



Signal Monitoring Project

Use of Apache Airflow for Historical Data Collection and Signal Monitoring Execution



Kasper Andersen, Zinur Charifoulline, Per Hagen, Michał Maciejewski, Christoph Obermair, Sivert Sagmo, Arjan Verweij
TE-MPE-PE

with technical help from:

TE-MPE-MS: Thibaud Buffet, Jean-Christophe Garnier, Markus Zerlauth

IT-DB-SAS: Prasanth Kothuri, Piotr Mrówczyński

BE-CO-APS: Grigorios Avgitidis, Krzysztof Krynicki, Marcin Sobieszek, Jakub Woźniak, Bartek Urbaniec

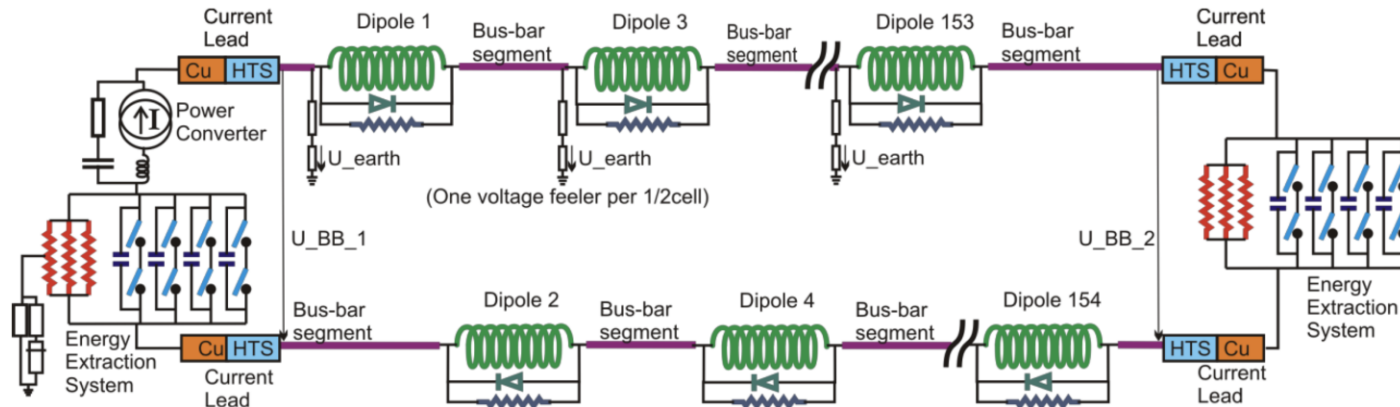


<https://cern.ch/sigmon>

We develop a general-purpose framework for monitoring:

1. Superconducting magnets and busbars
2. Circuit and magnet protection systems
3. Grounding networks
4. Current leads
5. ...

... with potential to extend to any other system with logged signals in PM and NXCALS* (e.g., Power Converters, Cryogenics Equipment, **Beam Instrumentation**).



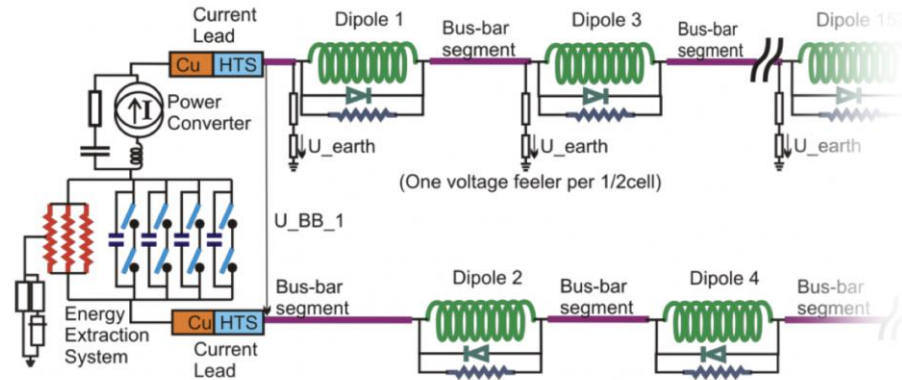
*NXCALS introduces a **paradigm shift** from local to cluster computing
The data should be processed where it is stored, i.e., on the cluster



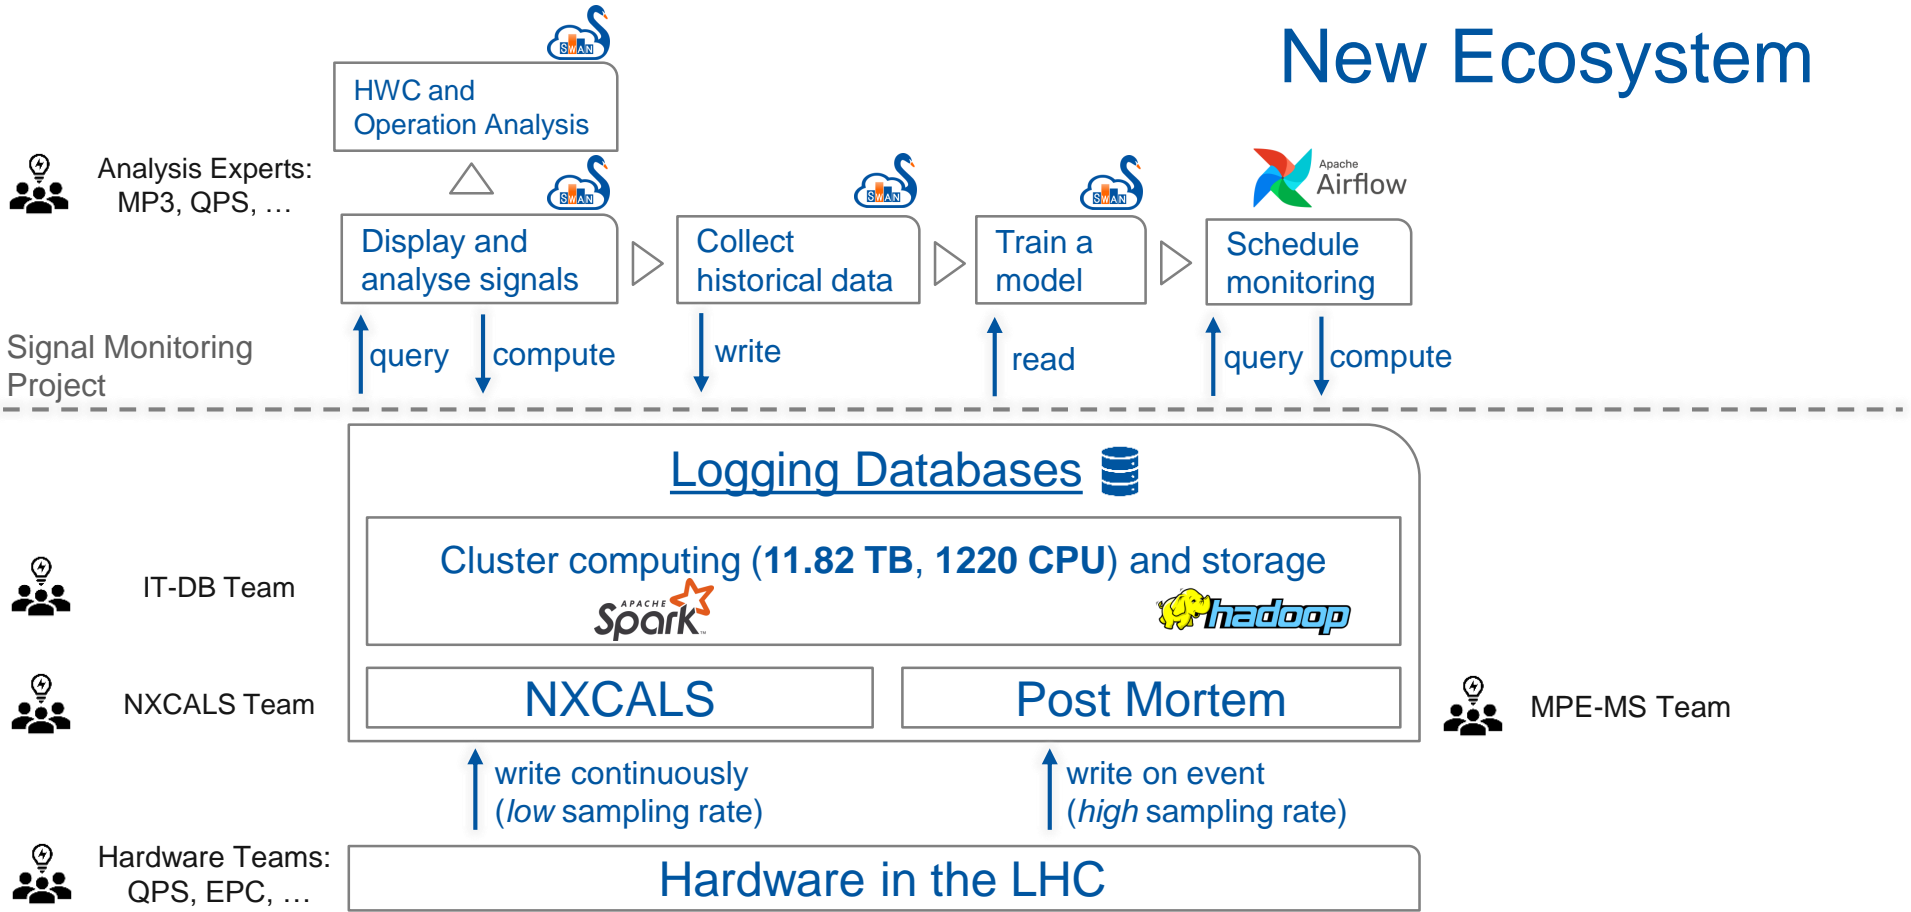
Requirements

- Many signals vary strongly during the operation cycle (ramp, FPA, flat-top, etc).
- Monitored signals should be compared to warning or alarm thresholds.
- Exceeding these thresholds should trigger an e-mail.
- User-friendly interface is needed to define signals, type of monitoring, time intervals, warning/alarm levels, e-mail address person to be notified.

→ Flexible triggering process
 → Intuitive user-interface



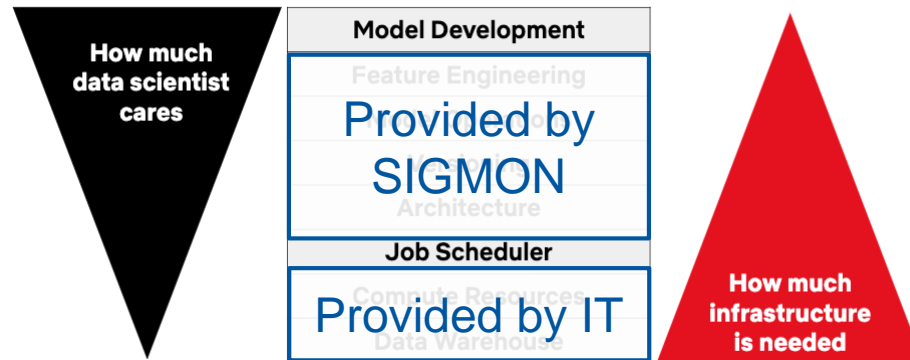
New Ecosystem



We aim at a **coherent approach** over all circuits, all systems, and *all* types of analysis
We use **python** due to the wealth of libraries for data analysis and modelling
We employ **SWAN notebooks** for development and communication across teams

Motivation

*The **new cluster** computing ecosystem and a **growing interest** in data-driven models call for a **standardized workflow and infrastructure** for data collection, modelling, and signal monitoring scheduling.*



Complex data workflows contribute to reproducibility crisis in science, Stanford scientists say

The main concerning takeaway from our study is that, given exactly the same data and the same hypotheses, different teams of researchers came to very different conclusions

- Russell Poldrack

While worrisome, Poldrack said the findings can help researchers assess and improve the quality of their data analyses moving forward. **Potential solutions** include ensuring that data is analyzed **in multiple ways**, as well as making data analysis workflows **transparent and openly shared** among researchers.



Outline

1. Signal Monitoring
2. Apache Airflow
3. Applications
4. Workflows
5. Summary

Signal Monitoring Workflow

Exploration → Data Collection → Modelling → Monitoring

Exploration – getting the signal features *right*

Creation of a notebook to explore a signal and compute characteristic features

	Feature 1	Feature 2	Feature ...	Feature n
event 1	0.078	980	...	10.4

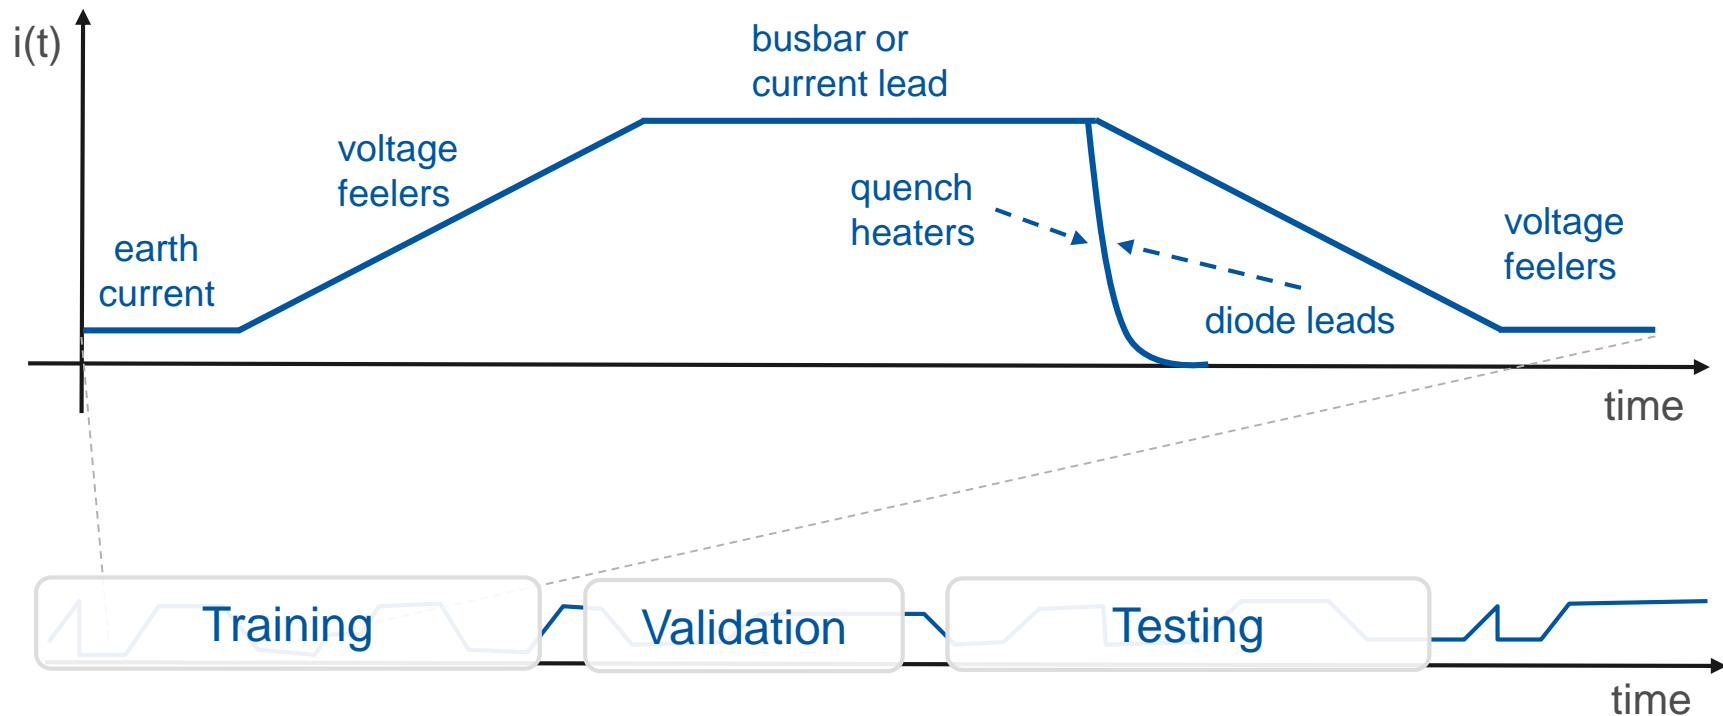
Data Collection – getting the *right* signal features

Execution of a notebook over past operation to collect data for numerical models

	Feature 1	Feature 2	Feature ...	Feature n
event 1	0.078	980	...	10.4
event 2	0.081	995	...	9.8
event
event m	.08	1000	...	10.1

Signal Monitoring Workflow


Exploration → Data Collection → **Modelling** → Monitoring



Data-driven models encode historical data in a compact way.

Signal Monitoring Workflow

Exploration → Data Collection → Modelling → **Monitoring**



Automatic execution of monitoring application depends on the state of operation:

- triggered by PM events (PC, QH, MAGNET) and AccTesting (test duration, circuit)
- triggered by change in the beam mode (GND, COLDBB)
- scheduled in regular intervals, e.g., every hour (DFB)

→ Our scheduler (Apache Airflow) can be triggered on demand through REST API

Limitations of SWAN

SWAN is a great tool for signal exploration, however:

- the duration of a user session is limited (data collection takes time)
- connection sometimes breaks
- broken analysis is not automatically restarted
- analysis can not be triggered on demand

We need a dedicated solution for scheduling of historical data collection and monitoring. Development of an in-house tool is beyond the scope and capabilities of our project.

A Solution - Cron

A standard solution for scheduling analysis jobs is cron.

IT department supports cron as a scheduler system.

However, it does not offer any user interface (a terminal application) and requires use of a special scripting language.

Each line of a crontab file represents a job, and looks like this:

```
# _____ minute (0 - 59)
# _____ hour (0 - 23)
# _____ day of the month (1 - 31)
# _____ month (1 - 12)
# _____ day of the week (0 - 6) (Sunday to Saturday;
#                               7 is also Sunday on some systems)
#
#
# * * * * * <command to execute>
```



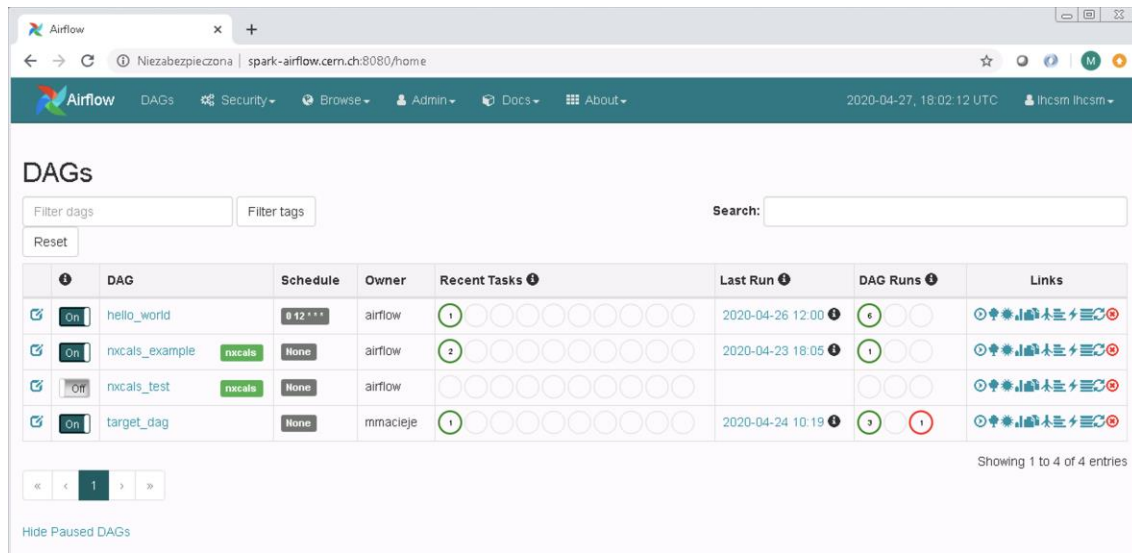
Source: <https://en.wikipedia.org/wiki/Cron>

A Solution – Apache Airflow

Airflow is designed under the principle of "configuration as code".

While other "configuration as code" use markup languages (XML), Airflow allows developers to import python libraries and classes.

In addition, Airflow provides an intuitive WebUI.



The screenshot displays the Apache Airflow WebUI interface. At the top, there's a navigation bar with the Airflow logo, 'DAGs', 'Security', 'Browse', 'Admin', 'Docs', and 'About' menus. The current date and time are shown as '2020-04-27, 18:02:12 UTC'. Below the navigation bar, the 'DAGs' section is visible, featuring a search bar and a table of DAGs.

	DAG	Schedule	Owner	Recent Tasks	Last Run	DAG Runs	Links
<input checked="" type="checkbox"/>	hello_world	@ 12 ***	airflow	1	2020-04-26 12:00	6	Refresh Refresh Refresh Refresh Refresh Refresh
<input checked="" type="checkbox"/>	nxcals_example	nxcals None	airflow	2	2020-04-23 18:05	1	Refresh Refresh Refresh Refresh Refresh Refresh
<input type="checkbox"/>	nxcals_test	nxcals None	airflow				Refresh Refresh Refresh Refresh Refresh Refresh
<input checked="" type="checkbox"/>	target_dag	None	mmacieje	1	2020-04-24 10:19	3	Refresh Refresh Refresh Refresh Refresh Refresh


Showing 1 to 4 of 4 entries




Source: https://en.wikipedia.org/wiki/Apache_Airflow



Apache Airflow is based on a concept of direct acyclic graphs to represent stages of application execution.



This is a DAG



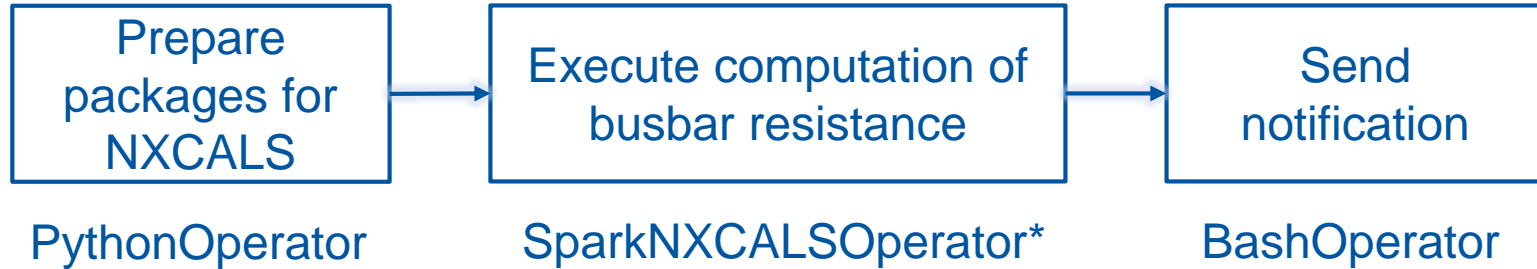
This is not a DAG

Each task is based on an operator, e.g. PythonOperator

For more details, please consult: <https://airflow.apache.org>

Sample DAG

Apache Airflow provides python scripting capabilities with an intuitive way for editing script parameters and an integration with the NXCALS cluster*.



With the SparkNXCALSOperator we can execute operations just like with SWAN.

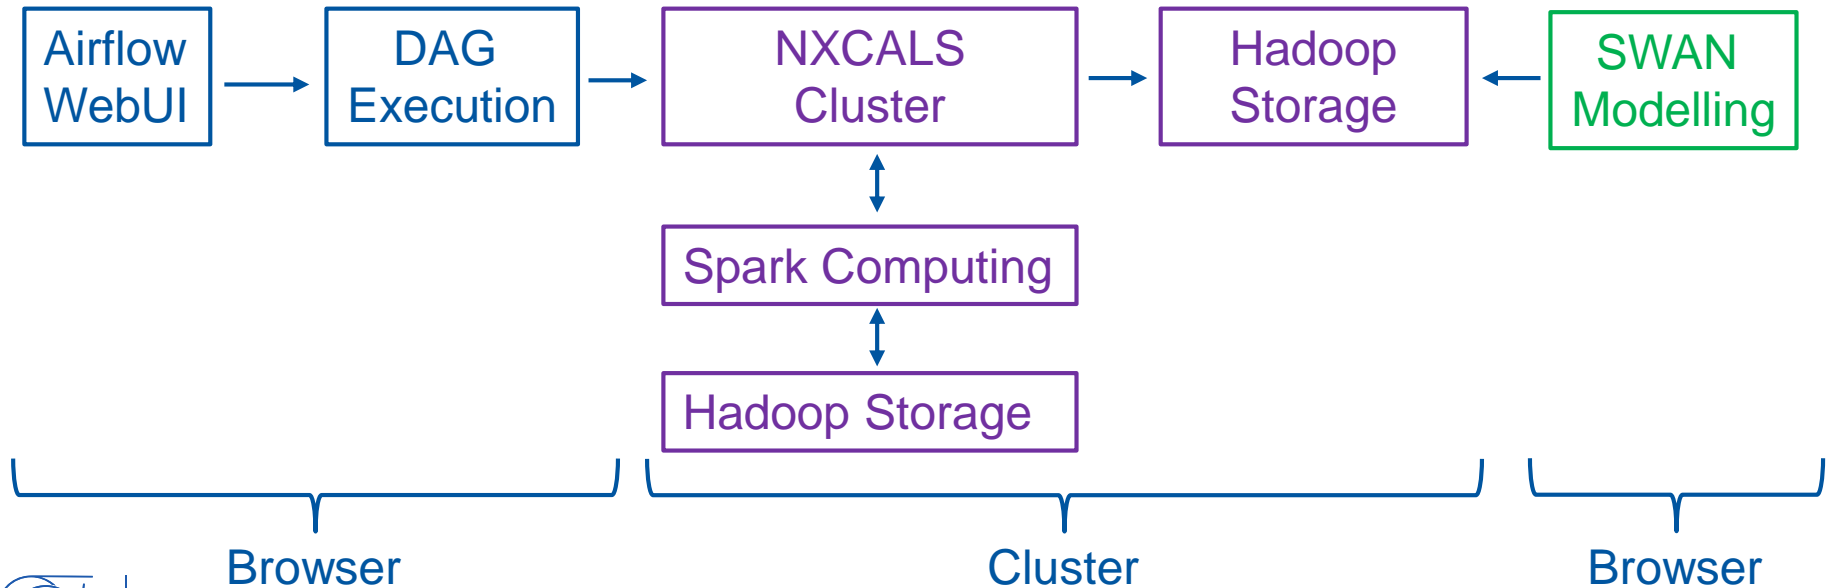
* https://gitlab.cern.ch/db/spark-service/airflow/-/blob/master/plugins/nxcals_plugin/operators/nxcals_operator.py



For more details please check a presentation from Piotr Mrówczyński:
<https://cernbox.cern.ch/index.php/s/8ypLPt6JTCIGIc7>

How does it work?

- ✓ A DAG is triggered from Airflow manually, on-demand, in regular intervals.
- ✓ An NXCALS computation is executed on the cluster (outside of Airflow environment).
- ✓ The computation stores results in Hadoop, which is accessible from SWAN



Demo

1. Hello World
2. NXCALS example
3. On-demand analysis
4. Computation of busbar resistance historical data
5. Computation of BLM historical data

- User-friendly interface is needed to define signals, type of monitoring, time intervals, warning/alarm levels, e-mail address person to be notified.

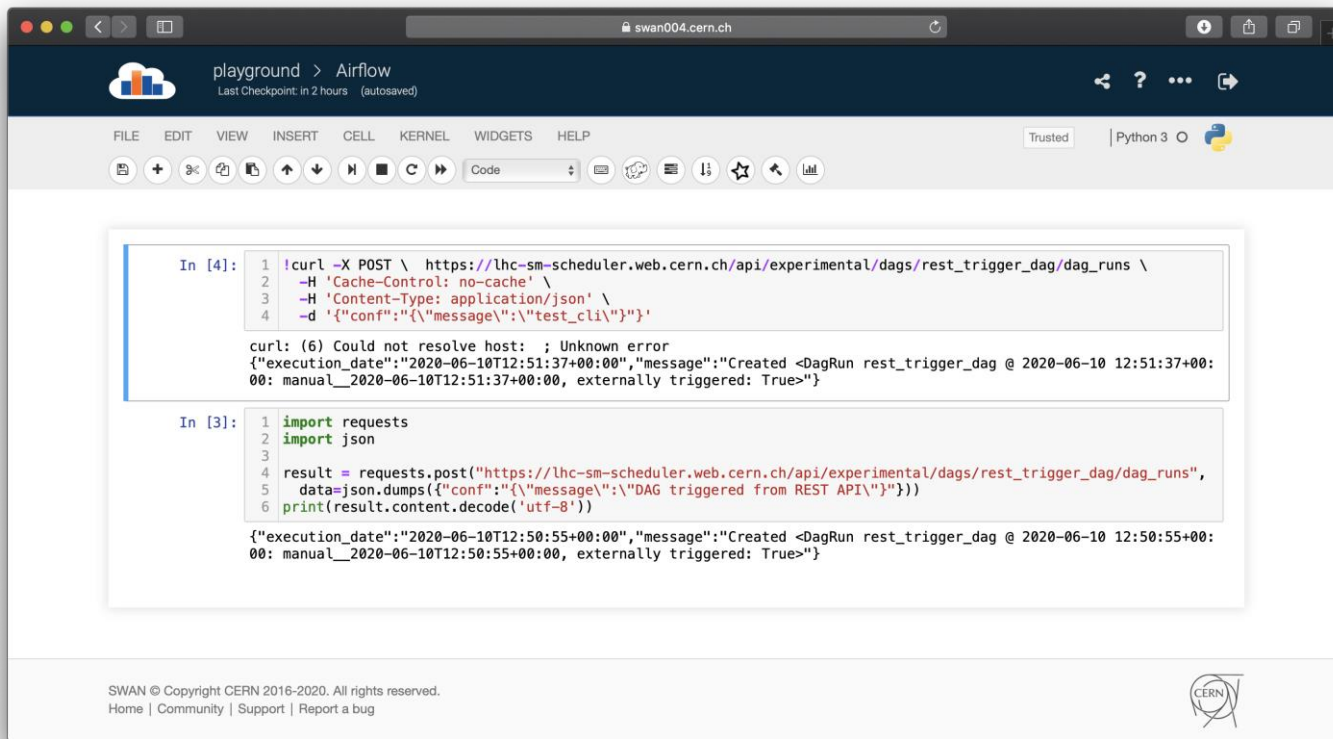
Demo

1. **Hello World**
2. **NXCALS example**
3. On-demand analysis
4. Computation of busbar resistance historical data
5. Computation of BLM historical data (work in progress)

- User-friendly interface is needed to define signals, type of monitoring, time intervals, warning/alarm levels, e-mail address person to be notified.

On-Demand Execution

Our analysis can be triggered directly by hardware, e.g. an FPA, a change in the beam mode, etc.



The screenshot shows a web browser window with the URL `swan004.cern.ch`. The page title is "playground > Airflow" and it indicates "Last Checkpoint: in 2 hours (autosaved)". The interface includes a menu bar (FILE, EDIT, VIEW, INSERT, CELL, KERNEL, WIDGETS, HELP) and a toolbar with various icons. The main content area displays two code blocks with their respective outputs.

```
In [4]: 1 !curl -X POST \ https://lhcs-sm-scheduler.web.cern.ch/api/experimental/dags/rest_trigger_dag/dag_runs \
2 -H 'Cache-Control: no-cache' \
3 -H 'Content-Type: application/json' \
4 -d '{"conf": {"message": "test_cli"}}'
```

```
curl: (6) Could not resolve host: ; Unknown error
{"execution_date": "2020-06-10T12:51:37+00:00", "message": "Created <DagRun rest_trigger_dag @ 2020-06-10 12:51:37+00:00: manual__2020-06-10T12:51:37+00:00, externally triggered: True>"}
```

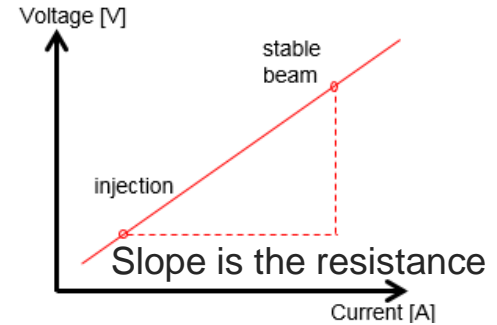
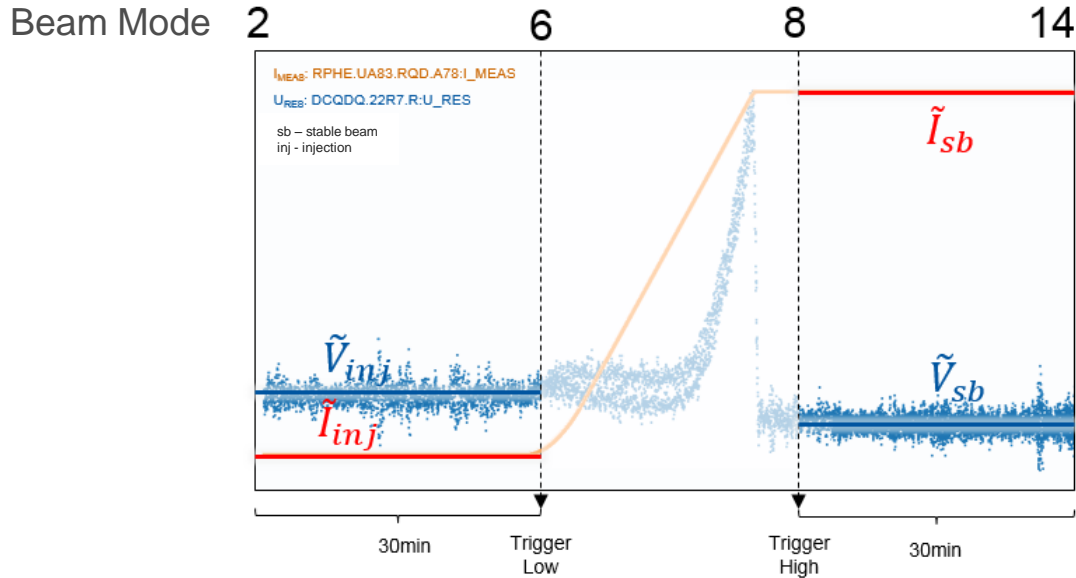
```
In [3]: 1 import requests
2 import json
3
4 result = requests.post("https://lhcs-sm-scheduler.web.cern.ch/api/experimental/dags/rest_trigger_dag/dag_runs",
5 data=json.dumps({"conf": {"message": "DAG triggered from REST API"}}))
6 print(result.content.decode('utf-8'))
```

```
{"execution_date": "2020-06-10T12:50:55+00:00", "message": "Created <DagRun rest_trigger_dag @ 2020-06-10 12:50:55+00:00: manual__2020-06-10T12:50:55+00:00, externally triggered: True>"}
```

At the bottom of the page, there is a footer with the text: "SWAN © Copyright CERN 2016-2020. All rights reserved. Home | Community | Support | Report a bug". The CERN logo is visible in the bottom left and right corners.

Busbar Resistance

Monitoring of busbar involves calculation of 1248 (RB) and 400 (RQ) resistances
The resistance is calculated from a linear fit of voltage and current at plateaus



Courtesy: C. Obermair



Z. Charifoulline, et al., "Resistance of Splices in the LHC Main Superconducting Magnet Circuits at 1.9 K,"

IEEE TAS, 28(3), Apr. 2018.

C. Obermair, SUMM Report: <http://cds.cern.ch/record/2639871?ln=en>

Busbar Resistance

We gather historical data for Run2

```
1 i_meas_feature_df = QueryBuilder() \  
2   .with_nxcals(spark) \  
3   .with_duration(t_start='2018-05-21 12:22:37', t_end='2018-05-21 13:49:12') \  
4   .with_circuit_type('RB') \  
5   .with_metadata(circuit_name='*', system='PC', signal='I_MEAS') \  
6   .feature_query(['mean', 'std'], function=translate_udf).sort_values(by='class').df  
7  
8 u_res_feature_df = QueryBuilder() \  
9   .with_nxcals(spark) \  
10  .with_duration(t_start='2018-05-21 12:22:37', t_end='2018-05-21 13:49:12') \  
11  .with_circuit_type('RB') \  
12  .with_metadata(circuit_name='*', system='BUSBAR', signal='U_RES', wildcard={'BUSBAR': '*'}) \  
13  .feature_query(['mean', 'std'], function=translate_udf).sort_busbar_location('RB', circuit_name='*').df  
14  
15 ResistanceBuilder().with_busbar_voltage(u_res_feature_df).with_busbar_current(i_meas_feature_df) \  
16  .calculate_mean_resistance('RB').convert_to_row(index=t_start_inj)|
```

▶ Apache Spark: 10 EXECUTORS 20 CORES Jobs: 36 COMPLETED						
DCBB.8L2.R_mean_inj	DCBB.8L2.R_std_inj	DCBB.8L2.R_mean_sb	DCBB.8L2.R_std_sb	DCBB.8L2.R_R_RES	DCBB.9L2.R_mean_inj	DCBB
1526898157236000000	0.000064	0.000043	0.000076	0.000022	1.164031e-09	-0.000022

NXCALS cluster computation of **1248** busbar resistances takes approximately as much time as query and local processing of **8** power converter currents

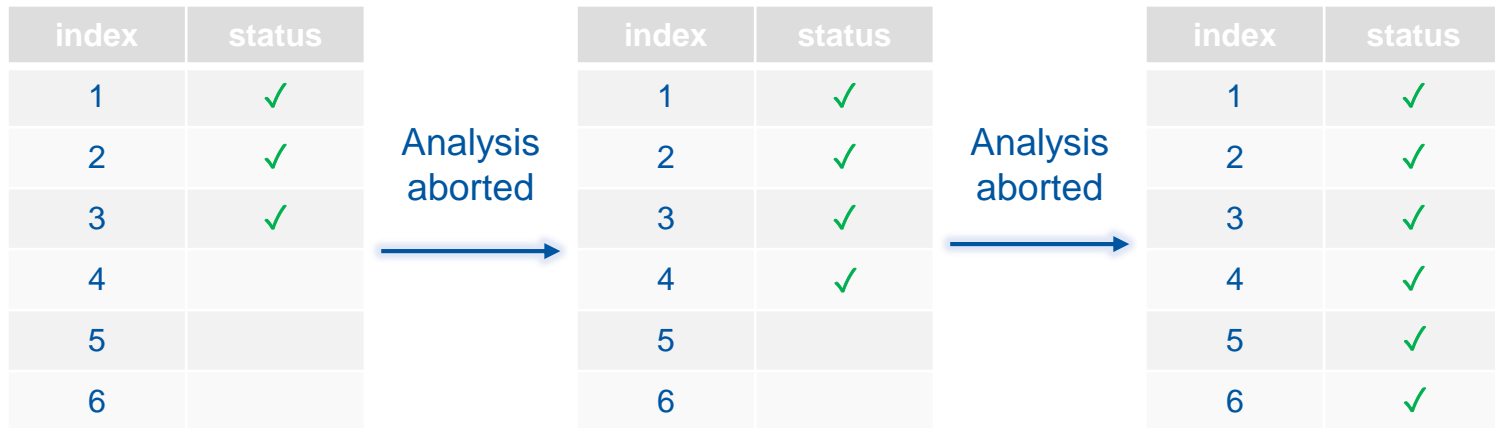


The selection of current plateaus is performed with a *user defined function* 22

Analysis with Checkpoints

Although we develop a solid software, the infrastructure can break leading to an abort of the computation.

In order to avoid calculating the same thing twice, we use checkpoints.



Saving Data with Partitions

While performing data collection over long periods of time, e.g. Run2 (2015-2018), it is advised to split the output dataset into partitions per year.

```
bash-4.2$ hdfs dfs -ls /project/lhc_signal_monitoring/busbar/rb
Found 4 items
drwxr-xr-x+ - lhcs m hdfs          0 2020-06-10 01:29 /project/lhc_signal_monitoring/busbar/rb/partition=2015
drwxr-xr-x+ - lhcs m hdfs          0 2020-06-08 14:24 /project/lhc_signal_monitoring/busbar/rb/partition=2016
drwxr-xr-x+ - lhcs m hdfs          0 2020-05-09 06:32 /project/lhc_signal_monitoring/busbar/rb/partition=2017
drwxr-xr-x+ - lhcs m hdfs          0 2020-05-06 03:30 /project/lhc_signal_monitoring/busbar/rb/partition=2018
```

This ensures that

- the table schema is consistent across a year
- each year can be loaded quickly
- all years (provided the schema is the same) can be read at once (by skipping partition)
- the data is stored at HDFS allowing for an immediate processing on Spark cluster



Exploration





Data Collection




Modelling





Monitoring


interactive notebook 


interactive compute


-  Jupyter Server
- python
- REST PySpark
- Spark Hadoop


output notebook


storage 


Tested notebook ready for scaling-up


parametrized notebook



parameterized script



scheduled compute

-  Apache Airflow
- python
- REST PySpark
- Spark Hadoop


storage 


Well-structured data ready for modelling

storage 



interactive notebook 


interactive compute

-  Jupyter Server
- python
- REST PySpark
- Spark Hadoop


storage 


Trained model ready for monitoring


parametrized notebook

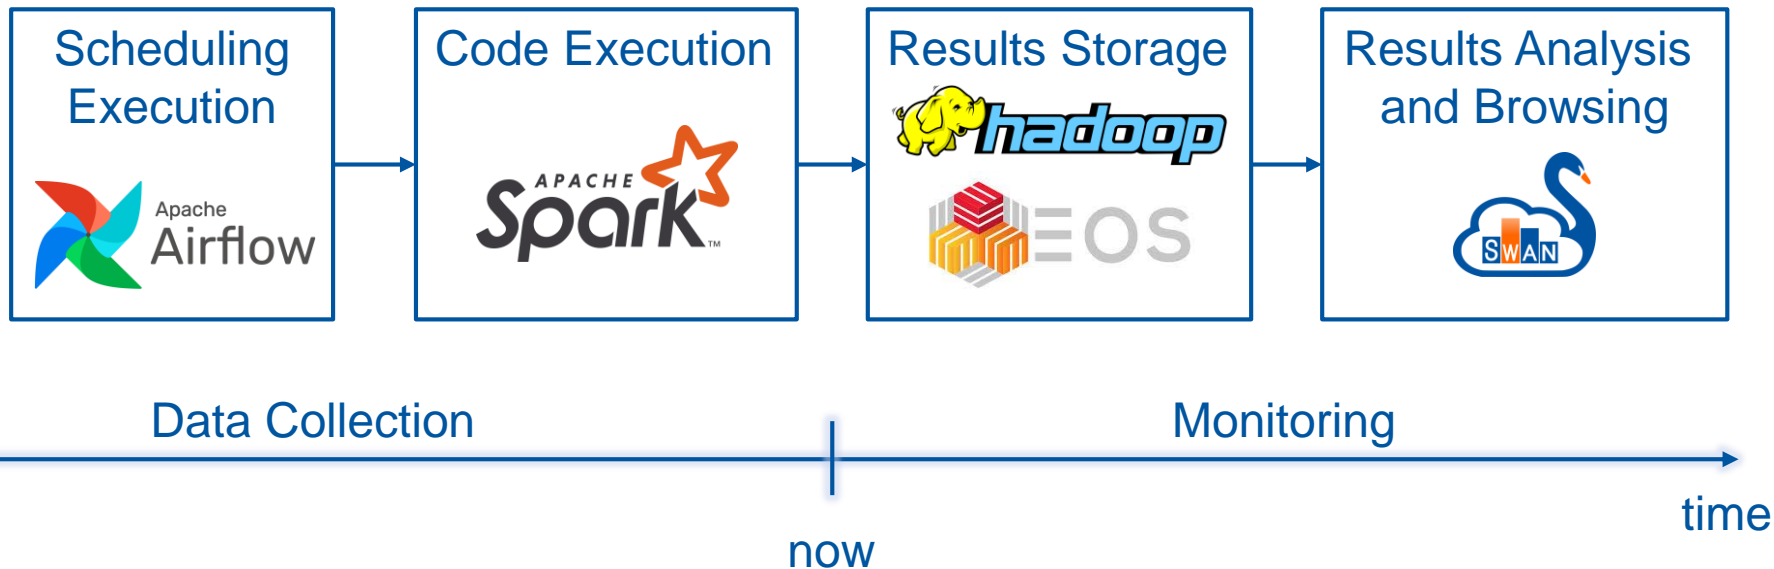

parameterized script

scheduled compute

-  Apache Airflow
- python
- REST PySpark
- Spark Hadoop

storage 

Signal Monitoring Pipeline



Single Statement Policy:

- at each stage only **one python** statement required to perform an operation
- this is **already** provided by Apache Airflow, Spark, and Hadoop
- we complement it with the **pyeDSL** (single statement to query, compute features, etc.)



The same infrastructure for data collection and monitoring scheduling

pyeDSL speaks AFT

Our pyeDSL provides a generic way of performing PM and NXCALS queries. Lately, the language has been extended to support AFT database:

- Context query

```
1 QueryBuilder().with_aft(session) \  
2   .context_query('faults/states')
```

	id	name
0	NON_BLOCKING_OP	Non-Blocking OP

- Fault query

```
1 QueryBuilder().with_aft(session) \  
2   .with_duration(t_start='2016-01-13T00:00:00Z', t_end='2016-05-13T00:00:00Z') \  
3   .fault_query(acceleratorId='LHC', accessNeeded=True)
```

	acceleratorName	acceleratorPropertyInstances	accessNeeded	description	displayLabel	duration	effectiveDuration	endTime
0	LHC	{{'propertyName': 'Time in Fill', 'value': '02...	False	access sectors 4,5,6,7 and LHCb indicate "blue...	None	4890000	4890000	2016-05-12T16:06:12Z

Quick development due to detailed documentation and a solid API: <https://aft.cern.ch/docs>
An example notebook with all types of queries:
<https://gitlab.cern.ch/LHCData/lhc-sm-api/-/blob/master/lhcsmap/dbsignal/aft/AFT.ipynb>

Summary

Apache Airflow provides a user-friendly WebUI for defining (in python) signals, type of monitoring, time intervals, warning/alarm levels, e-mail address of person to be notified. Also:

1. **long-running** data collection jobs (PM, NXCALS, **AFT**) with auto-restart in case of failures
2. **on-demand** execution of monitoring applications
3. **scheduled** execution of monitoring applications
4. *scheduled notebook execution and report generation (e.g. FPA analysis) - in progress*



Find out more at: <https://cern.ch/sigmon/about>

Summary

Apache Airflow provides a user-friendly WebUI for defining (in python) signals, type of monitoring, time intervals, warning/alarm levels, e-mail address of person to be notified. Also:

1. **long-running** data collection jobs (PM, NXCALS, **AFT**) with auto-restart in case of failures
2. **on-demand** execution of monitoring applications
3. **scheduled** execution of monitoring applications
4. *scheduled notebook execution and report generation (e.g. FPA analysis) - in progress*

The results of data collection are stored in **Hadoop** enabling system modelling in **SWAN**.

Our datasets will be openly available on our website: <https://cern.ch/sigmon/datasets>

Apache Airflow **is not yet** officially supported by IT. However, the NXCALSOperator and a Docker image are already backed up by our colleagues.



Find out more at: <https://cern.ch/sigmon/about>



Notebook-Centric Infrastructure

