ATLAS Phase-II Upgrade Project

SoC for ATLAS subsystems: Requirements Document for HL-LHC

Abstract

This document describes the high-level functional and performance requirements of the SoC for all the deliverables of ATLAS Phase-II Upgrade Project (UPR) and their interfaces. While the conceptual design and implementation of the ATLAS detector upgrades is described in the individual sub-detectors TDRs, this document provides the UPR systems and sub-systems with a common framework for requirements capture and the specification of interfaces, allowing detailed designs to proceed concurrently.

ATLAS-SoC Interface Requirements Document	
ATLAS Doc:	ATU-GE-ES-0005
EDMS Id:	2373932
EDMS Url:	https://edms.cern.ch/document/2373932
Version:	1.2
Created:	May 7, 2020
Last modified:	July 24, 2020

Prepared by:	**Checked by:**	**Approved by:**
Revital Kopeliansky Wainer Vandelli	ATLAS Upgrade Steering Committee	ATLAS Upgrade Steering Committee

[INTENTIONALLY BLANK PAGE]

REVISION HISTORY

Revision	Date	Description
1.0	09/07/2020	First release
1.1	21/07/2020	Acknowledgement statement added
1.2	24/07/2020	Minor wording and spelling fixes

ACKNOWLEDGEMENT

We acknowledge the help and contributions of the ATLAS Upgrade community, the ATLAS TDAQ Phase-II community and the ATLAS TDAQ system administration team. This document is the result of fruitful discussions with and valuable input from many members of these communities.

TABLE OF CONTENTS

1

INTRODUCTION

1.1 SCOPE AND PURPOSE

The objective of this document is to define System-on-Chip (SoC) related requirements for all the Phase-II ATLAS UPR systems deliverables and to provide guidelines to the UPR system and sub-system managers for the redaction of their requirement documents.

The document applies to all deliverables including SoCs designated for direct communications with external devices.

1.2 APPLICABLE STANDARDS

This document is based on standards defined in the following references, which shall be used to develop requirement documents in the ATLAS Phase-II UPR.

REFERENCES

[1] David D. Walden et al., eds. *Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities*. 4th ed. Hoboken, NJ: Wiley, 2015. ISBN: 978-1-118-99940-0 (cit. on p. 2).

[2] ISO/IEC/(IEEE). *ISO/IEC (IEEE Std) 24765 : Systems and Software engineering - Vocabulary systems*. 2010 (cit. on p. 2).

[3] IEEE. “Adoption of ISO/IEC 15288:2002 Systems Engineering-System Life Cycle Processes”. In: *IEEE Std 15288-2004 (Adoption of ISO/IEC Std 15288:2002)* (2005). DOI: 10.1109/IEEESTD.2005.96287 (cit. on p. 2).

1.3 CONTENTS

The document is organized as follows:

- Overall hardware architecture
- SoC functions
- Requirements associated to hardware, functionality and integration

1.4 Frequently Used Terms

Among the systems and sub-systems' earliest priorities on a project is to define a precise terminology in order to establish unambiguous communication and definition of the system and its elements, functions, operations, and associated processes [1].

The glossary provided in Appendix A is based on the definitions reported in [2, 3] and nomenclature used in the ATLAS Phase-II project.

2

GENERAL COMMON REQUIREMENTS

2.1 SYSTEM HARDWARE ARCHITECTURE

In the ATLAS Phase-II UPR it is anticipated that SoC will play a major role in the configuration, monitoring and operation, mainly in custom ATCA electronics but potentially could also be used in other cases.
It is anticipated that the overall number of SoC devices to be used in ATLAS will be at the order of 1500 and therefore requires a proper planning and coordination of it's usage which is the purpose of this document.

SoC will allow implementing several different functions, as discussed in Section 2.2. In particularly in the case of an ATCA usage, these functions will require dedicated hardware connections to the onboard components as suggested in 2.3. The SoC will master one or more specific bus dedicated to the monitoring of onboard components such as optical transceivers, FPGAs and ASIC chips, etc... According to the ATCA compliance [4], these are considered as non-critical components. Conversely, critical components, like power management modules, DC-DC converters, FPGA power modules, temperature sensors near high-power density regions, will be monitored via the onboard IPMC on at least one dedicated bus.

Remark 2.1: *TDAQ SoC user requirements document*

Being the largest consumer of SoC devices in ATLAS, TDAQ has compiled a SoC user requirements document that is serving as baseline for the ATLAS document. As such, some of the more detailed requirements & recommendations have been generalized in order to reflect the overall usage in ATLAS. Further detailed can be found in [5].

2.2 ATLAS PHASE-II SOC FUNCTIONS

The SoC in the Phase-II UPR are expected to cover several different and independent functionalities. Onboard functions will include slow control and monitoring, and spy-buffer collection. Externally, the SoC will be communicating with DCS and run control infrastructures and perform local data-taking for commissioning periods. SoC will enable interactive testing and debugging of the boards. In some cases, realtime processing of trigger data is being considered.

2.3 Hardware Requirements

The following requirements and recommendations were established based on different sources: common discussions within the ATCA community that includes the systems coordinators, system designers, system engineers and external services coordinators (Technical-coordination, System Administrators, DCS, DAQ, etc...). Furthermore, several surveys where distributed among the TDAQ ATCA subsystems, and conclusions deduced from the lessons learned from Phase-I upgrades.

Remark 2.2: *Discussions within the ATLAS ATCA community & related surveys*

The brainstorming discussions mentioned where coordinated and held within the TDAQ community with minor representation of other ATLAS systems. The survey was distributed amongst TDAQ systems only. However, the conclusions were found to be general enough to apply for all other ATLAS systems planning to use a SoC device.

Definition 2.1: *Hardware requirements*

Hardware requirements describe the essential hardware features of SoC.

As of the preparation of this document, the majority in ATLAS prefers the usage of Xilinx Zynq UltraScale+ FPGAs with ARM 64bit CPUs. Some systems are also planning to use Intel Stratix-10 devices. Both FPGA models include microprocessors and programmable logic.

Commonality has several advantages, among them are the benefit of being part of a community for sharing knowledge and tools, as well as support from external services. In the case of ATLAS, the choice of a common vendor might not be possible to achieve, although recommended. Within the anticipated exceptions, the document still aims to apply for at least some level of commonality.

Recommendation 2.1: *Common vendor*

Sourcing SoC from a common vendor is anticipated to provide numerous advantages, especially looking forward to the long-term operation and maintenance of the ATLAS system. Common tooling, OS image management, software, ... will allow for common, reusable support. As such the systems are encouraged to deploy chips from the agreed upon vendor.

Remark 2.3: *Common vendor choice*

As of the day this document is being written, the majority in ATLAS (i.e. TDAQ, Tile calorimeter and LAr calorimeter in partial LATS and LASP-RTMs) are planning to use Xilinx as the SoC vendor, with the exception of the main LASP ATCA blade that will rely on Intel Stratix-10 SoC.

The majority of ATLAS ATCA systems are planning for front connection from the SoC device to the ATLAS Technical and Control Network (ATCN) through a top-of-the-rack Ethernet switch, with an exception of LAr that are planning to use a back-plane switch in their ATCAs that will connect to a Gateway PC.

In the later case, where a Gateway PC is the only device that is connected directly to the ATCN, the regular CERN IT rules for connecting PC servers to the network apply. However, for the first case, although it is foreseen that a usage of a single Ethernet port in the SoC should satisfy the anticipated functionality, it is essential to have a spare as a replacement in case of an unexpected fault, or if proven over time that a second interface is needed from whatever reason.

Recommendation 2.2: *Ethernet Connectivity*

SoC that is planned to be directly connected to the ATCN switch should be equipped with two (one plus spare) physically independent Ethernet ports. They should be capable of operating at least at GbE speed.

Ethernet networks require a unique MAC address associated to each physical port. Therefore SoC that will interface directly with ATCN must have a MAC address associated to each port.

Requirement 2.1: *MAC Address*

The MAC address associated with the Ethernet ports shall be either:

- uniquely allocated by the Vendor
- individually configurable for each SoC and each port

Operating systems have minimal requirements, in particular with respect to memory and storage, which the SoC should comply with.

Requirement 2.2: *OS Memory and Storage*

The SoC shall be equipped with at least:

- 1 GB of RAM per logical core
- 10 GB of local storage (e.g. SD card)

Additional local storage and memory needs required by the SoC functions and operation must be accounted on top of OS needs. Local storage must allow the definition of a read-write file-system (supposedly ext4).

Remark 2.4: *SD card write-cycle limitation*

Given the limited write cycles allowed by SD cards, the usage should be limited mainly to read-only functions requiring rare updates (e.g. operating system images). High-frequency write functions (e.g. log files) must not use the SD cards.

In case the SoC is required by the subsystem to connect directly to ATCN or interface directly with external services, it will have to be centrally managed. In order to benefit from such support, it is essential that the device will be upgradable during the HL-LHC lifetime, which could be a challenge considering the rapid evolution of operating system and software environments. At the day this document is been written, the only known approach to ensure upgradability is the use of a System-on-Module (SoM). In cases where there is no need to connect the device directly to the ATCN, upgradability may not be a concern. However this may imply that the centrally managed benefits won't be available either. In this scenario, subsystems will be responsible for interfacing with external services and for the maintenance of the operating system and software stack and the operational monitoring of the SoC.

Requirement 2.3: *SoC lifetime requirements*

Upgradability of the SoC OS is required for maintaining a secure interface into the network during its operation lifetime. In order to ensure long term functionality, systems including SoC in their hardware should make sure the device is integrated as a SoM.
In case the implementation of a SoM is not technically feasible, the SoC will not be replaceable and therefore may become outdated and unsupported during the ATLAS Phase-II lifetime. In this scenario the device will have to be isolated from the ATLAS control network. While different isolation scenarios are possible, it must be possible to operate the system under the most constraining one.
In this case the system would become accessible exclusively through a gateway machine and it shall be ensured that neither software nor hardware limitations (e.g. missing rack space for additional components) would prevent this. Moreover, the appropriate subsystem will be responsible for the development and maintenance of any additional software layer required in order to interface with common software tools (online-software libraries, Open Platform Communications Unified Architecture (OPC UA) server for DCS, etc.).

Remark 2.5: *SoM requirements indicated by TDAQ systems*

Based on the survey [6] several common characteristics for the SoM were identified:

- 12V powering
- onboard power supplies for the SoC core supply voltages with an interface to the ATCA blade power-sequencing circuitry
- support for different I/O supply voltages for the GPIOs, including I/O supply voltages powered by the ATCA blade
- I/O interfaces: UART, I2C (I^2C), SPI, GPIO, JTAG, AXI, Ethernet
- local flash and SD card interface

A sample platform for the SoM, as indicated also by Xilinx representatives, are the devices provided by Enclustra (https://www.enclustra.com/en/products/system-on-chip-modules/).

Remark 2.6: *Considering next generation of SoCs*

Future evolution of the Xilinx portfolio, for example the planned Versal family, may imply that multiple SoCs will be installed on each ATCA blade, whereas today most systems plan to include only one. As such, the need and scaling of ATCN connectivity should be coordinated in advance with the sub-detector Technical Coordination. Also, if these components will be expected to interface individually with external services (such as DCS, RunControl, etc...), planning in advance with the related coordinators should take place.

2.4 Functional Requirements

Definition 2.2: *Functional requirements*

Functional requirements describe qualitatively the system functions or tasks to be performed in operation by each deliverable.

One of the planned usages of the SoC is to perform monitoring (and in some cases may be also control) of other onboard components. In this case the SoC is usually defined as a master of a dedicated bus. Following

several discussions based on the Phase-I experience it has been concluded that master arbitration in ATCA blades is discouraged, hence a separation between the ATCA blade IPMC bus and the SoC bus is required. When establishing the separate buses, one has to bear in mind several main restrictions: the limited IPMI bandwidth between the onboard IPMC and the ATCA shelf-manager IPMC, the limited processing resources of the shelf-manager IPMC, and the requirement to fulfill the ATCA compliance [4]. Following this and after several discussions within the TDAQ ATCA engineers, a careful decision was made where the IPMC bus shall be dedicated to the critical components required by the ATCA compliance [4] and for the board safety. All additional monitoring and control aspects shall be handled via the SoC, taking advantage of greater resources and flexibility.

Requirement 2.4: *Monitoring of non-critical components*

All non-critical components (e.g. optical transceivers, LED status registers, extra temperature sensors), as defined by the ATCA compliance [4], shall be monitored through a SoC. The suggested implementation includes a dedicated bus mastered by the SoC to which the components are connected. Equivalent schemes implementing a communication link between the SoC and another devices with control of the monitoring components can be envisaged as long as they are compliant with the ATCA specification [4].

Remark 2.7: *ATCA compliance - critical components*

The critical components according to the ATCA compliance are specified under Section 3.9.3.1 on page 3-197 of PICMG 3.0 R3.0_withErrata001 [4] and are detailed over REQ 3.735-3.743. These include at least one temperature sensor per blade and it is also recommended to incorporate power-supply sensors.

Requirement 2.5: *ATCA blades onboard critical components list*

The baseline design of the different subsystems ATCA blades is rather common in terms of the basic chosen onboard components. The following list of common components specifies the ones that are considered critical according to both the ATCA compliance [4] and the blades safe functionality, and as such should be included in the IPMC bus, rather than on the SoC bus:

- Power-management chips
- DCDC converters
- FPGA power-modules
- At least one temperature-sensor per FPGA
- At least one onboard temperature sensor, near regions of high-power density in the PCB

Remark 2.8: *Critical components monitoring while limited resources*

In case the total of number of critical components to be monitored exceeds the input, output or processing capabilities of the IPMC, a careful decision needs to be made with regards to the frequency of the monitoring, while giving priority to power devices.

2.5 REQUIREMENTS ON SYSTEM INTEGRATION AND SOFTWARE

Definition 2.3: *Operational requirements*

Operational requirements define the operational conditions or properties that are required for the system to operate or exist. This type of requirement includes: human factors, ergonomics, availability, maintainability, reliability, and security.

Definition 2.4: *Modes and/or state requirements*

Modes and/or state requirements define the various operational modes of the system in use and events conducting to transitions of modes.

The SoC as a computing device must follow the CERN Computing Security rules [7]. In particular this is required for devices physically connected to the ATCN network that rely on unrestricted IP-level communication with network peers.

Requirement 2.6: *Direct ATCN connectivity*

Direct connection of computing devices to the ATCN requires compliance with different rules and practices.

- the Operating System must be approved by CERN
- it must be possible to regularly update the system within reasonable time
  - severe security flaws may require prompt patching, compatibly with the operation needs

According to the CERN Computing Security rules, computing devices in the technical networks must be centrally managed. In the case of ATLAS, SoC directly connected to ATCN should be fully managed by ATLAS TDAQ System Administrators. This will also guarantee a smooth integration in the overall P1 computing infrastructure in addition to the compliance with CERN computing rules. In turn, support from CERN and TDAQ System Administrators in the SoC management implies the use of a common OS. As such, the inputs [6] collected from the systems on their OS needs indicates a unanimous preference for a Linux-based OS. In particular CentOS and Petalinux appear to be the possible options.

Requirement 2.7: *The SoC Operating System*

All ATLAS sub-systems implementing SoCs devices (with an aim to be connected directly to the ATCN) should utilize a common OS. The choice of the OS to be deployed on the SoCs falls under the responsibility of the ATLAS Phase-II Upgrade Coordinator, after consulting with the sub-detector Upgrade Project Leaders, TDAQ System Administrators, CERN IT and ATLAS management at large.

Remark 2.9: *Operating System Definition*

A centrally managed, common operating system will include:

- kernel version and common driver set
- system tools and libraries (e.g. init system, compiler, ...)
- user-space libraries and software not distributed with CERN or ATLAS releases

Centrally managing a large number of computing devices requires mechanisms to network boot them.

This is needed to either distribute the OS images or to proceed to local installation in a coordinated fashion. In the commercial computing domain the de-facto standard is PXE or the equivalent functionality in more modern UEFI motherboards.

Requirement 2.8: *Common network booting procedure*

A common network boot mechanism, PXE or equivalent, must be available for the SoC.

The actual network boot mechanism will be identified as part of the device and operating system down-selection.
The ATCN is a shared network infrastructure which is necessary for the stable operation of ATLAS subsystems. It supports low-level hardware devices that may not implement high-level protocols and are therefore sensitive to network overload and packet drops. Hence the ATCN throughput of SoC may have to be constrained, especially for potentially high-throughput functions like local data-taking. The actual throughput limitations depend on the network layout and traffic patterns.

Recommendation 2.3: *ATCN Throughput*

Sustained high-throughput (>100 Mbps) functions should limit the traffic to the local switch.

The SoC will be one of the two interfaces to the DCS back-end. An OPC UA server will be operated either on the SoC or on the DCS server, using the Quick OPC UA Server Generation Framework (quasar) tool developed by the central DCS team. In both cases a mid-layer software will mediate between the sub-system hardware and the OPC UA server. Monitoring of the SoC I^2Cbus components and registers is established via communication between the OPC UA server and a corresponding client configured on the Local Control Station (LCS) machine. In case the OPC UA server will be deployed on the SoC, Central DCS will ensure compatibility of quasar with the chosen common SoC OS.

Requirement 2.9: *OPC UA server*

An OPC UA server built using a tool supported by central DCS (currently quasar) shall be used in order to establish a monitoring interface between the SoC and the DCS backend.

Remark 2.10: *OPC UA server installation*

As of the time this document is being written, the assumption is that the OPC UA server will be deployed on the SoC itself. Alternatively, the deployment could happen on a remote DCS server. The merits and drawbacks of both approaches will have to be evaluated and quantified by the subsystems. In both cases a mid-layer software, developed and deployed by the subsystem responsible of the SoC, is either recommended or needed in order to mediate between the server and the hardware monitoring, in order to ensure safe operation of the device.

It is anticipated the SoC will play a major role in configuring and operating the custom boards. In this sense, they will need to interface with the common ATLAS RunControl infrastructure in order to e.g. be aware of the DAQ Finite State Machine (FSM), implement recovery actions, connect to the monitoring and event sampling services, ...

Requirement 2.10: *SoC RunControl software*

A common software library, developed as part of TDAQ UPR Online Software, will be provided for communication between the ATCA blades SoC and DAQ RunControl.

Requirement 2.11: *SoC Actions*

Any action implemented as part of the SoC functions whose outcome may affect data-taking or triggering functionalities shall be coordinated with ATLAS operation. In particular the effects on data-quality, DAQ FSM, ... should be considered on a case-by-case basis.

Requirement 2.12: *Implementation of board-specific functions*

The implementation of board-specific functions and interfaces is responsibility of the corresponding sub-systems. This includes the development, testing and maintenance of the associated software, at least for the duration of HL-LHC operation.

REFERENCES

[1] David D. Walden et al., eds. *Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities*. 4th ed. Hoboken, NJ: Wiley, 2015. ISBN: 978-1-118-99940-0 (cit. on p. 2).

[2] ISO/IEC/(IEEE). *ISO/IEC (IEEE Std) 24765 : Systems and Software engineering - Vocabulary systems*. 2010 (cit. on p. 2).

[3] IEEE. "Adoption of ISO/IEC 15288:2002 Systems Engineering-System Life Cycle Processes". In: *IEEE Std 15288-2004 (Adoption of ISO/IEC Std 15288:2002)* (2005). DOI: 10.1109/IEEESTD.2005.96287 (cit. on p. 2).

[4] *PICMG® 3.0 Revision 3.0 AdvancedTCA Base Specification Errata Notification 3.0-R3.0-ERRATA001*. URL: https://cdsweb.cern.ch/record/1159877?ln=en (cit. on pp. 3, 7).

[5] R. Kopeliansky, W. Vandelli. *TDAQ Phase-II SoC user requirements document*. Tech. rep. AT2-D-ER-0001. URL: https://edms.cern.ch/document/2188675/5 (cit. on p. 3).

[6] R. Kopeliansky, W. Vandelli. *TDAQ subsystems - SoC Questionnaire - Summary*. Tech. rep. AT2-D-RG-0001. URL: https://edms.cern.ch/document/2176280/1 (cit. on pp. 6, 8).

[7] *CERN Computing Security Rules*. URL: https://security.web.cern.ch/security/rules/en/index.shtml (cit. on p. 8).

[8] P. Allport, D. Francis, G. Herten, N. Hessey, F. Lanni, M. Nessi, L. Pontecorvo. *Upgrade Organization*. Tech. rep. ATU-ORG-MM-0001, (EDMS: 1093133). URL: https://edms.cern.ch/document/1093133/5 (cit. on p. A.3).

Appendix A

Terms, Definitions and Glossary

List of Definitions

List of Requirements

List of Specifications

List of Recommendations

List of Remarks

List of Tables

List of Figures

A.1 Glossary

ATCA Advanced Telecommunications Architecture. 3, 4, 6, 7, 9, A.1, A.2

ATCN ATLAS Technical and Control Network. 4–6, 8, 9, A.1

ATLAS A Toroidal LHC Apparatus. i

DAQ Data Acquisition System. 9, 10

DCS Detector Control System. 6, 9

FSM Finite State Machine. 9, 10

I^2C I2C. 6, 9

IPMC IPMI Management Controller. 3, 7

LASP LAr Signal Processing Board. 4

LATS LAr Timing System. 4

LCS Local Control Station. 9

OPC UA Open Platform Communications Unified Architecture. 6, 9, A.1

OS Operating System. 6, 8

quasar Quick OPC UA Server Generation Framework. 9

RTM Rear Transition Module. 4

SoC System-on-Chip. i, v, 1, 3–10, A.1, A.2

SoM System-on-Module. 5, 6, A.2

TDAQ Trigger-DAQ. 3, 4, 6–9, A.2

TDR Technical Design Report. i

UPR Upgrade Project: see [8] for further details. i, 1, 3, 9