the necessary upgrades to cope with the higher data rate of the DVCS setup but also the deployment of the crosspoint switch.

Due to its versatility and scalability, the FDAQ is also suitable for other high-energy physics experiments. Recently, it has been chosen by a second experiment at the SPS which is searching for light dark matter [14].

References

[1] COMPASS collaboration, P. Abbon et al., *The COMPASS experiment at CERN*, *Nucl. Instrum. Meth. A* **577** (2007) 455 [hep-ex/0703049].

[2] F. Gautheron et al., *COMPASS-II proposal*, CERN-SPSC-2010-014, CERN, Geneva Switzerland (2010) [SPSC-P-340].

[3] B. Grube, *A trigger control system for COMPASS and a measurement of the transverse polarization of Λ and Ξ hyperons from quasi-real photo-production*, Ph.D. thesis, Technische Universität München, Munich Germany (2006).

[4] T. Schnmidt, *A common readout driver for the COMPASS experiment*, Ph.D. thesis, Albert-Ludwigs-Universität Freiburg, Freiburg Germany (2002).

[5] B. Grube et al., *Architecture of a common GEM and silicon readout for the COMPASS experiment*, in proceedings of 7th *International Conference on Advanced Technology and Particle Physics (ICATPP-7)*, Como Italy October 2001.

[6] S. Bartknecht et al., *GANDALF — design of a modular high resolution transient recorder for high energy physics*, *IEEE Nucl. Sci. Symp. Conf. Rec.* **N42-2** (2009) 2221.

[7] H.C. van der Bij et al., *S-LINK, a data link interface specification for the LHC era*, http://hsi.web.cern.ch/HSI/s-link/introduc/rt97.ps, (1997).

[8] M. Bodlak et al., *New data acquisition system for the COMPASS experiment*, in proceedings of *Topical Workshop on Electronics for Particle Physics (TWEPP)*, Oxford U.K. September 2012 [2013 JINST **8** C02009].

[9] T. Uchida et al., *New communication network protocol for a data acquisition system*, *IEEE Trans. Nucl. Sci.* **53** (2006) 286.

[10] C. Ghabrous Larrea et al., *IPbus: a flexible ethernet-based control system for xTCA hardware*, 2015 *JINST* **10** C02019.

[11] M. Bodlak et al., *FPGA based data acquisition system for COMPASS experiment*, *J. Phys. Conf. Ser.* **513** (2014) 012029 [arXiv:1310.1308].

[12] C. Gaspar and M. Dönszelmann, *A distributed information management system for the DELPHI experiment at CERN*, in proceedings of 8th *Conference on Real-Time Computer applications in Nuclear, Particle and Plasma Physics*, http://dim.web.cern.ch/dim/papers/dim.ps, Vancouver Canada June 1993.

[13] Vitesse, *VSC*3144-02 6.5 *Gbps* 144 × 144 *asynchronous crosspoint switch*, (2011).

[14] S. Andreas et al., *Proposal for an experiment to search for light dark matter at the SPS*, arXiv:1312.3309.