INTRODUCING BIG DATA ANALYSIS IN A PROTON THERAPY FACILITY TO REDUCE TECHNICAL DOWNTIME

P. Fernandez Carmona[†], Z. Chowdhuri, S. G. Ebner, F. Gagnon-Moisan, M. Grossmann,
J. Snuverink, D. C. Weber, Paul Scherrer Institut, 5232 Villigen PSI, Switzerland

Abstract

At the Centre for Proton Therapy of the Paul Scherrer Institute (PSI) about 450 cancer patients are treated yearly using accelerated protons in three treatment areas. The facility is active since 1984 and for each patient we keep detailed log files containing machine measurements during each fraction of the treatment, which we analyse daily to guarantee dose and position values within the prescribed tolerances. Furthermore, each control and safety system generates textual log files as well as periodic measurements such as pressure, temperature, beam intensity, magnetic fields or reaction time of components. This adds up currently to approximately 5 GB per day. Downtime of the facility is both inconvenient for patients and staff, as well as financially relevant. This article describes how we have extended our data analysis strategies using machine archived parameters and online measurements to understand interdependencies, to perform preventive maintenance of ageing components and to optimize processes. We have chosen Python to interface, structure and analyse the different data sources in a standardized manner. The online channels have been accessed via an EPICS archiver.

INTRODUCTION

The Paul Scherrer Institute (PSI) in Switzerland started treating tumours using accelerated protons in 1984. Since then more than 9000 cancer patients have been treated at its fixed beamline for eye irradiation and three gantries.

The facility continuously produces large amounts of data originating from the particle accelerator, beamlines and control and safety systems needed for the dose delivery. Part of this data, especially everything directly related to patient irradiation, gets stored for online and future analysis. Patient treatment log files get analysed daily throughout the treatment duration to guarantee the delivery standards. Some sensor data is used for forensics after particular machine malfunctions, and other sources are used as surrogates to estimate when a part needs to be replaced due to ageing or wear.

Downtime due to unexpected failure of components is both inconvenient for patients and personnel, and expensive in financial terms due to unused clinical resources and paused revenue. With the goal in mind to reduce downtime and to improve preventive maintenance we started a pilot to introduce structured big data analysis techniques in our facility using the extensive available data. In the present paper we will describe the data available, our steps to classify it and process it as well as the first promising results.

DATA SOURCES

PROSCAN, the facility dedicated to treating patients at PSI, is a complex set of interconnected but largely independent and heterogeneous subsystems. There are several control, safety and monitoring systems generating status data at different rates and formats. The first part of this work was to list and attempt to classify all the available data for its future analysis.

Machine Data

A superconducting cyclotron provides protons to the treatment areas by means of 5 beam lines. These contain 17 steering magnets, 45 quadrupoles, 8 deflecting dipoles, 47 beam monitors, 14 beam blockers and several other auxiliary elements each with its set point and actual status values. Most of these values are permanently available, some on request as they are beam-disrupting.

The usage and sharing of the beam across the different areas also gets monitored and archived. It is always possible to know at any given time which area had mastership, that is, control and access to the beam, and at which energy and intensity. This can be useful for forensics, but also for statistics, see Figure 1, and to identify potential inefficiencies.
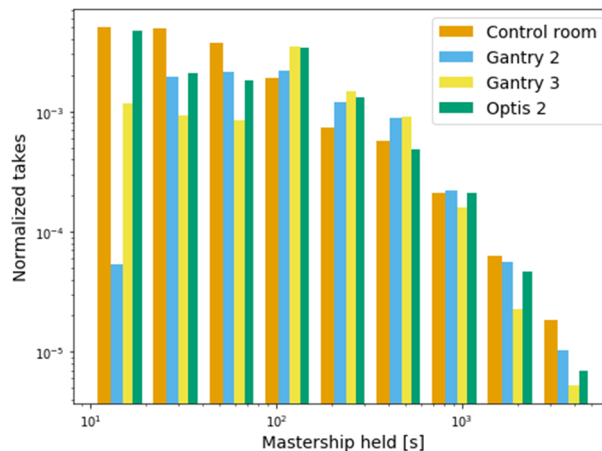

Figure 1: Histogram of mastership duration per area in 2019.

Additionally each treatment area stores dosimetry-relevant parameters not limited to but including humidity, pressure and temperature at different locations. The control systems also store technical values from electronics and sensors, namely supply voltage, power consumption or internal temperature.

Patient Data

The treatment of a patient is typically divided in fractions, out of which there are between 4 and 37, depending on the indication. Each fraction includes a prescribed dose distributed into one or more fields to be sequentially applied. At PSI gantries we use Proton Beam Scanning (PBS) technology, a superposition of single pencil beams that added together cover the full volume of the tumour achieving a high conformity [1]. Each beam application is divided into spots, a predefined number of protons at a given energy and location. The Therapy Control System (TCS) requests the right machine setting for each spot and verifies its correctness in a short dead time between spots. A typical field in our Gantry 2 contains 20'000 spots, each of which is individually logged. The logs contain information of position and dose, as well as the settings of all the involved beamline elements. These files are stored in clear text and in a Structured Query Language (SQL) database and kept indefinitely.

Safety Systems

The safety of the dose delivery in each treatment area is guaranteed by a Patient Safety System (PaSS) [2]. This is a Field Programmable Gate Array (FPGA) based interlock system in place to monitor the facility and to interrupt the treatment whenever any subsystem detects a hazardous condition. There are three levels of interlock severity, which deflects the beam, pauses the acceleration of protons by reducing the High Frequency (HF) power or switches off the accelerator respectively by shutting the HF.

It is connected to beam monitors, beam blockers, the control system and other monitoring devices, whose status samples at 1 MHz. Several supervision functions were built in, including measuring the reaction times of all connected elements. The status of all signals and interlock outputs, as well as all measurements are published via EPICS [3] with 0.1 to 1 s resolution and archived for later use. Additionally after each interlock event, a log file with the sequence of events with a 1 μs time resolution is generated in text format.

Interlocks are, together with device failures, the main sources of downtime in the facility. It is for this reason that understanding them, and eventually being able to predict them is of the upmost importance.

METHODS

Before the start of this project some sort of data analysis was already performed in individual systems for different purposes. It is the case of the patient Quality Assurance (QA) log file analysis [4]. Other analyses were made only for research projects or as part of a forensic investigation after an incident.

At first we analysed the facility topology and talked to the different experts to obtain a list of all the available data, its format and how to access it. We calculated that 5 GB of data are generated on an average treatment day. This gets down-sampled for daily storage, and again yearly, to save space.

Unified Access Architecture

We chose Python as the programming language to process the data as it provides excellent analysis and visualization tools, namely Pandas [5] and Matplotlib [6]. The goal was to provide a unified access tool to all the heterogeneous data sources.

PSI has a centralized network time protocol (NTP) server to which most systems are synchronized. For the data sources that are not synchronized we attempted to find the difference to the central NTP server and apply it as an offset to the corresponding data. For the missing data points we chose either interpolation or extrapolation of the last known value, depending on its nature.

Some data sources have a continuous nature, like temperature or beam intensity. Others however are discrete, such as the interlock events. In order to be able to look for correlations between them we created variables such as $\frac{interlock}{minute}$. In Figure 2 we demonstrate the unified access to related variables of different formats and sources: The instantaneous current out of the accelerator from the machine archiver and the interlock events due to overcurrent detected, from the safety system data.
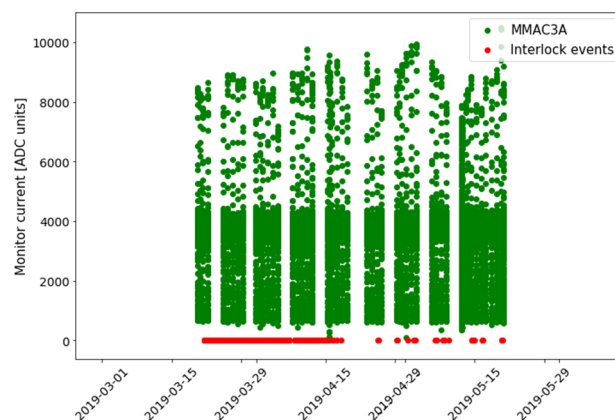

Figure 2 Display of accelerator output current and the interlocks that are produced when it is too high for Gantry 2.

EPICS and Archiver

Most of the hundreds of magnets, monitors, collimators and diverse sensors distributed along the accelerator, beamlines and gantries publish regularly their status via EPICS. This data is online available at different data rates and can be used for online visualization, accessible to the whole laboratory. In addition to this there are a short term archiver which keeps the full resolution of data for one year and a long term archiver with down sampled data.

The archived data can be accessed using a python API and it is time stamped with a unified time reference.

Patient Data Analysis

All machine log files from fields that are applied to patients in G2 and O2 are analysed daily using Matlab™ software tools that were developed in-house. Parameters such

as dose and spot-position deviations, and number of interruptions due to interlocks, among others, are evaluated. A report detailing the results is generated per patient per day, and in addition, summary spreadsheet files are appended with values of particular interest.

We have developed a python interface to access these files and to make its content available in a Pandas dataframe format, with time stamps synchronized to the rest of the facility data. This includes the precise status of each beam line element at the time any patient was being treated. Figure 3 shows the daily average spot position deviation for all patients treated in Gantry 2 in the current year, which is well below the 0.5 mm allowed tolerance.
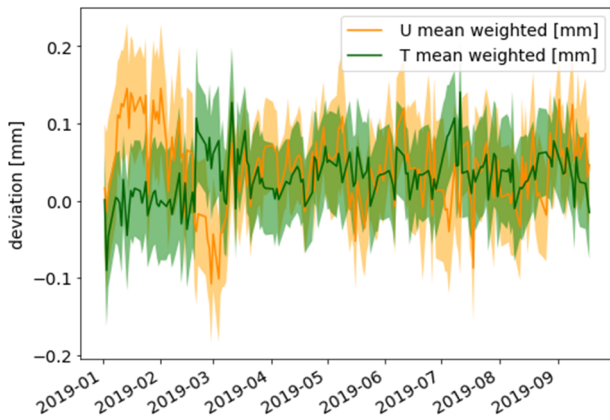

Figure 3: Daily average position residuals ±σ interval over a year for both transversal directions.

RESULTS

The goal of this project was to be a proof of concept to show the benefits of big data analysis in our small size facility. For this reason we chose a few representative use cases to put the first efforts in.

Usage Overview

As seen in Figure 4 the distribution and type of interlocks is clearly visible from the logs from Gantry 2 safety system. Low level interlocks are uniformly distributed along the week, while there seems to be an unproportioned number of high severity interlocks on Thursdays. This is because of quality assurance tests of the emergency buttons are typically scheduled on these days.

From the mastership allocation we can process and display area usage overview, as seen in Figure 5. The granularity is configurable according to the needs of the visualization and could be extended with number of patients treated or interruptions due to interlocks.

Prediction Models for Preventive Maintenance

AMAKI, the deflecting magnet which switches the beam on and off at the output of the cyclotron is a central facility element. For this reason its performance is permanently monitored using magnetic sensors. If the magnet does not switch off in less than 200 µs an interlock is triggered, halting the acceleration of protons and placing beam blockers

in the beam path. The sensors are located inside the accelerator's bunker and therefore suffer continuous radiation
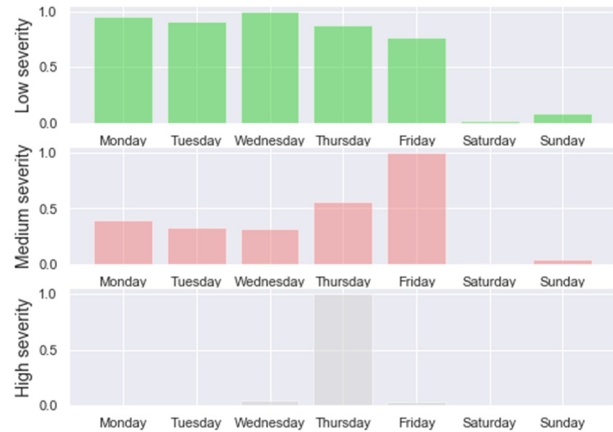

Figure 4: Distribution of the three levels of interlocks per week day in 2019 in Gantry 2.
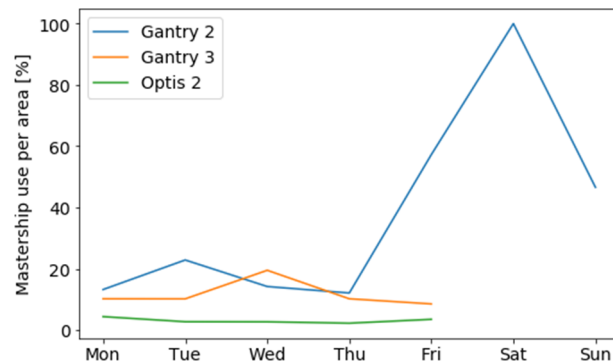

Figure 5: Treatment beam sharing time in a week where Gantry 2 was used for experiments during the weekend.

damage which makes them slower, causing false alarm interlocks. Its replacement requires manual intervention of experts inside the bunker.

Changing the sensor when it unexpectedly fails causes a downtime of 4 hours, half a treatment day. On the other hand replacing the sensor too often leads to unnecessary dose to personnel.

The introduction of the latest PaSS upgrade made available a precise measurement of the AMAKI switching time, which is both published via EPICS and archived. In Figure 6 one can see a clear pattern of the performance degradation of the sensor. It is evident that the replacement frequency is sub optimal.

We created a regression model to estimate the average switching time in the future, based on all the measurements since the last sensor exchange. To this we can add a margin to include the past time variability with 5σ and project when it will cross the interlock line (accounting for the measurement cable latency). In Figure 7 one can see the prediction for the next interlock, and how we can recommend a replacement after 8 weeks instead of the currently 6 weeks planned.
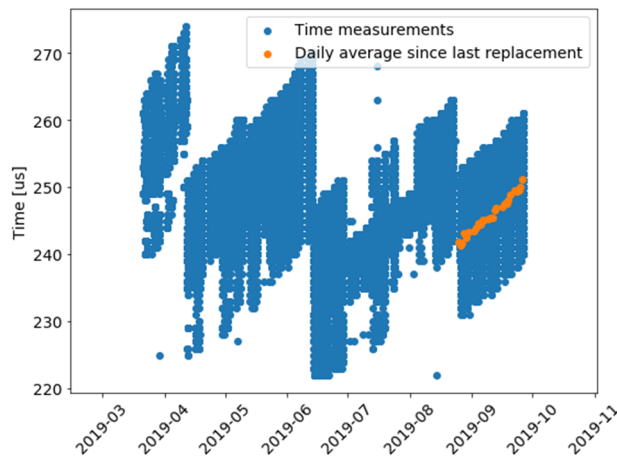
Figure 6: Switching time measurements from AMAKI magnetic sensors since data started being recorded.
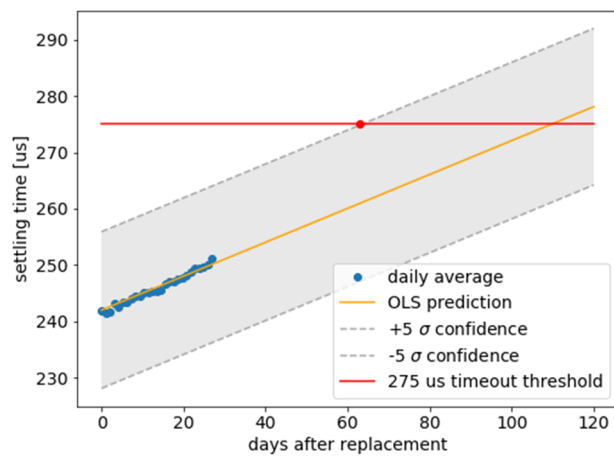

Figure 7: Linear regression model from AMAKI magnetic sensor switching times, projecting the future failing point.

Finding Interdependencies

We have started analysing the interdependencies of all the available data with some basic correlations. Figure 8 displays the readings from the slow control system archived data for temperature sensors distributed along Gantry 2. One can see what sensors are close to each other, as their temperature is similar and their correlation high. The same analysis can be later applied for less evidently related variables.

SUMMARY

We have described the main different data sources at PSI's Centre for Proton Therapy: Machine status, patient logs and from the safety systems. We then presented the approach to classify and make all data available for analysis in a unified way and with a common time reference. Later we discussed how we started analysing the available data and showed the first results, which however reduced suggest a great potential for better understanding the facility, optimizing its usage and planning preventive maintenance to reduce downtime.
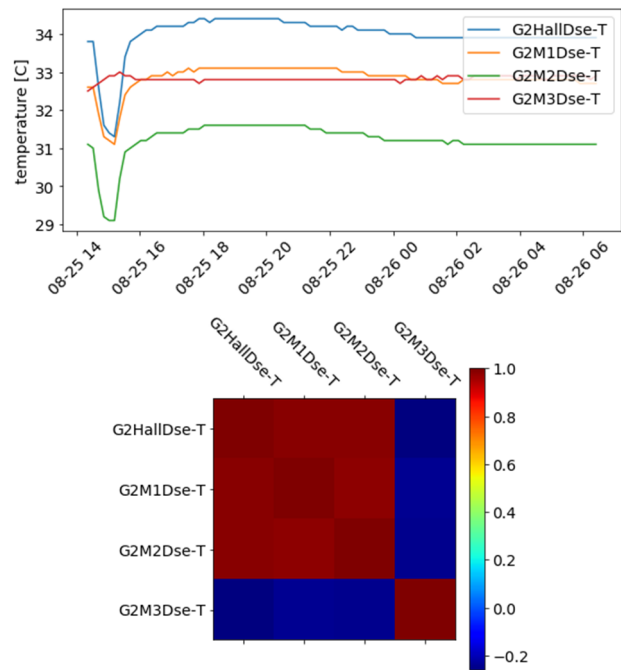

Figure 8: Temperature sensors distributed in Gantry 2 area (above) and its cross correlation over a day (below).

FUTURE WORK

The presented project was a proof of concept of exploratory nature and we plan to further develop the analysis of the facility data in two major lines: First the analysis and visualization need to be improved with more advanced techniques. Secondly we will create prediction models for early failure detection and eventually train an artificial intelligence expert system to monitor online status variables of the facility and suggest small correction interventions.

REFERENCES

[1] A. J. Lomax *et al.*, "Treatment planning and verication of proton therapy using spot scanning: Initial experiences", *Med. Phys.*, vol. 31, no. 11, pp. 3150-3157, 2004. doi:10.1118/1.1779371

[2] P. Fernandez Carmona, M. Eichin, M. Grossmann, E. Johansen, A. Mayor and H. A. Regele, "Reusable patient safety system framework for the Proton Therapy Centre at PSI", in *Proc. 15th Int. Conf. on Accelerator and Large Experimental Control Systems (ICALEPCS'15)*, Melbourne, Australia, Oct. 2015, pp. 549-553. doi:10.18429/JACoW-ICALEPCS2015-TUC3004

[3] EPICS, https://epics-controls.org

[4] D. Scandurra *et al.*, "Assessing the quality of proton PBS treatment delivery using machine log files: comprehensive analysis of clinical treatments delivered at PSI Gantry 2", *Phys. Med. Biol.*, vol. 61, no. 3, pp. 1171-1181, 2016. doi:10.1088/0031-9155/61/3/1171

[5] Python Data Analysis Library, https://pandas.pydata.org

[6] Matplotlib, https://matplotlib.org