THE MEDAUSTRON ACCELERATOR CONTROL SYSTEM

J. Gutleber, CERN, Geneva, Switzerland
A. Brett, R. Moser, M. Marchhart, C. Torcato de Matos, EBG MedAustron, Austria
J. Dedic, Cosylab, Slowenia

Abstract

This paper presents the architecture and design of the MedAustron accelerator control system. This ion therapy and research facility is currently under construction in Wr. Neustadt, Austria. The accelerator and its control system are designed at CERN. This class of machine is characterized by rich sets of configuration data, real-time reconfiguration needs and high stability requirements. The machine is operated according to a pulse-to-pulse modulation scheme. Each beam cycle is described in terms of ion type, energy, beam dimensions, intensity and spill length. The control system is based on a multi-tier architecture with the aim to achieve a clear separation between front-end devices and their controllers. In-house developments cover a main timing system, a light-weight layer to standardize operation and communication of front-end controllers, fast and slow control of power converters and a procedure programming framework for automating high-level control and data analysis tasks.

INTRODUCTION

The MedAustron particle accelerator [1] is intended to deliver beams of various light ions for research and ion particle therapy. It features multiple ion sources, a Linac, a synchrotron and five beam lines including a proton-gantry (see Fig. 1). Upon request, virtual accelerators that represent beam paths from an ion source to an irradiation room supply beam cycles with pre-configured characteristics. A client system requests beam cycles with certain characteristics and the accelerator control system timely re-configures the front-end components. A request for a next beam cycle is issued while a cycle is generated (pipelined operation), requiring that the control system operates in soft real-time to keep dead-times between beam cycles low. Accelerator subsystems provide pulse-to-pulse modulation operation on a best effort basis, requiring a requesting client system to verify the characteristics of the delivered beam cycles.

Key features of the MedAustron control system are:

- Possibility to partition the accelerator into multiple, independent virtual accelerators and working sets to allow tuning of machine components and servicing concurrently with beam commissioning and operation.
- Optimized operation for pipelined pulse-to-pulse modulation supporting several hundreds of thousands of beam cycle configurations, being able to switch configuration in soft real-time without dead-time for cycle durations down to one second.
- High density remote power converter controller with sub-microsecond timing precision and value precision up to 24 bits operating at up to 50 KHz.
- Highly accurate and flexible real-time event distribution network as main timing system with 100 nanosecond GPS timestamping capabilities.
- Possibility to operate multiple machine partitions concurrently, supporting several modes and operation security levels.
- Role-based authentication and authorization.
- An RDBMS based system to store and organize configuration data and to create and trace versions of configuration data for operation.

This article gives an overview of the accelerator control system architecture and design decisions that govern its implementation.

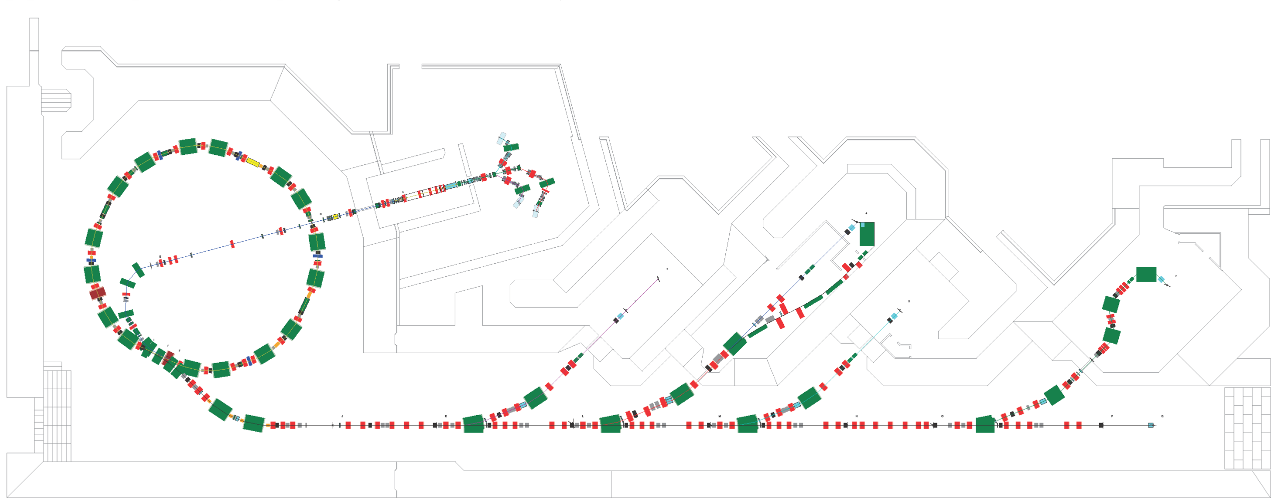

Figure 1: Layout of the MedAustron particle accelerator. Gantry beamline is indicated in bottom right part.

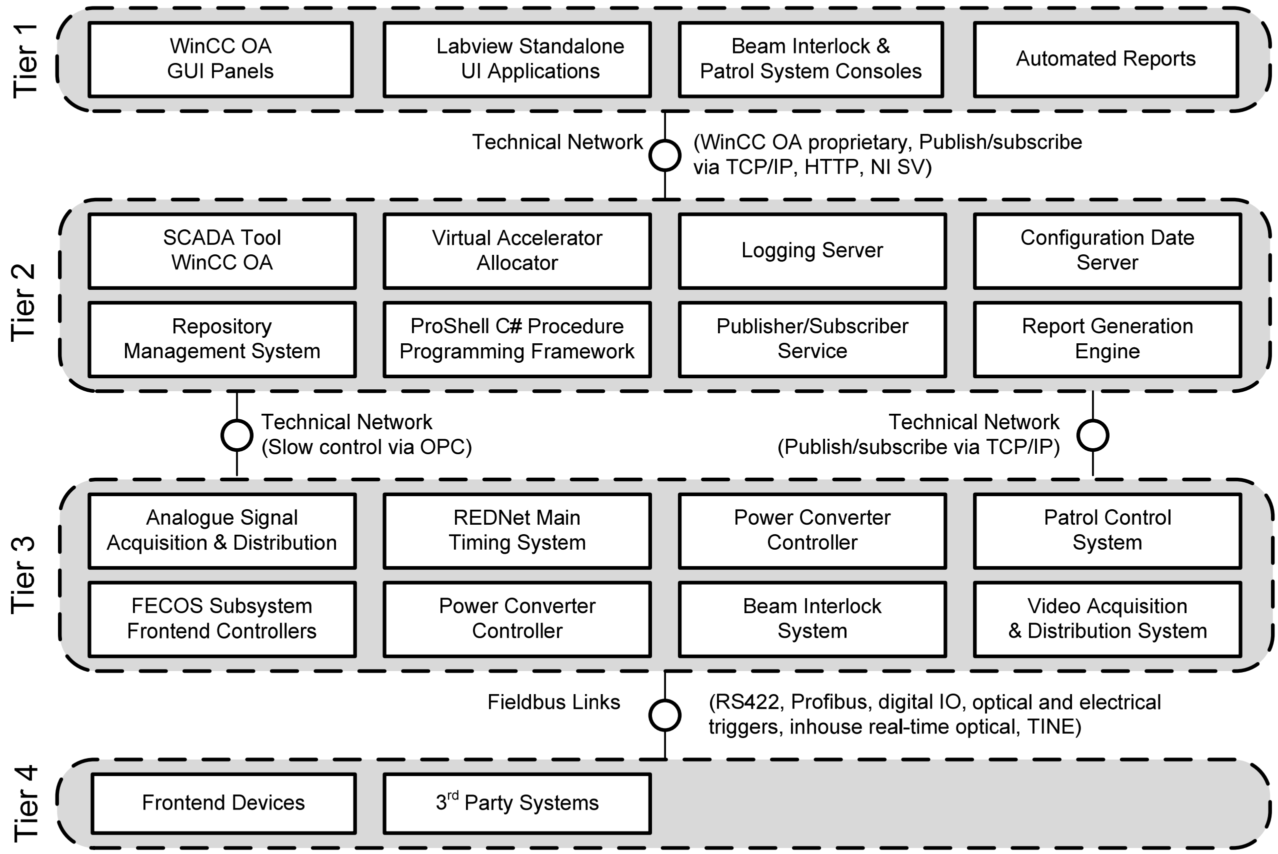

Figure 2: 4-tier accelerator control system architecture. Tier 4 comprises low-level front-end devices or third party control systems that must be integrated via a proxy front-end controller at tier 3. A single physical technical network permits communication among tiers in a secured network environment. Applications at tier 1 and 2 are virtualized.

ARCHITECTURE

Overview

The architecture (see Fig. 2) extends the industry best practice, 3-tier model [2,4] in accordance with [3]: (1) *presentation* tier, (2) *processing* tier, (3) *equipment* tier and (4) *frontend* tier. Components in separate tiers that are distributed over a number of processing devices communicate with each other through a dedicated Ethernet network. Communication between equipment tier and frontend tier may also be achieved through dedicated field-bus and custom links, depending on the imposed constraints. Subsequent sections describe each tier, starting with the frontend tier that is closest to the equipment under control.

Frontend Tier (T-4)

Front-end devices that control the actual accelerator hardware are categorized according to the project's work package organization:

The ion source supplier provides low-level API libraries for a COTS real-time computing platform. They serve implementing a front-end controller (FEC) at tier three that autonomously operates a single ion source. That system is replicated for each of the foreseen ion sources.

Conventional magnet temperature sensors open a circuit to indicate an over-temperature condition. A single PLC using the Siemens Safety Matrix tool [5] drives the circuits of all 262 magnets. This system interfaces also directly to all power converters to deliver machine protection interlock functions. Special magnets feature a separate, in-house developed, Siemens PLC based positioning and cooling monitoring system that interfaces directly to tier 2 via the S7 protocol over TCP/IP.

Power converters come in 2 different flavours: The first family is digitally controlled via RS422 lines and a purpose defined, UART based protocol for current control, slow controls and monitoring. All other power converters are treated as voltage sources that are controlled by in-house designed, analogue regulation boards. Slow controls for status monitoring and state changes are again performed via the specified RS422 line protocol. Each power converter features a trigger input to acquire measurements at dedicated points in time, synchronized within one microsecond. Hence, for slow-controls a fully homogeneous design has been achieved, including the power converters of special magnets.

Injector and synchrotron RF amplifiers feature supplier-provided Siemens PLCs for status monitoring. They interface directly to the SCADA tool at tier-2.

The injector low-level RF system features an embedded, supplier-provided computing device with few modifications to ease pulse-to-pulse modulation operation and cycle-based tagging of measurement data. This Tine-based system is interfaced via the Tine [6] TCP/IP application-level protocol and a custom-built real-time bus for generating stability and duty pulses.

The synchrotron low-level RF system is currently under development at CERN. Its communication interface remains yet to be defined.

Vacuum equipment controllers are self-contained, autonomously operating devices that feature serial interfaces and few digital IO points. They are directly interfaced to the SCADA tool at tier-2 via an OPC server that connects to several terminal servers and a few serial/Ethernet remote IO devices. The same concept applies to the beam interception devices.

Beam diagnostics devices are characterised by a rich set of different front-end electronics and interfaces. A description of the system at tier-4 goes beyond the scope of this article.

Equipment Tier (T-3)

This tier serves two purposes: first, it generates a uniform view of the diverse devices at tier-4 by introducing autonomously working front-end controllers (FEC) that all follow one design pattern. Secondly, the front-end controllers locally keep all configuration data for timely applying configuration changes to achieve efficient pulse-to-pulse modulation operation. Tier-4 devices are virtually represented at this tier either in physical FECs or directly in the SCADA tool at tier-2. FECs are National Instruments PXIe systems running the MS-Windows or Labview-RT operating systems, depending on subsystem specific constraints. FECs are programmed with Labview, using a light-weight layer called Front End Controller Operation System (FECOS) that shields programmers from dealing with communication details, providing design patterns for achieving efficient pulse-to-pulse modulation operation, standard commanding and configuration mechanisms.

Physical FECs communicate with a SCADA system at tier-2 via the OPC protocol. Direct graphical user interfaces for service purposes interface directly to FECs via the NI Shared Variable engine. For high-rate/high-bandwidth measurement data distribution, FECs use an application-level publisher/subscriber protocol based on NI's STM over TCP/IP. FECs may communicate with each other via that protocol to implement non real-time distributed applications or via dedicated low-level signals for real-time synchronization.

Notable physical FECs are the Power Converter Controller (PCC), the REDNet Main Timing System Generator (MTG), the ion source FEC, the injector and synchrotron LLRF FECs to interface to the respective low-level RF control systems.

The main timing system called REDNet [7] distributes events in real-time to synchronize the actions of the FECs. It processes "runs" consisting of lists of pre-defined beam cycles. A beam cycle is represented by a 64bit wide data structure that encodes possible beam characteristics. The MTG acts as single service access point via an application-level TCP/IP protocol for clients that have already successfully allocated a machine partition and that possess a universally unique session identifier (UUID) for this allocation. Before a cycle can be started, the cycle identifier is broadcast to all FECs via the MTG. Once a client has requested the start of a cycle, individual events are broadcast either via a dedicated, optical real-time network and/or the publisher/subscriber mechanism to the frontend applications. Timing system receiver cards generate real-time responses directly to devices via PXI/PXIe real-time bus lines or dedicated hardware interfaces using MRF's Universal IO [8].

The PCC is a PXIe based embedded systems hosting numerous FlexRIO FPGA cards. Each FPGA card can drive up to 6 power converters via an in-house developed, optical real-time link adapter for the NI FlexRIO platform (see Fig. 3). such high-density real-time control of all 260 power converters can be achieved.

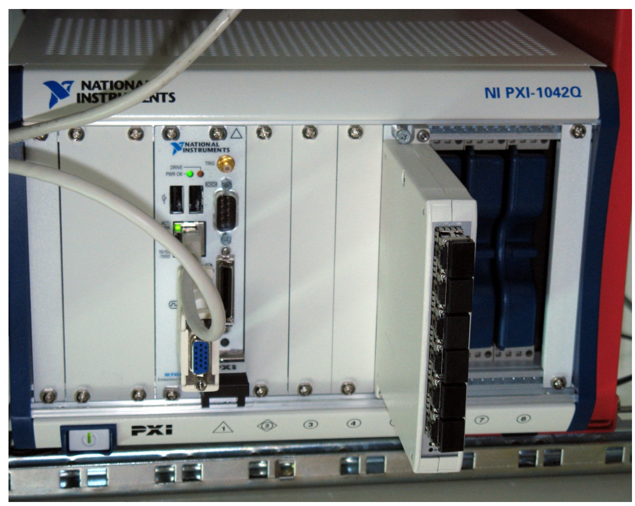

Figure 3: PCC test setup with 1 FlexRIO FPGA board and mounted 6-way optical real-time link adapter.

Finally, this tier hosts also an analogue signals acquisition and distribution system (SADS) that acts as a virtual oscilloscope for up to two concurrent users for about 100 analogue signals. This system is integrated with the MTS so that operators can easily correlate signal over time displays with distributed timing events.

Processing Tier (T-2)

The processing tier (T-2) configures and supervises equipment and frontend tiers. The tier runs supervisory procedures that include machine quality assurance and performance tuning programs. The following list summarizes functions that are located in this tier:

- virtualization of all physical devices as data point structures including a single, uniform state machine for each device that requires state-driven operation,

- creation of a uniform system image by representing all machine partitions, devices, and software components as data points,
- enforcing of role based authentication and access control,
- configuration of front-end controllers via OPC, HTTP and FTP protocols,
- persistent storage of system configuration and device parameters via an RDBMS,
- synchronization and regulation of access to shared resources,
- alarm and error handling, collection of log messages,
- distribution of on-line acquired process data,
- operation data logging and report generation,
- automated execution of supervisory control procedures to acquire machine and beam-related data, to perform quality assurance and performance tuning tasks and to generate new configuration data.

The core system at this tier is a distributed WinCC OA system from ETM. It is the only path to interact with devices during operation periods. Device configuration data may include simple or complex values. They may be static or cycle dependent. A cycle dependent configuration value is assigned a cycle mask and filter, (two 64bit key values). If a software component at FEC level receives the request for a beam cycle, the FECOS framework will look up configuration values based on the cycle mask and filter and provide them to the front-end application. At any time, configuration values can be inspected and modified via WinCC OA. Configuration data are either directly provided to FECs via the OPC protocol or are made available via a dedicated WEB server from which the FECOS executive picks the data.

A particular software component called Virtual Accelerator Allocator (VAA), written entirely in C# is the single entity that accepts requests for machine partition allocations and requests to start a run. This software component interacts solely with the WinCC OA system to achieve its task. Clients issue requests for machine partitions with QoS parameters such as required timing system precision, priority, run mode (medical, medical physics, machine physics or service), name of desired virtual accelerator or working set. If a request cannot be satisfied immediately, the VAA keeps it queued until the client either renews the request or the request times out. All applications in this tier art virtualized using the VMWare server virtualization infrastructure.

Presentation Tier (T-1)

User interfaces are implemented with WinCC OA panels. Enriched user experience such as accelerator synoptics and on-line presentation of measurement data are achieved via re-usable, in-house developed Qt widgets embedded in WinCC OA panels. Data obtained from MAPS can be directly processed these widgets. A panel navigator has been created to allow hierarchical organization of user interface panels, to perform panel-level access control for the defined user roles and to ensure that critical panels cannot be opened multiple times by multiple users at different locations. The CERN/JCOP defined security component [9] is used for implementing fine-grained role based access within panels. A colour and visual language optimized for low-ambient control room use has been devised from best practices in industry and EU norms. A claim mechanism implemented in cooperation between GUI panels and FECOS allows standalone expert user interfaces such as Labview panel applications to be used concurrently with WinCC OA panels to gain direct access to FECs for service purposes without jeopardizing operation safety.

SUMMARY

This article presented the architecture and selected architectural significant components of the MedAustron accelerator control system. This system is based on a 4-tier architecture with the aim to cleanly separate user interfaces, configuration, monitoring and system state processing services, real-time processing and proprietary and third-party supplied accelerator control devices. This control system is an evolution of the CNAO accelerator control system that has recently entered operation. A full vertical column of the system including user interfaces, the main timing system, the configuration repository and the power converter controller has been demonstrated and has proven the validity of the approach. Hence, this architecture may serve as a blueprint for future particle accelerators with comparable operation requirements.

REFERENCES

[1] M. Benedikt and A. Wrulich, “MedAustron – Project overview and status”, Eur. Phys. J. Plus (2011) 126:69
[2] W. W. Eckerson, “Three Tier Client/Server Architecture: Achieving Scalability, Performance and Efficiency in Client Server Applications”, Open Information Systems (1995) 1:10
[3] H. Winick, “Synchrotron Radiation Sources: A Primer”, (1994), pp. 218
[4] L Evans and P Bryant (eds.) “The CERN Large Hadron Collider: Accelerator and Experiments”, JINST (2008) 3, chapter 9, pp.98
[5] H. Pavetits, “Evaluation of the Siemens Safety-Matrix tool with safety PLC for a person and machine protection system of a particle accelerator”, Bachelor thesis FH Wiener Neustadt (2011).
[6] P. Duval and S. Herb, “The TINE Control System Protocol: How to Achieve High Scalability and Performance”, in Proc. PCaPAC 2010, Oct. 5-8 2010.
[7] J. Dedic et al., “Timing System for MedAustron Based on Off-The-Shelf MRF Transport Layer”, in Proc. IPAC 2011.
[8] MRF, Universal IO Specification, last seen at http://www.mrf.fi/index.php/universal-io-modules.
[9] P. Golonka, M. Gonzales-Berges, “Integrated Access Control For PVSS-Based SCADA Systems at CERN”, in Proc. ICALEPCS 2009.