PAUL SCHERRER INSTITUT

PSI

Michele Brambilla :: Software Engineer :: Paul Scherrer Institute

Data streaming at SINQ

NOBUGS 2022 — New Opportunities for Better User Group Software

Motivations

Until very recently standard data-acquisition electronics at neutron research facilities have saved data as a serie of bins in time and detector position (pixels): **histogramming mode**

Limiting factors

- Modest neutron flux
- Speed of acquisition electronics, network speed and disk space
- Lack of/outdated DAQ paradigms (histogramming)

During histogramming all **time information is lost**

New paradigm: **event-based** recording

For each neutron store
- pixel position
- timestamp

Each neutron can be correlated with state of instrument and sample environment (requires a data acquisition clock)

This enables **dynamic experiments**

PAUL SCHERRER INSTITUT

BrightnESS

"*The BrightnESS programme is designed to ensure that **key challenges are met** in order to build an ESS that can deliver **high-impact scientific** and **technological** knowledge*."

- WP 5: maximise the scientific output of the ESS by enabling **live** (real time) **processing** of the data taken on ESS instruments.
    - WP 5.3: develop the software to **aggregate** the sources of data into one data stream and make it available so that other software packages can use the data for processing.

Requirements:

- scalability
- reliability
- redundancy
- handling of clients and sources connecting and disconnecting

PAUL SCHERRER INSTITUT

PSI

The outcome of BrightnESS has the following design

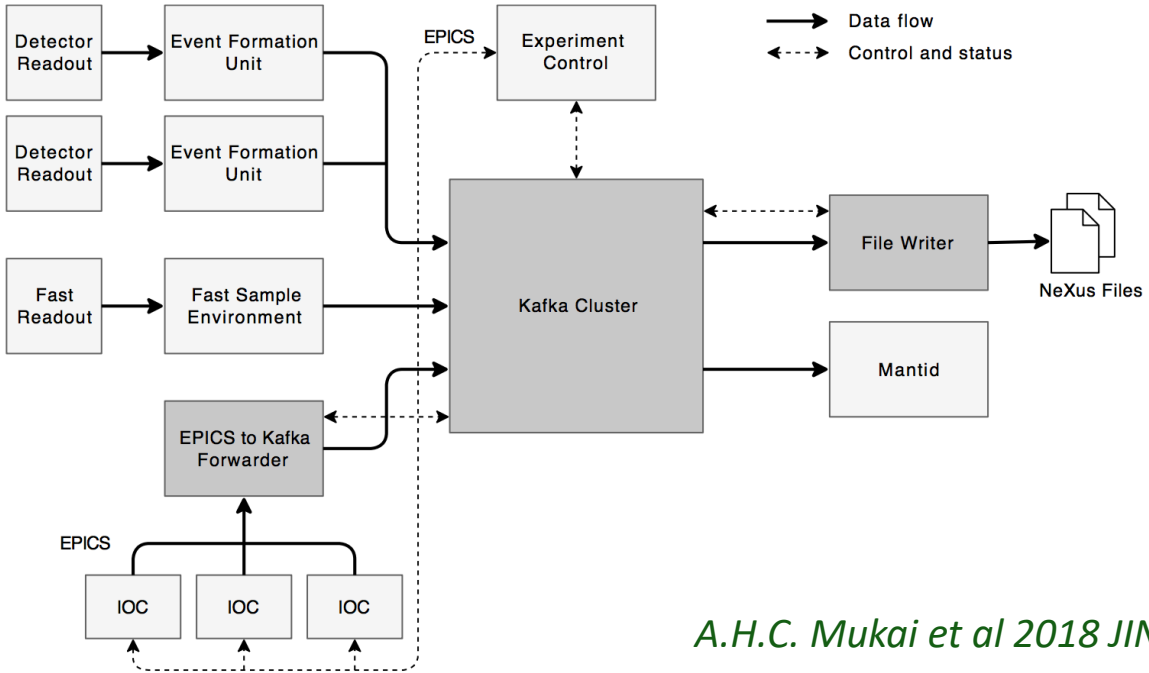

A.H.C. Mukai et al 2018 JINST 13 T10001

PAUL SCHERRER INSTITUT

Apache Kafka

Kafka (https://kafka.apache.org/) is a pub-sub communication system. Key concepts:

- a stream of messages of a particular type is defined by a **topic**

- **producer** can publish messages to a topic

- the published messages are then stored at a set of servers called **brokers**

- **consumer** can subscribe to one or more topics from the brokers, and consume the subscribed messages by pulling data from the brokers.

These functionalities are “provided in a distributed, highly scalable, elastic, fault-tolerant, and secure manner”.

PAUL SCHERRER INSTITUT

Other actors

EPICS: set of software tools and applications which provide a software infrastructure for use in building distributed control systems to operate devices
(https://epics-controls.org/)

Google flatbuffers: cross platform serialisation library
(https://github.com/google/flatbuffers)

Mantid: framework that supports high-performance computing and visualisation of materials science data
(https://www.mantidproject.org/)

PAUL SCHERRER INSTITUT

event-formation-unit: processing of neutron detector event data into neutron events and publish to Kafka
(https://github.com/ess-dmsc/event-formation-unit)

kafka-to-nexus: writes NeXus files from experiment data streamed through Kafka
(https://github.com/ess-dmsc/kafka-to-nexus)

forwarder: forwards EPICS PVs to Kafka
(https://github.com/ess-dmsc/forwarder)

just-bin-it: lightweight program for histogramming neutron event data
(https://github.com/ess-dmsc/just-bin-it)

Benchmarks

Tests performed on Intel Xeon E5-2690V4, 256 GB, GPFS via 4x Infiniband FDR

Producer Throughput vs Message Size

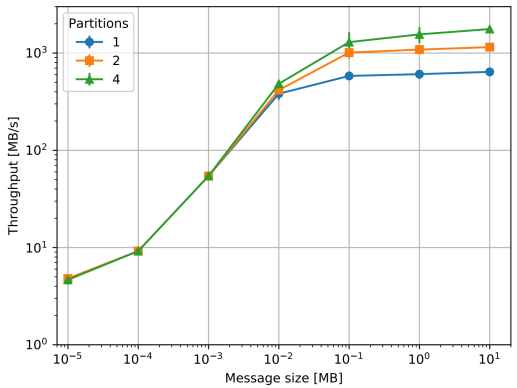

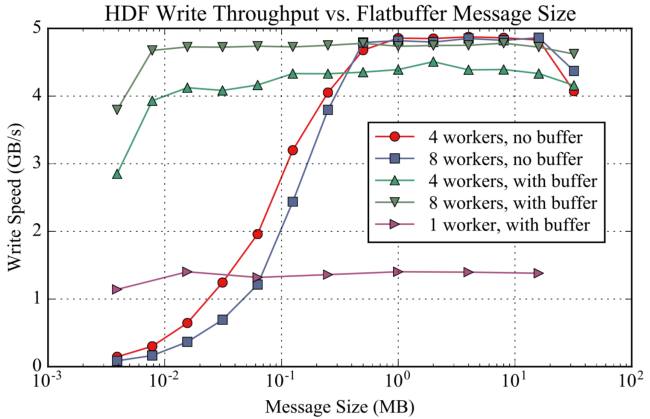

[D.Werder et al, JACoW-ICALEPCS2017-THPHA167]

Pipeline Throughput vs Number of brokers

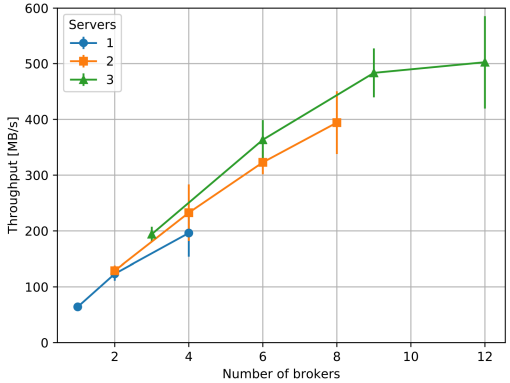

Implementation at SINQ

Two instruments at SINQ are capable of event-stream: **AMOR** and **DMC**

One (+ 1) instrument supports "2nd generation" detector electronics and is in principle compatible: **CAMEA** (+ **SANS-LLB**)

- **Redpanda** (https://redpanda.com) replaces Kafka (implements the Kafka API)

- there is no **global timing system**. We rely on NTP (PSI servers)

- sample environment uses **SECoP**, no direct connection to ~~Kafka~~ Redpanda

AMOR is **in production** since 2018 (currently stopped due to rebuilding)

DMC is **in production** since December 2021

Architecture at PSI

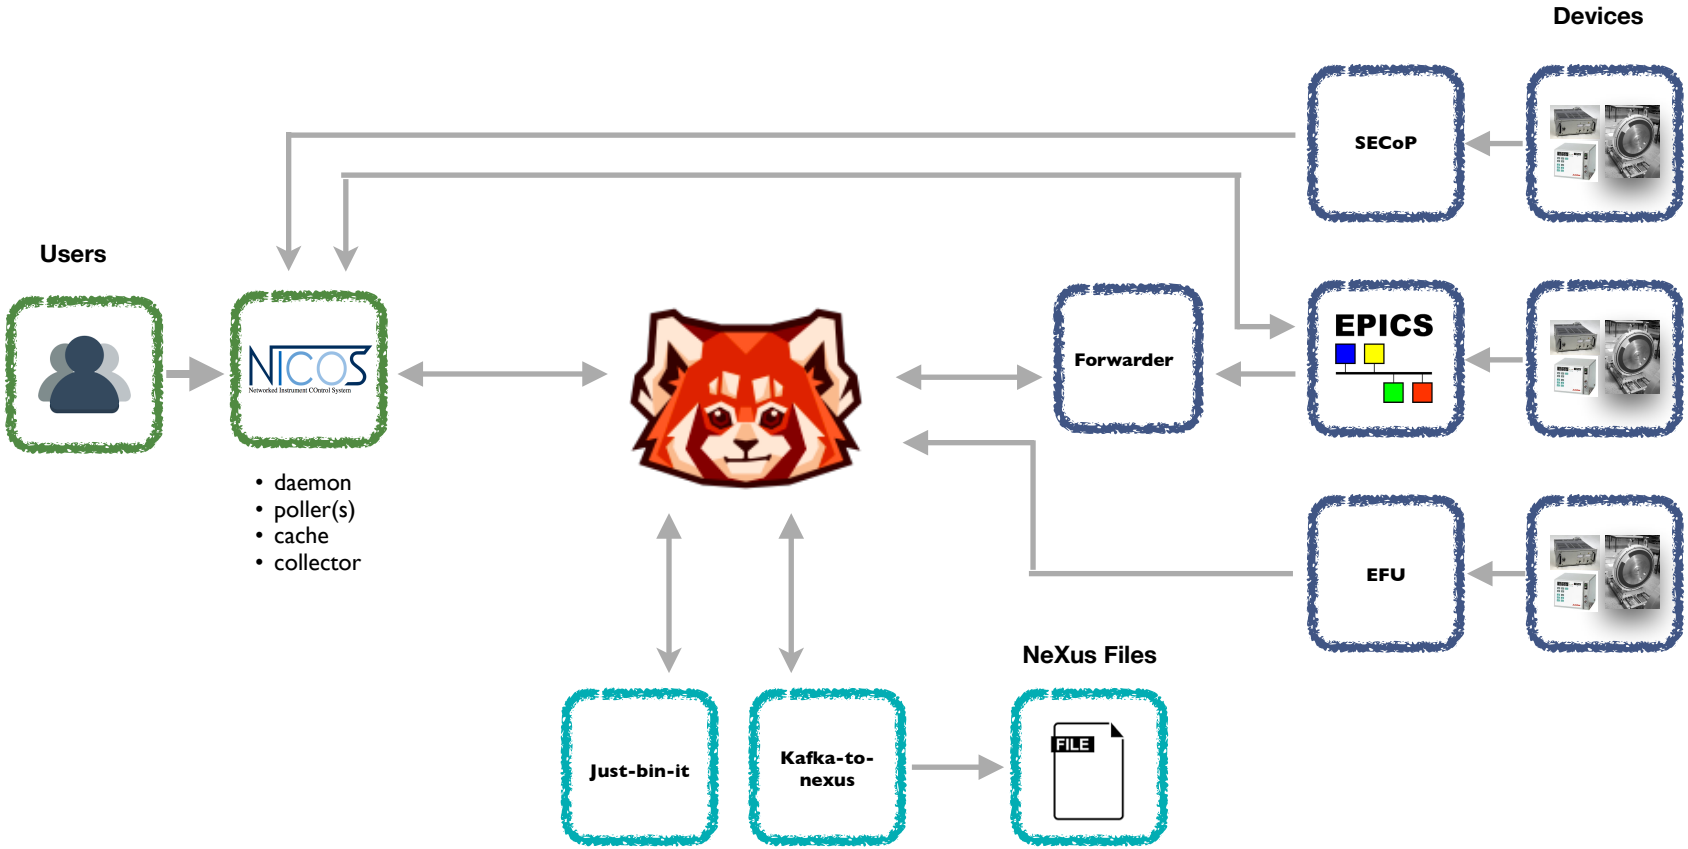

For a developer data looks like

From the point of view of the user

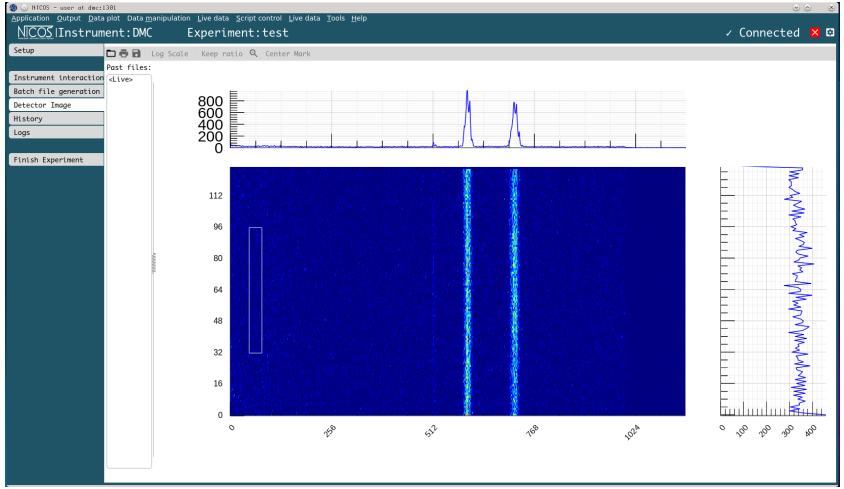

Scattering patterns @ DMC

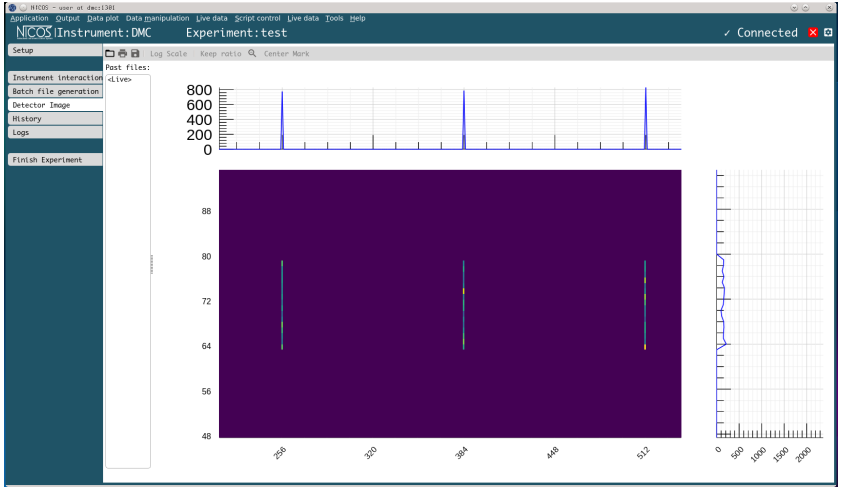

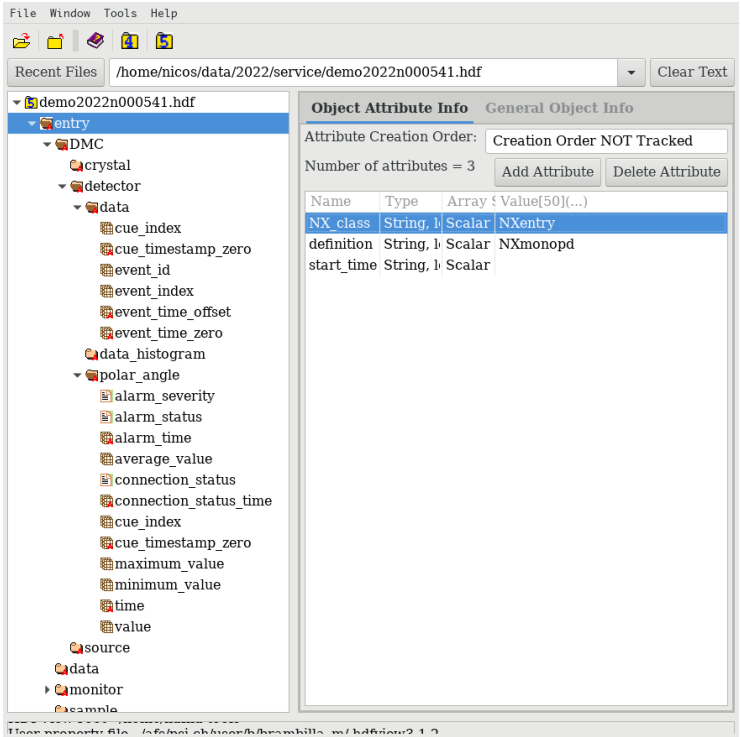

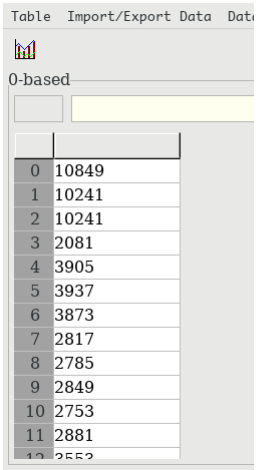

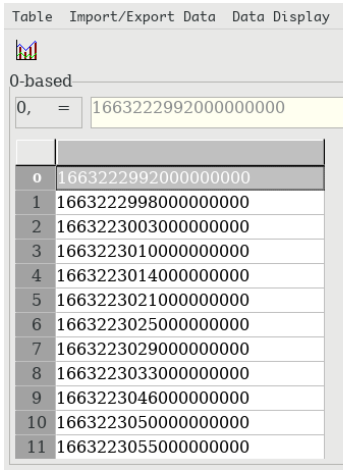

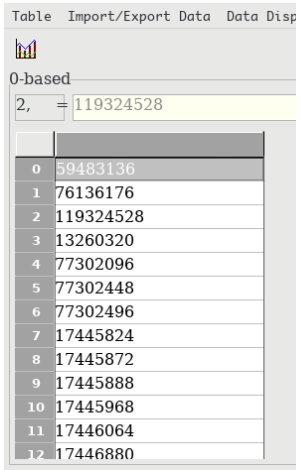

- the “hit” position are recorded as 1-D pixel array

- the timestamp/ToF is reconstructed from a time_zero and a time_offset

Conclusion

As member of the BrighnESS collaboration PSI contributed to the design and implementation of a new paradigm for data acquisition and storage

Thanks to the SINQ upgrade program and new detectors we are able to store experimental data in the form of a data stream and store it to a NeXus file

To achieve a complete event-based experiment we still miss

- hardware for event-based sample environment
- a global time system to ensure synchronisation of timestamp with high resolution