



Cognitive plants through proactive self-learning hybrid digital twins

DT-SPIRE-06-2019 (870130)

Deliverable Report

Deliverable ID	D1.4	Version	V1.1
Deliverable name	A complete digital twin enabled with cognitive elements for the Non-Ferrous pilots		
Lead beneficiary	HYDRO, ELKEM		
Contributors	SINTEF, CYBERNETICA		
Due date	28-February-2023		
Date of final version	23-November-2023		
Dissemination level	Public Demonstrator		
Document approval	Frode Brakstad	23.11.2023	



The COGNITWIN project has received funding from the European Union's Horizon 2020 research and innovation programme under GA No. 870130

PROPRIETARY RIGHTS STATEMENT

This document contains information which is proprietary to the COGNITWIN consortium. The document or the content of it shall not be communicated by any means to any third party except with prior written approval of the COGNITWIN consortium.

Executive Summary

This document is a public deliverable “D1.4: A complete digital twin enabled with cognitive elements for the Non-Ferrous pilots” of the project COGNITWIN and describes the work carried out for two of the pilots. It is accompanied by digital twin pipeline descriptions in the COGNITWIN Toolbox and also with Demonstrator videos in the COGNITWIN YouTube channel.

In the **Hydro** pilot, the main aim is to evenly adsorb HF in the primary alumina feed. This is because the variation of fluoride in alumina fed to electrolysis cells influences the electrolysis heat balance, which in turn influences the yield of the cell. Even fluoride content in alumina will be achieved through logging/predicting over time the HF evolution (fluoride to be adsorbed) and adjusting the alumina feed such that the weight fraction of adsorbed fluoride becomes as uniform as possible. Several non-process factors will have to be included in the predictions, including: alumina quality, humidity and temperature. The temperature can be adjusted in the pilot and sought to be held as constant as possible ($\sim 90^{\circ}\text{C} \pm 5^{\circ}\text{C}$), to ensure even and optimal conditions for adsorption. The pilot also includes active regulation of the three main fans, which in total can consume up to 3,6 MW when on max load. By regulating these fans based on the present operating conditions and an even mass flow of off gas from the electrolysis process, one can save energy, here stipulated to 5% reduction compared with today. The Hydro pilot has been built on three cases, and the results from these can be seen below:

Case I; has resulted in a digital twin predicting HF evolution, hence able to optimise the feed of primary alumina to the GTC accordingly. The twin utilises data's from process system, meteorological servers and also from the flow control/measurement in Case III. The twin is based on a first-principles dynamic model. This model has been extended to be partially data-driven (hybrid), and lastly operators' experience has been incorporated into the optimisation system, i.e. a cognitive element.

Case II; the case of temperature regulation has, for all practical purposes met its objective, i.e. keeping the temperature in the reactor $90 \pm 5^{\circ}\text{C}$. The achieved range was $91.8 \pm 5,9^{\circ}\text{C}$ – a huge improvement over non-optimised performance and very much within the preferred range for HF adsorption in GTCs.

Case III; has demonstrated its value in saving energy simply by harvesting the changes in temperature. The savings during testing/measuring were found to be 7%. Moreover, the control of the main fans enables the suction rate to be lowered down towards the leakage limits of the cells, resulting in a further 15% potential reduction of energy use. The case targets have therefore been met.

Hydro consists of several production sites and will be able to bit-by-bit include elements from the pilot in all of our plants after COGNITWIN ends. Moreover, the methodologies developed/refined in COGNITWIN Hydro pilot will also be applied to other purposes; for example, meteorological data will be used for pot room ventilation measurements, which is a vital part of emission control.

In the **Elkem** pilot, the main goal is to develop an online model for refining of ferrosilicon alloys combining state-of-the-art sensors and a dynamic model for mass balance and heat transfer. By adjusting to real-time measurements, the refining and alloying of liquid ferrosilicon will be more precise, and this will eventually lead to higher yield, improved accuracy and lower specific energy consumption. The development and implementation of a cognitive twin will help to reduce the process variations caused by human interpretation as well as enhancing the decision support for the

operators. Furthermore, the installation of 3 infrared cameras at different locations in the process line will enable additional measurements of the process and give valuable information that can be automatically processed and fed into the online model for improved predictive capability.

D1.4 adds the following

This report is the result of the second stage (M30-M36 (M42)) of the project development and builds in particular on the previous confidential deliverables "D1.2: A Database and digital platform for the Non-ferrous pilots", and "D1.3: Hybrid models with cognitive elements for the Non-Ferrous pilots" In particular. Following the revision of D1.3, D1.4 is a public deliverable that adds the following:

For the Hydro pilot

- Improved data quality and new data measurements used by the pilot in M18-M30
- Online architecture of Hybrid Digital Twin now running on plant servers
- Current and planned Digital Twin functionalities
- Status and further deployment plans for automation
- Final public Demonstrator video for Hydro

For the Elkem pilot

- A description of the IoT platform and pipeline of data, model and application components
- Installation of IR camera for tapping and casting processes.
- Streaming of live video from cameras
- Status and plans for further deployment
- Final public Demonstrator video for Elkem

1 Table of Contents

- 1 Hydro pilot..... 11
 - 1.1 Introduction to Hydro & Process description..... 11
 - 1.1.1 Hydro & Process description 11
 - 1.1.2 Hybrid and Cognitive Digital Twins for the Hydro Pilot..... 12
 - 1.2 Overview case status and current challenges 12
 - 1.2.1 Case I – Match primary alumina feed to HF content in alumina..... 13
 - 1.2.2 Case II – Temperature control..... 15
 - 1.2.3 Case III – Main fan control..... 17
 - 1.3 Development of data connectivity and system architecture 22
 - 1.3.1 Digital twin migration to new permanent Hydro server 22
 - 1.3.2 Virtual PC for running MET Norway weather interface 22
 - 1.4 Online hybrid digital twin 23
 - 1.4.1 Data connectivity and system architecture..... 23
 - 1.4.2 First-principles model..... 24
 - 1.4.3 Hybrid digital twin - model adaptivity..... 25
 - 1.5 Cognitive digital twin..... 28
 - 1.5.1 Optimisation to match primary alumina feed to HF-content in alumina..... 28
 - 1.5.2 Cognitive silo level constraint for avoidance of segregation effects..... 29
 - 1.5.3 User interface / dashboard 30
 - 1.5.4 Autoencoder for anomaly detection..... 31
 - 1.6 Future work (post-COGNITWIN)..... 34
 - 1.6.1 Testing of automatic control 34
 - 1.6.2 Container solutions 34
 - 1.6.3 Integration of Autoencoder..... 34
 - 1.6.4 Follow-up on KPIs 35
- 2 Elkem Pilot..... 35
 - 2.1 Introduction and background..... 35
 - 2.1.1 Motivation 35
 - 2.1.2 Process description..... 36
 - 2.2 Use case and challenges 37
 - 2.3 Installations of IR cameras in production line 39
 - 2.3.1 IR-camera – refining 40

2.3.2	IR camera – tapping.....	40
2.3.3	IR camera – casting.....	40
2.4	Image analysis and algorithm development	40
2.4.1	Introduction and background.....	40
2.4.2	Refining camera.....	41
2.4.3	Tapping camera	43
2.4.4	Casting camera	44
2.4.5	Infrared image analysis	45
2.5	Online cognitive digital twin.....	47
2.5.1	Data connectivity and system architecture.....	47
2.5.2	Digital twin.....	48
2.5.3	Hybrid digital twin	51
2.5.4	Cognitive digital twin.....	51
2.5.5	Operator support system	51
2.6	Future activities after project ends	52
2.6.1	Camera installation.....	52
2.6.2	Hybrid/cognitive twin.....	53
2.6.3	Follow-up on KPIs	53
3	Conclusions.....	54
3.1	Conclusions Hydro Pilot.....	54
3.2	Conclusions Elkem Pilot.....	55
3.2.1	Cognitive twin vs hybrid twin	55
3.2.2	Summary of challenges	56
3.2.3	Summary and conclusive remarks.....	57
4	References.....	58

List of Figures

Figure 1. Flow Sketch of aluminium electrolysis process, with Gas Treatment Centre (GTC).	11
Figure 2. Hydro pilot Cases.....	12
Figure 3. Raw gas HF concentration from a production site (not pilot case).....	13
Figure 4. Covariation between LOI and HF evolvment.....	14
Figure 5. Absolute humidity in the ambient air covariating with HF concentration in raw gas from electrolysis. 14	
Figure 6. Adsorption capacity for SGA as function of humidity and temperature [1]	15
Figure 7. Measured filter reactor temperature for GTC facilities.	15
Figure 8. Normal distribution of reactor temperatures in GTC.....	16
Figure 9. Adsorption capacity for SGA as function of humidity and temperature [1], with indicated (red ellipse) operational area for Hydro Pilot	16
Figure 10. Left, a), 3D drawing of main raw gas duct, right b), CFD model geometry.....	18
Figure 11. Placing the dP and temperature points by utilising the CFD simulations.	18
Figure 12. Placement and installation of dP sensor and HF laser (red circle) in main raw gas duct in the GTC.	19
Figure 13. Result of CFD calibration of gas flow measurement system.....	19
Figure 14. Generated example of surplus energy usage of constant mass flow (Nm ³ /h) with set point at 15 °C ambient temperature.....	20
Figure 15. Main fan control, gas flow and corresponding set point.	20
Figure 16. Zoomed period of main fan regulation, gas flow versus set point.	21
Figure 17 Sum of energy usage from 4 main fans. Red, with manual load adjustment. Green, autoregulated with flat normal cubic meter flow (Nm ³ /h].....	21
Figure 18: Schematic showing the retrieval and parsing of raw data by the MAPI interface, the preparation and formatting of the data by the Python service and the transfer of the data to the digital twin via Cybernetica OPC UA server.	23
Figure 19: Data flow (including external sources) to the server running the digital twin.	24
Figure 20: IoT architecture for the Hydro pilot showing interconnectivity of data and toolbox components.....	24

Figure 21: Diagram of dynamic GTC model control volumes and process flows. Alumina flows are shown in black (low fluoride) and green (high fluoride). Gas flows are shown in red (high fluoride) and blue (low fluoride). Control volumes are marked with dashed lines. 25

Figure 22: Dynamic model results for several weeks in January 2023 with adaptive LOI-parameter estimation. 26

Figure 23: Model predictions vs. measurements of HF emissions from the GTC with adaptive parameter (saturation concentration of HF in alumina) estimation..... 27

Figure 24: Model vs. sensor predictions of mass in the secondary alumina silo with adaptive parameter (primary alumina input to the GTC) estimation. 27

Figure 25: Schematic of cognitive (cyclical) silo level constraint. The total period for a silo cycle is configured to last for 28 days..... 30

Figure 26: Preview (screenshot) of web-based dashboard for the Digital Twin operator interface. ... 31

Figure 27. Example of the layers' structure in an autoencoder. 32

Figure 28. Detection of anomalies in the HF laser measurements (Monitor). The shaded red areas are the ones flagged by the autoencoder as anomalies. The reconstructed monitor data is also shown, together with the model data from Cybernetica's model (here displayed only for comparison, the Model data is not used by the autoencoder). We see that the autoencoder flags areas where the Monitor data behaves unexpectedly, and significantly differs from the model data. 33

Figure 29. Detection of anomalies in the HF laser measurements (Monitor). The shaded red areas are the ones flagged by the autoencoder as anomalies. We notice that the autoencoder correctly flags the moment the HF-laser loses connection. 33

Figure 30 - Production of ferrosilicon. The focus of this project is the treatment and control of liquid FeSi after tapping of the submerged arc furnace. 36

Figure 31 - Tapping, refining/alloying and casting of ferrosilicon. The ladle may (as here) or may not be equipped with a bottom plug for addition of oxygen/air. 37

Figure 32. Depiction of how the slag thickness is estimated, with a simplified view of how the system appears to the camera to the right. 42

Figure 33. Surface slag divided into different temperature regions, with an uncalibrated volume estimate..... 42

Figure 34 - Infrared image from tapping furnace #2..... 43

Figure 35. Depiction of a slag-chunk within the liquid alloy flow field. Arrows indicate the image order. 44

Figure 36 - Casting process as viewed from the control room (left) and through IR camera (right). ... 45

Figure 37. Proposed IR camera architecture..... 46

Figure 38: An overview of the system architecture for the Elkem pilot. All components are currently installed at the plant and running online except the machine learning slag model.....	47
Figure 39: Digital twin pipeline for the Elkem pilot.....	48
Figure 40: A histogram of the error distribution of the modelled temperature.....	49
Figure 41: A histogram of the error distribution of the modelled aluminium concentration.	50
Figure 42: A histogram of the error distribution of the modelled calcium concentration.	50
Figure 43: A mockup of the Operator Support System.	52

List of Tables

N/A.

List of Acronyms

GTC	Gas Treatment Centre
HF	Hydrogen Fluoride
MW	Megawatt
NmP ³	Normal cubic meter, at 1013 mbar and 273,15°K (DIN 1343)
GE	General Electric
IHEX	Internal Heat Exchanger
SGA	Smelter-Grade Alumina
IoT	Internet of Things
KPI	Key Performance Indicator
KTP	Karmøy Technology Pilot
IPR	Intellectual Property Rights
ML	Machine Learning
AI	Artificial Intelligence
OPC-UA	Open Platform Communication- Unified Architecture
TDL	Trusted Data Layers
API	Application Programming Interface
APICS	Aluminium Production Control System
CFD	Computational Fluid Dynamics
LSTM	Long short-term memory
VM	Virtual machine
IR	Infrared
ROI	Region of interest
GUI	Graphical user interface
DAQ	Data acquisition system
MLP	Multilayer perceptron
LSTM	Long-short term memory
CNN	Convolutional neural networks
k-NN	k-nearest neighbour
PLS	Partial least square
NMPC	Nonlinear model predictive control
PTH	Post taphole

List of Authors and Reviewers

Are Dyrøy (Hydro) – Lead Author
 Kjetil Hildal (Elkem) – Lead Author
 Kasper Linnestad (Cybernetica) – Author
 Sudi Jawahery (Cybernetica) – Author
 Stein O. Wasbø (Cybernetica) – Author
 Filippo Remonato (SINTEF) – Author
 Stein Tore Johansen (SINTEF) - Author

Anders Hansen (SINTEF) – Author

Sølve Eidnes (SINTEF) – Author

Karl-Fredrik Hansen (Elkem) - Author

1 Hydro pilot

Some of the information in this chapter was introduced in the confidential deliverables D1.1, D1.2. and D1.3. It is also included here for the readability and context of this D1.4 public deliverable.

1.1 Introduction to Hydro & Process description

1.1.1 Hydro & Process description

The raw materials for electrolysis of Aluminium are Aluminium oxide (Al_2O_3), anode carbon (C), Aluminium fluoride (AlF_3), and electrical current. The electrolysis process with its Gas Treatment Centre (GTC) consists of two major material flows, a gaseous flow from the cell, and a bulk solids flow to the cell. In Figure 1 a sketch of the process flows is given. The alumina is discharged from a buffer or main storage (A) and fed to the dry scrubber reactor in the GTC (B). Here it adsorbs the HF and some SO_2 and the alumina is separated from the gas in a filter and collected in the filter house hopper (C). From this hopper the reacted alumina (secondary alumina) is sent to the electrolysis cell (D), either directly or via buffer storages before and after alumina distribution to individual electrolysis cells. The gas is collected from the electrolysis cell by means of suction (1), then drawn through the reactor (2) where it is mainly stripped for the HF gas (~95%) and moved further to the filter bag house where the filter cake adsorbs the rest of the HF gas. In total the efficiency of HF removal by the dry stage is 99.9%. After the dry stage the gas is pulled further by fans placed after the dry scrubber (3). The gas is after the fans pushed either directly to a stack (4), or into a wet scrubber stage for removal of SO_2 and rest HF. The wet scrubber stage uses either sea water or $\text{NaOH}+\text{H}_2\text{O}$ solution. The liquid absorbing the SO_2 and rest HF is returned to sea (5), often via make-up and conditioning and settling basins.

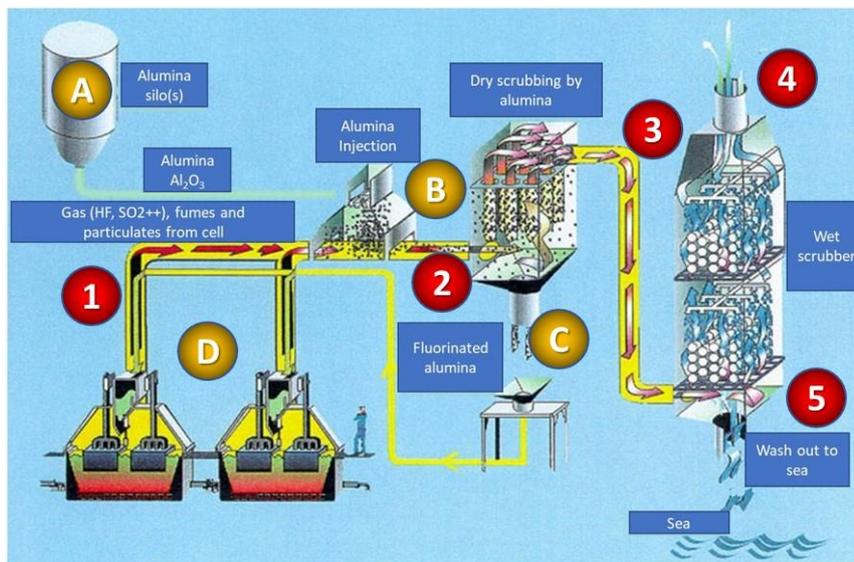


Figure 1. Flow Sketch of aluminium electrolysis process, with Gas Treatment Centre (GTC).

GTCs will vary in size. The capacity depends on the number of electrolysis cells it is set to serve. This can vary from a few cells (6-20), and all the way up to several hundreds, meaning that the capacity might go from gas capacities of ~200 000 Nm³/h to ~2 000 000 Nm³/h at 80 °C.

The GTC facilitating the Hydro pilot for the COGNITWIN project consists of 9 chambers of GE (now REEL) type ABART 600 with internal heat exchangers (IHEX). This gives a capacity of approximately 710 000 Nm³/h at 90 °C, and an alumina feed capacity above 20 t/h.

1.1.2 Hybrid and Cognitive Digital Twins for the Hydro Pilot

The COGNITWIN key phrases “Digital Twin”, “Hybrid Twin” and “Cognitive Twin” take on specific meanings in the scope of the Hydro pilot.

- “**Digital Twin**” is a basic analytical model representing the process – in the Hydro pilot context, this refers to the first-principles model first presented in D1.2 and updated in section 1.4.2
- “**Hybrid Twin**” is a data-adjusted analytical model to fit the actual process response and sensitivity – in the Hydro pilot context, this refers to the specific tuning and adjustments made to represent KTP data *as well as* the online model adaptivity (via data-driven state and parameter estimation) discussed in section 1.4.3.
- “**Cognitive Twin**” is an extension of a Hybrid Twin capable of recognising manual corrections and/or pre-learned feed forward information resulting in early onset of corrective actions – in the Hydro pilot context, this refers to the nonlinear model predictive control capabilities discussed in 1.5.1 and 1.5.2, which result in corrective actions in order to avert anticipated disturbances to the target variable (HF-content in secondary alumina).

1.2 Overview case status and current challenges

The Hydro pilot consists of 3 cases (as indicated with roman numerals in Figure 2).

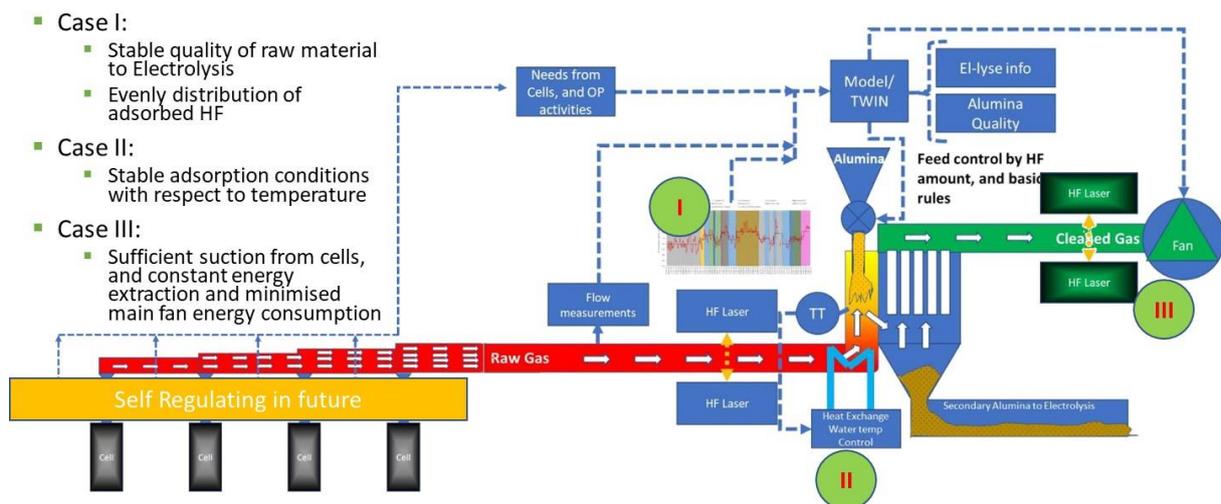


Figure 2. Hydro pilot Cases.

Each case has its own target. Each of the targets aim to optimise stability in electrolysis, and overall through this reduce energy consumption.

- Case 1. Matched and even distribution of HF to primary alumina feed, by demonstrating primary feed matching HF mass flow (calculated from logged operational data).
- Case 2. Keep constant temperature for best possible adsorption, i.e. 90°C ±5°C (from logged data)
- Case 3. Reduce the power consumption on the 3x 1 200 kW fans by 5%, measured by logged energy consumption from fans before and after activation of Case 3.

The common factor in all of our cases is the meteorological data, humidity and temperature.

1.2.1 Case I – Match primary alumina feed to HF content in alumina

Aluminas used for fluoride recovery in electrolysis, so called Smelter Grade Alumina (SGA), will vary in quality. The parameter called LOI (Loss of Ignition) is defined at 300-1 000 °C, and represents the weight fraction that escapes (evaporates or sublimates from) a heated sample. This means effectively the crystalline water in the alumina that is released when entering the electrolysis cell. Water, in the alumina and free moisture from the ambient air, is the main cause of HF evolution from the cell. The rate of HF will, in turn, affect the saturation of fluoride in the alumina from the GTC. By predicting and measuring the fluoride content in the raw gas and controlling the primary alumina fed to the GTC, we aim to distribute the fluoride evenly in the alumina returned to electrolysis.

In Figure 2 we see an example of this HF evolution and how the concentration in the raw gas changes through time/ seasons and alumina type used.

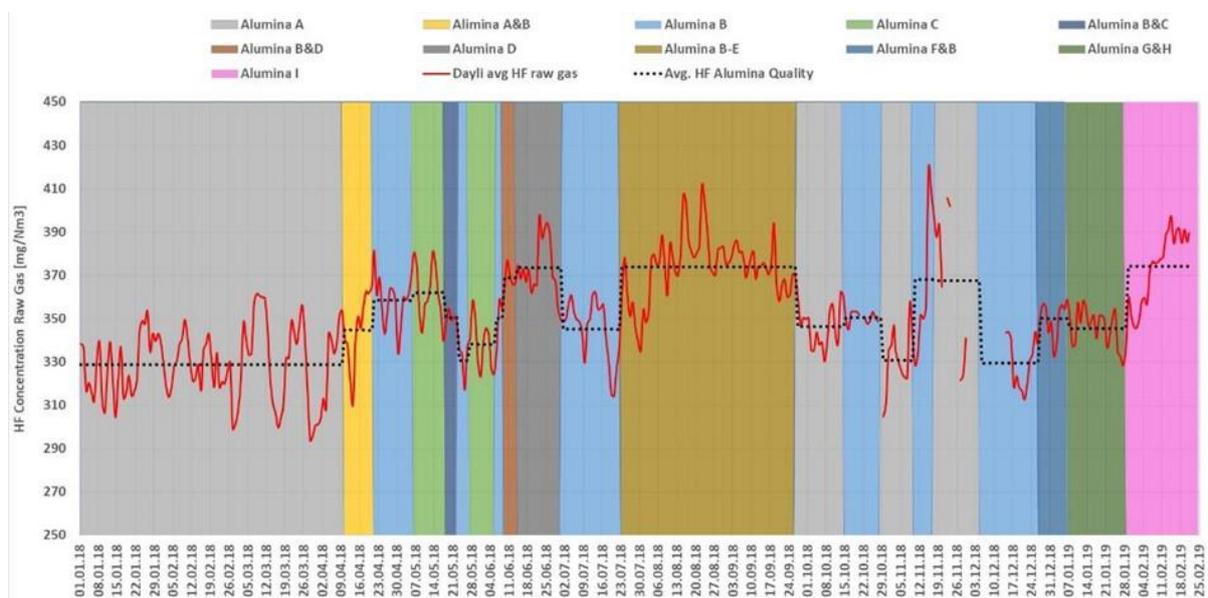


Figure 3. Raw gas HF concentration from a production site (not pilot case)

As mentioned, the variables found to influence the evolution of HF gas from the pot room are:

- LOI Crystalline water in the alumina, released between 300-1 100 °C
- MOI Moisture adsorbed during handling of alumina in transport and GTC operation, mostly from ambient air or compressor/blower air

- The amount of secondary alumina feed to the electrolysis cells

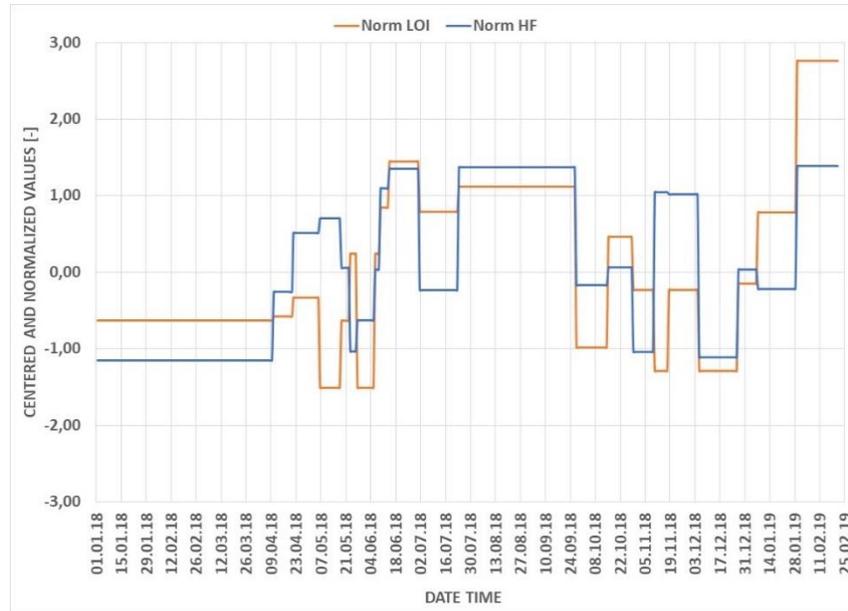


Figure 4. Covariation between LOI and HF evolvement

Although not directly correlated, the covariation between the absolute humidity in the air and the evolved HF in the raw gas can be seen in Figure 5 below. This indicates a linking influence from MOI on the HF evolvement.

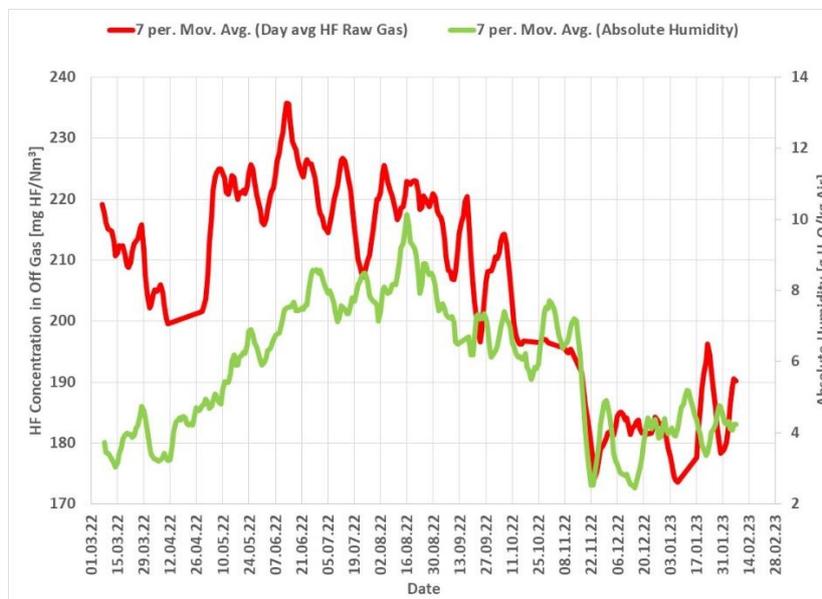


Figure 5. Absolute humidity in the ambient air covariating with HF concentration in raw gas from electrolysis.

Based on variables such as ambient humidity, LOI from alumina certificate, gas flow and alumina feed combined with literature on the topic, the digital twin will be able to predict the amount of HF to be adsorbed into the secondary alumina returned to electrolysis (see 1.4)

1.2.2 Case II – Temperature control

When the raw gas enters the reactor of the GTC, it meets a mix of primary and re-circulated alumina. The gaseous fluoride is then adsorbed to the alumina, this adsorption process is very temperature and humidity dependant, see Figure 6. Hence, control of the conditions is essential. By using the described IHX's the temperature is controlled, and later the internal re-circulation of alumina will be used to compensate for varying humidity.

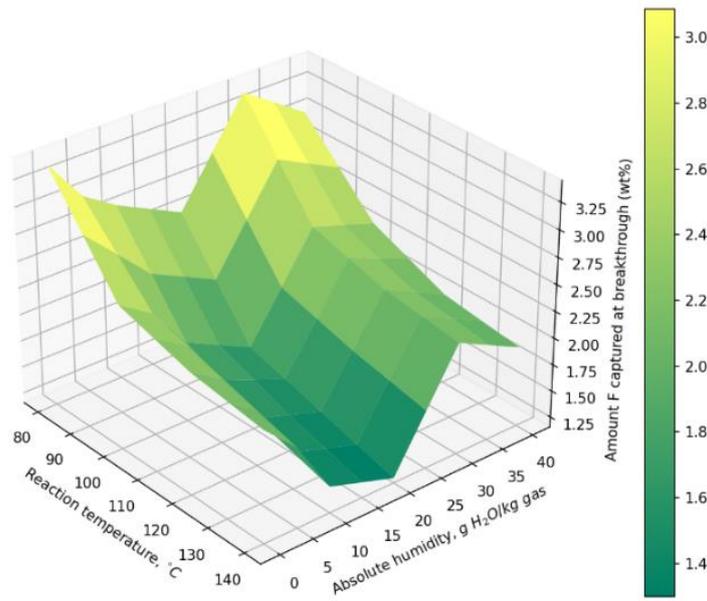


Figure 6. Adsorption capacity for SGA as function of humidity and temperature [1]

To measure the success in optimising the adsorption the reactor temperatures have been logged and are shown in Figure 7.

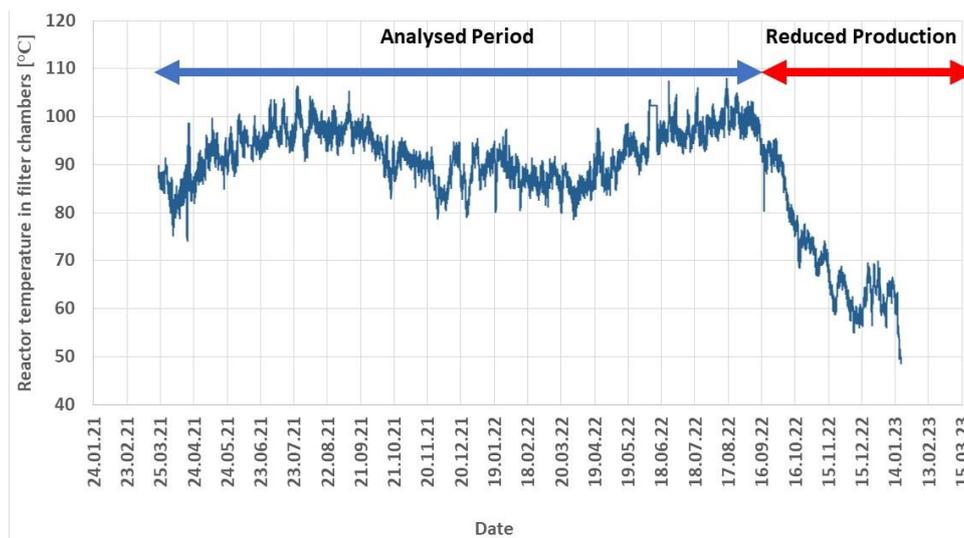


Figure 7. Measured filter reactor temperature for GTC facilities.

As can be seen in Figure 7, the temperature drops towards the end of the shown period. This is due to production curtailing due to the energy situation in Norway/Europe. This reduction in production leads to less energy to generate heat in the raw gas.

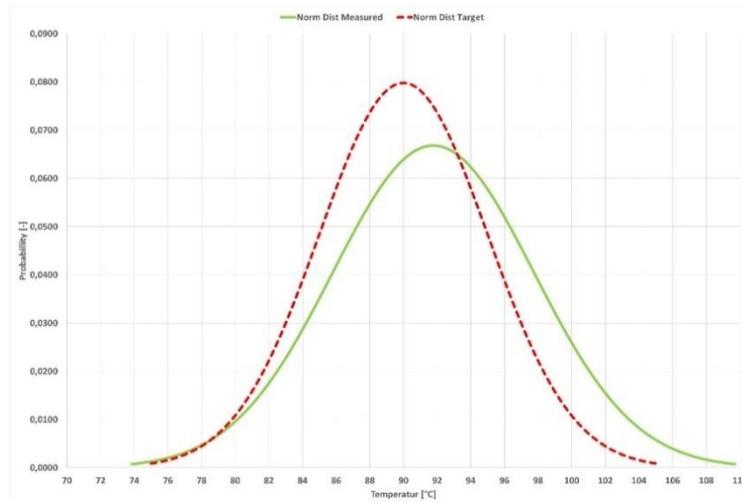


Figure 8. Normal distribution of reactor temperatures in GTC.

Figure 8 shows the normal distribution for the measured average temperature in the filter chamber reactors. The target was set to be $90^{\circ}\text{C} \pm 5^{\circ}\text{C}$. The range achieved was $91.8^{\circ}\text{C} \pm 5.97^{\circ}\text{C}$. Keeping in mind the harshness and volatility of the electrolysis process, this is great step forward. Moreover, from the measured data one can calculate that 96% of the sample population of temperatures measured in the analysed period is within the range of 80-110°C. This range is significant to the process because: 1) at prolonged temperatures below 70°C one runs the risk of condensation of acidic gasses, leading to corrosion of gas ducts and all other steel material, and 2) above 110°C, the adsorption capability of the primary alumina becomes much less effective.

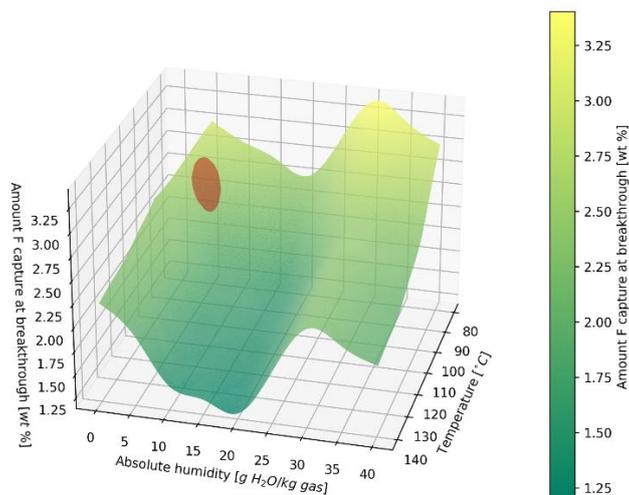


Figure 9. Adsorption capacity for SGA as function of humidity and temperature [1], with indicated (red ellipse) operational area for Hydro Pilot

As can be seen from Figure 9, the measured temperature range in combination with standard ambient conditions (absolute humidity) corresponds to an expected high adsorption potential, which is the desired target.

1.2.3 Case III – Main fan control

The main fan control is dependent on actual online gas flow measurements. This is further developed into the COGNITWIN system seen in 1.2.3.1 below.

1.2.3.1 Gas flow sensor system

The basic system of simplified gas flow measurement is done by utilising the basic equation given below:

$$\Delta p = \frac{1}{2} \cdot C \cdot \rho \cdot v^2$$

Where:

- Δp pressure difference measured [Pa]
- ρ The density of gas/ air at a given temperature [kg/m³]
- v The velocity of gas/ air [m/s]
- C A constant, typically representing a resistance in the system

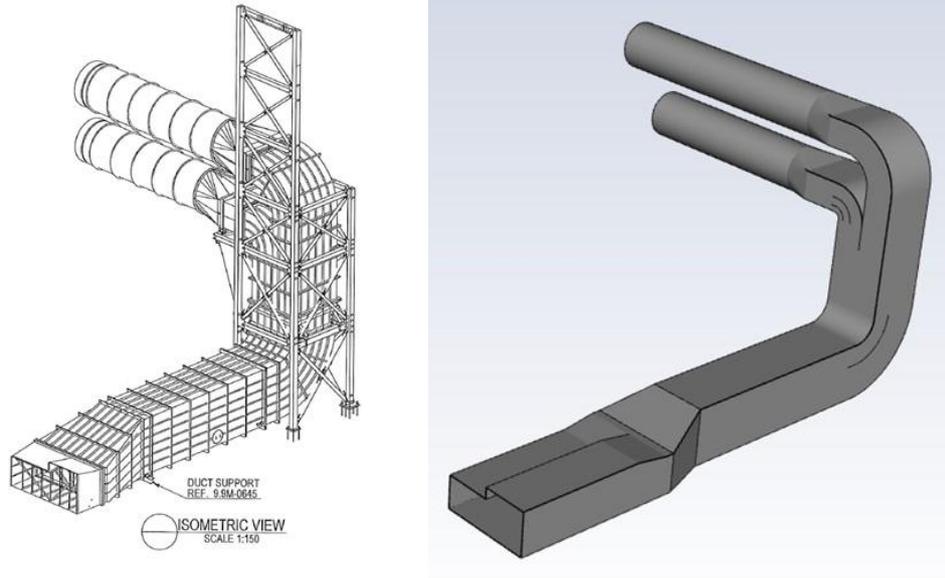
From this and the geometry of the duct, the gas flow will be found by equation below:

$$q \left[\frac{m^3}{h} \right] = v \cdot A \cdot 3600$$

Where:

- q Gas/ air flow [m³/h]
- v The velocity if gas/ air [m/s]
- A Cross sectional area of gas duct [m²]

In the delivered GTC system there is already a large number of sensors and estimators, including an estimator for the total gas flow. Gas flows are in general trivial to measure and monitor. However, in the sized ducts, with short distance to and from the last bend, it becomes more cumbersome. As an addition to the built-in estimator for flow, an added sensor system comprised of dP-cell and temperature has been installed and combined with Computational Fluid Dynamics (CFD) models to check the validity of the estimator. Figure 10 shows the drawing a), and the geometry b) used for CFD simulations.



a)

b)

Figure 10. Left, a), 3D drawing of main raw gas duct, right b), CFD model geometry.

Below, in Figure 11, the predicted results for the pressure and flow in the main raw gas duct are used for placing the connection points of the dP cells, and the temperature sensors. The model results are then used to establish the relations between measured pressure differences, temperatures and gas mass flow rates. The same methodology will be applied for the three intermediate gas ducts after the dry scrubber.

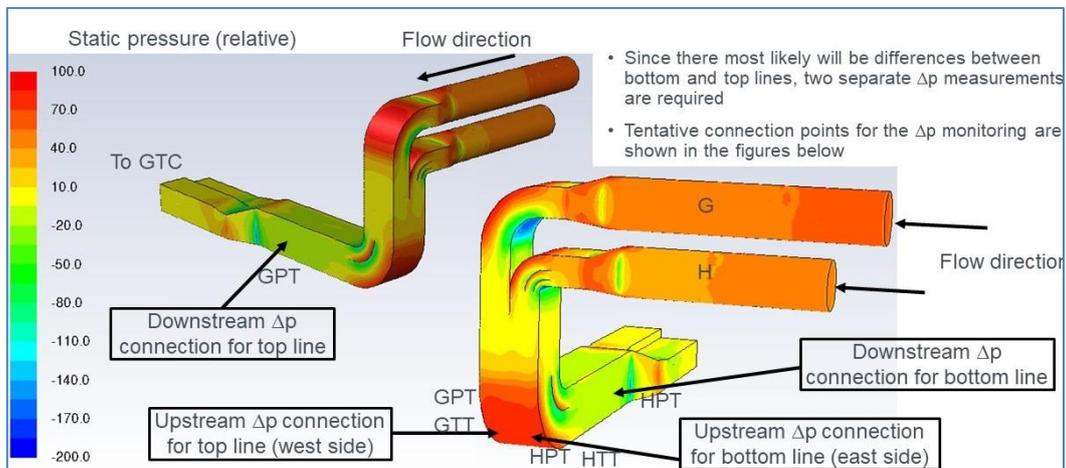


Figure 11. Placing the dP and temperature points by utilising the CFD simulations.

In Figure 12 below the full GTC is shown and the placement for the new dP based gas flow installation is pointed out in the main intermediate gas ducts. This will grant the opportunity to control the gas flow to each of the SO₂ scrubbers at the same time as one have the total flow available.

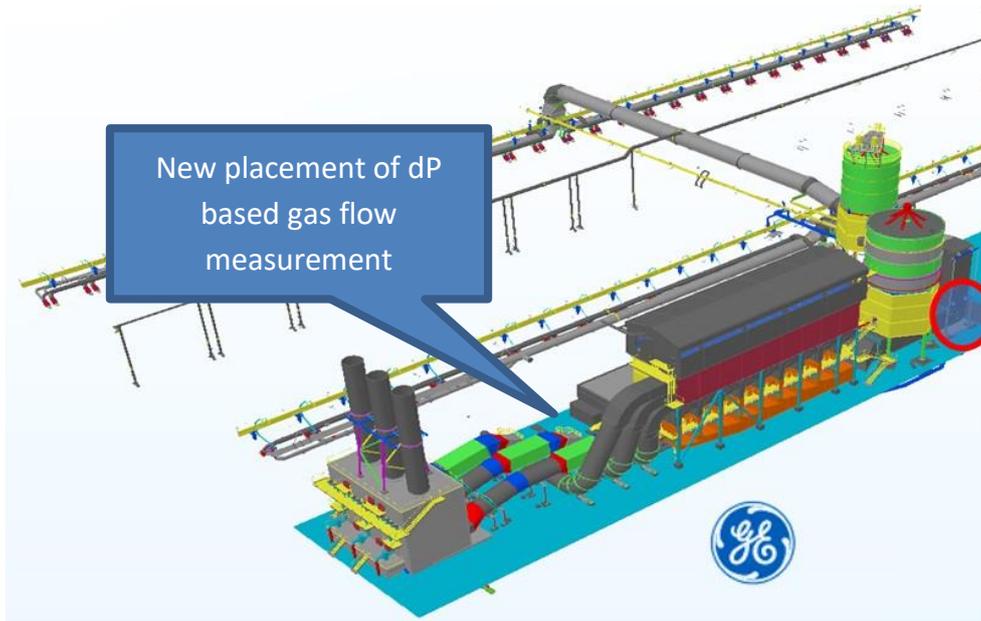


Figure 12. Placement and installation of dP sensor and HF laser (red circle) in main raw gas duct in the GTC.

For the new placement of the gas flow measurement, a CFD simulation is carried out to determine the respective resistance constants C_i as shown in Figure 13 below.

Pressure drop coefficients for each duct

- Velocity at the outlet (longest duct) is shown vs ΔP
- The three outlets have very similar velocities e.g. for $\dot{m} = 800\,000 \left[\frac{Nm^3}{hr} \right]$
 - outlet 1 (shortest duct) $V = 31.63 \left[\frac{m}{s} \right], P = 334 [Pa]$
 - outlet 2 (intermediate duct) $V = 33.04 \left[\frac{m}{s} \right], P = 399.1 [Pa]$
 - outlet 3 (longest duct) $V = 32.91 \left[\frac{m}{s} \right], P = 398.3 [Pa]$
- Outlet area: $3.097 [m^2]$
- Pressure drop Coefficients for each:
 - outlet 1 (shortest duct) $C_1 = 0.706$
 - outlet 2 (intermediate duct) $C_2 = 0.770$
 - outlet 3 (longest duct) $C_3 = 0.768$

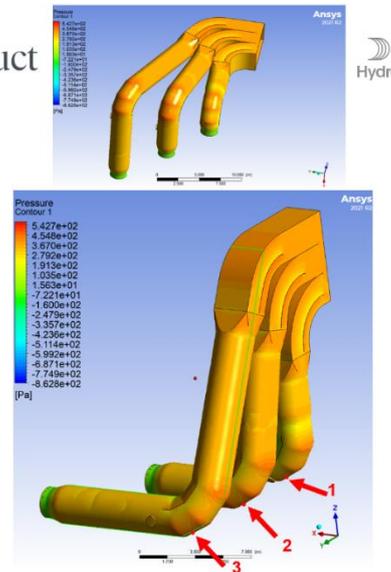


Figure 13. Result of CFD calibration of gas flow measurement system

Each of the fans at the pilot plant are controlled by frequency converters, making it possible to adjust to operate around a given set point for Normal Cubic Meter flow, which is a defined mass flow rate of gas.

1.2.3.2 Results of main fan control

When exhausting the electrolysis cells, a constant extraction of air mass flow is desired, since it is the air mass flow that conveys energy. Since fans operate with volumetric flow and not mass flow (measured in normal cubic meters), a fixed load on the fan will give variation in energy flow, due to air density varying with local weather conditions.

The accumulated energy savings is the integral of the blue curve in Figure 14. Thus there is fan energy to be saved if one could follow suit with the temperature.

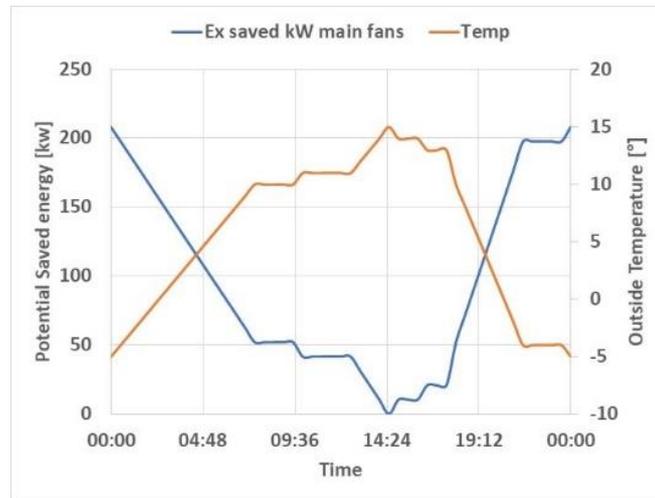


Figure 14. Generated example of surplus energy usage of constant mass flow (Nm³/h) with set point at 15 °C ambient temperature.

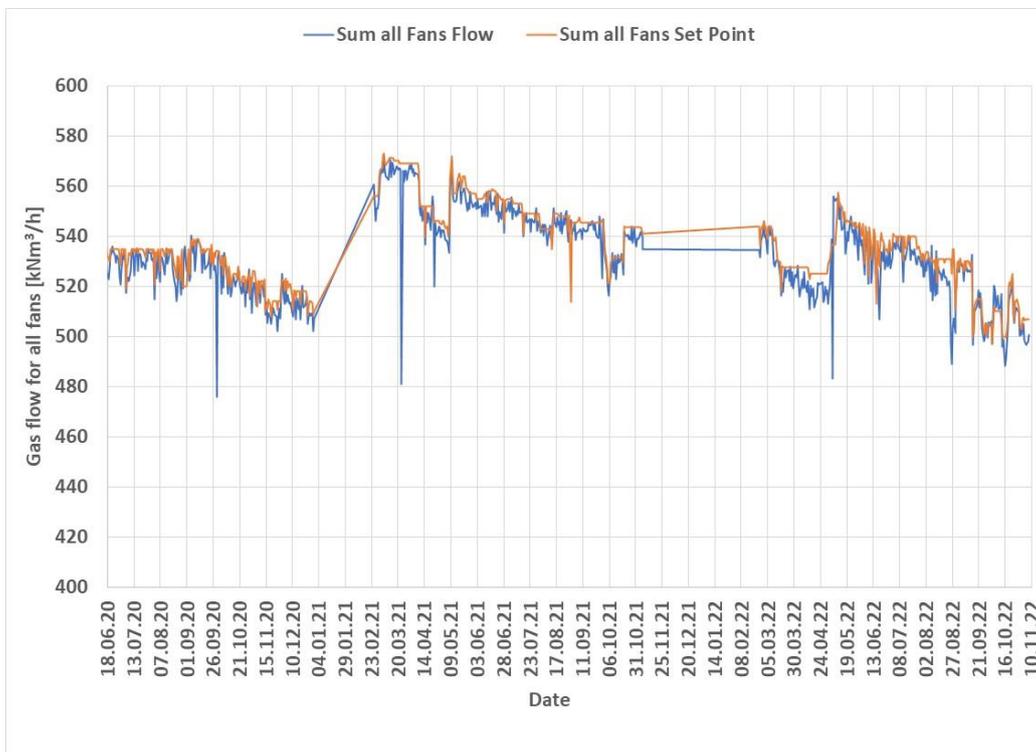


Figure 15. Main fan control, gas flow and corresponding set point.

For the main fans control, one can see long term data in Figure 15. If zooming in to check the measured gas flow follow the set point (see Figure 16), one can see that the gas flow follows the set point fairly

accurately. The rapid change in set point is due to ongoing maintenance and up-grade work in the GTC in general.

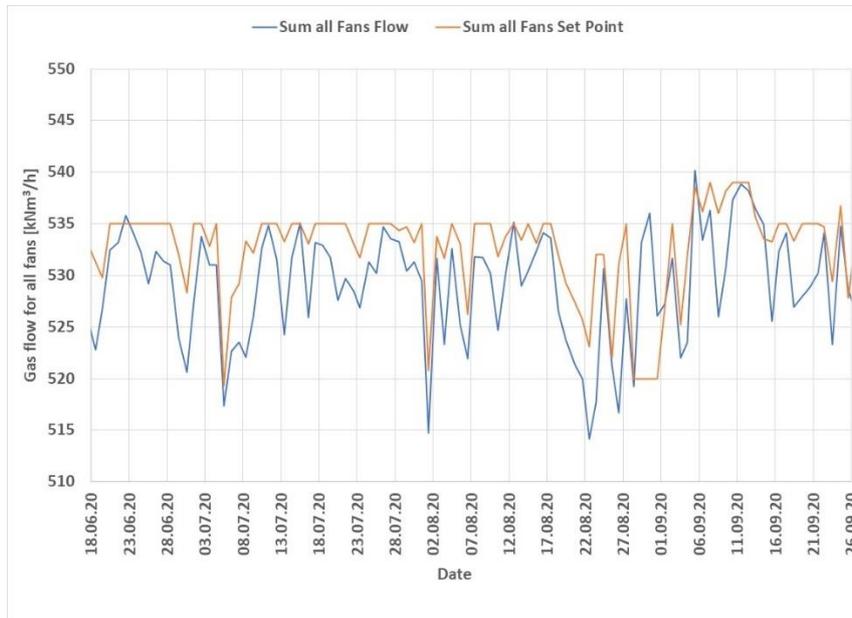
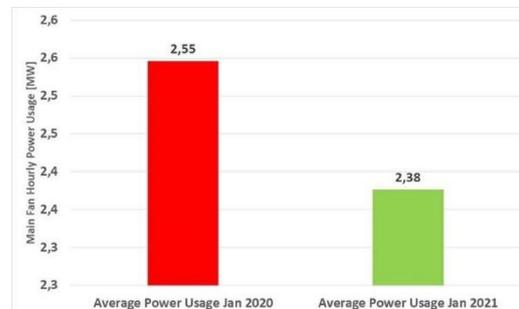
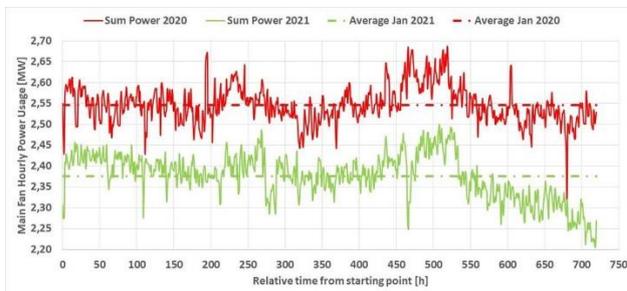


Figure 16. Zoomed period of main fan regulation, gas flow versus set point.

As reported in D1.2 (Figure 17), by fixing the mass flow through set fixed normal cubic meter (Nm^3/h), the saved power consumption are reduced by ~7%.



a) Comparing fan power with and without fixed mass flow of gas (fixed Nm^3/h)

b) Comparing fan power with and without fixed mass flow of gas (fixed Nm^3/h), average over 790 h.

Figure 17 Sum of energy usage from 4 main fans. Red, with manual load adjustment. Green, autoregulated with flat normal cubic meter flow (Nm^3/h).

From the above, the calculation of what this would mean for Hydro as a whole can be done. Main fans in Hydro Norway adds up to a nameplate effect of 52.4 MW, these are typically set to operate on ~70%, giving a constant effect of 37.4MW. By saving 7% of this the saving becomes 2.62 MW. For the full year it ads up to ~23 GWh.

The control of main fans also enables the operation of the electrolysis cells closer to limits for gas leakages. This due to knowing and controlling the gas flow opens for active adjustments. As a result, the gas suction from the cells can be reduced by at least 15%. Since the effect used by the fans are directly proportional (see equation below) with the gas flow (in m^3/h), this reduction in energy consumption can be added to the savings.

$$E = \frac{\Delta p_{total} \cdot q_{total}}{\eta}$$

Where:

- E is the power used [W]
- Δp is the total pressure drop met by the fan [Pa]
- q is the total gas flow [m³/s]
- η is the efficiency factor for the fan [-]

When considering the energy situation of late, this is valuable energy that can be released to other business or private consumption, and that is when only considering the main fans in Hydro Aluminium Norway. The results from COGNITWIN can apply to many large-scale fans in Europe.

1.3 Development of data connectivity and system architecture

1.3.1 Digital twin migration to new permanent Hydro server

Previously, the GTC digital twin was installed and run on a backup server for other soft-sensor applications where it lacked continuous access to internet-based weather data. This installation was always intended as a temporary solution while Hydro's KTP infrastructure was upgraded and designated GTC resources were made available. These milestones were reached during the summer of 2022 – in June, the digital twin and associated infrastructure were migrated to a permanent server that is both within Hydro's closed process network and is up-to-date with current Hydro KTP technology standards. Updated data connectivity to the digital twin is described in 1.4.1; the solution for accessing weather-based data from within the process network is described in 1.3.2.

1.3.2 Virtual PC for running MET Norway weather interface

Previous deliverables have described the development of the MET-API Interface (MAPI) in order to retrieve historical weather measurements and weather forecast data. At the end of 2021, MAPI was integrated into a Python service that continuously retrieves weather data and prepares a file containing the most recent weather measurements and the up-to-date weather forecast for use by the digital twin.

Since MAPI requires internet access, collecting weather data and then making it available for the digital twin running on Hydro's closed process network has been an ongoing and looming challenge. The original, temporary solution during the initial phase of online digital twin testing was to run the MAPI Python service on a network PC and then transfer a file including the weather data to the server running the digital twin. In June 2022, a permanent solution was established: MAPI is run on a virtual PC within the process network where the firewall allows access to data from the Norwegian Meteorological Institute's web servers. The current weather measurement and forecast is then read and made available for the digital twin by Cybernetica's OPC UA server.

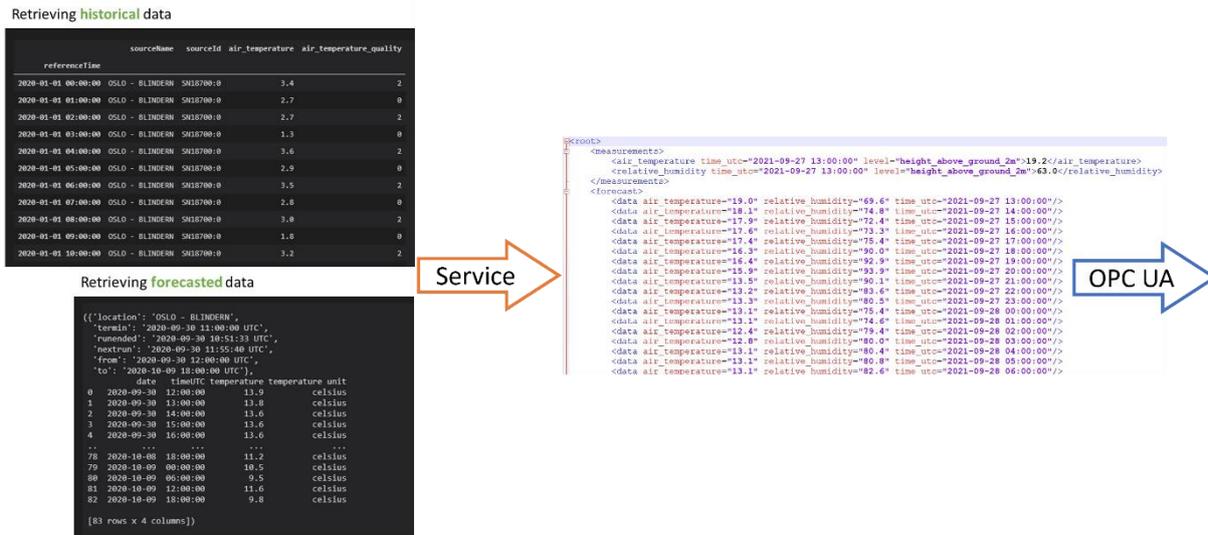


Figure 18: Schematic showing the retrieval and parsing of raw data by the MAPI interface, the preparation and formatting of the data by the Python service and the transfer of the data to the digital twin via Cybernetica OPC UA server.

The virtual PC-solution securely accesses data from the internet without exposing the Hydro process network to risks. The virtual PC is included in the network connectivity schematic in 1.4.1

1.4 Online hybrid digital twin

1.4.1 Data connectivity and system architecture

The digital twin and its supporting COGNITWIN toolbox components have been established on a server located within the process network at KTP. Most of the input data to the digital twin originates on Hydro plant servers (Hydro OPC UA, GTC OPC UA). A notable exception is the weather measurements and forecast. The MET-API Interface (MAPI) requires internet access in order to retrieve weather data and is therefore, due to security concerns, normally not accessible from within the process network, posing a challenge to the pilot as weather data is necessary for real-time digital twin operation. As described in 1.3.2, a permanent solution using a virtual PC with a selectively opened firewall has been implemented. This data retrieved by MAPI is processed by a toolbox component (Cybernetica OPC UA server, with a custom weather data extension) and the data is made available on OPC UA tags to be read by the digital twin. The data flow schematic in Figure 19 has been updated to reflect the permanent installations.

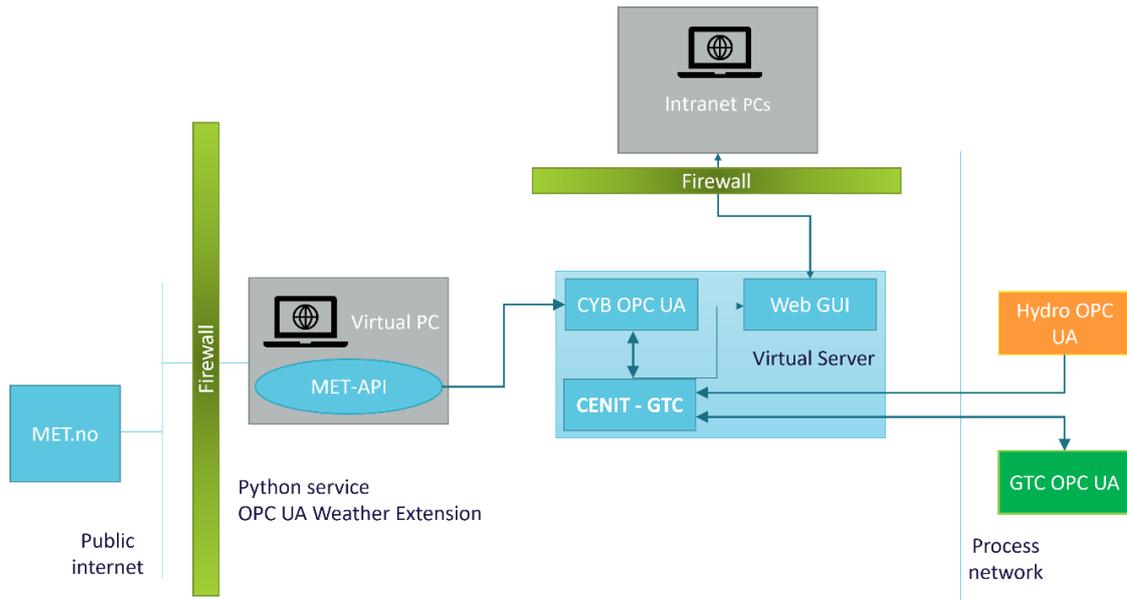


Figure 19: Data flow (including external sources) to the server running the digital twin.

Figure 20 was introduced in previous deliverables and shows the interconnectivity and handling of data by COGNITWIN toolbox components.

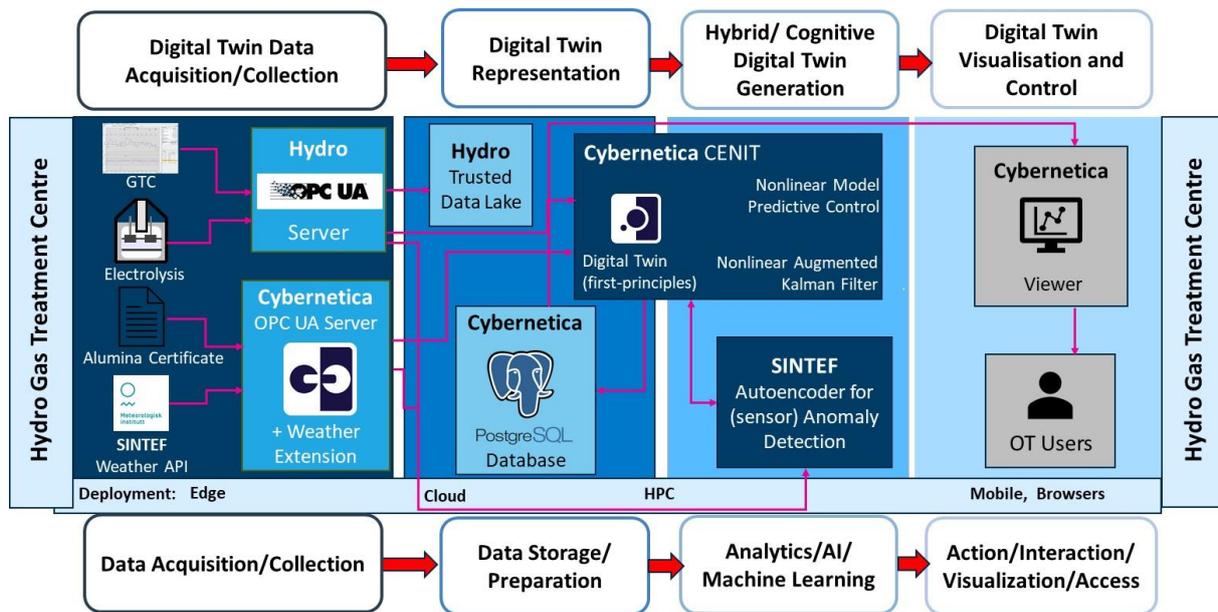


Figure 20: IoT architecture for the Hydro pilot showing interconnectivity of data and toolbox components.

1.4.2 First-principles model

A dynamic model for the GTC was introduced in Deliverable 1.2. This model was developed for online implementation as part of the pilot work in order to address the need for accurate real-time predictions of HF emissions. This model is written in C/C++ and is a continuous time model formulated as a set of differential equations for the evolution of dynamic states. As the model is intended for use in online digital twin-based optimization and control, the model must necessarily be fast to evaluate as well as numerically smooth (continuously differentiable).

Figure 21 shows a diagram of the control volumes and process flows included in the dynamic model. Conceptually, the model can be divided into main components that capture different phenomena: 1) the fumes, where HF evolution takes place, 2) the filter chamber (adsorption reactor and bag filter), where HF adsorption takes place, and 3) the alumina silos, which account for most of the time delay in the GTC. The model for the fumes is a combination of earlier literature modelling [2], experimental [3, 4] studies of HF evolution and representations of GTC components developed as part of this case.

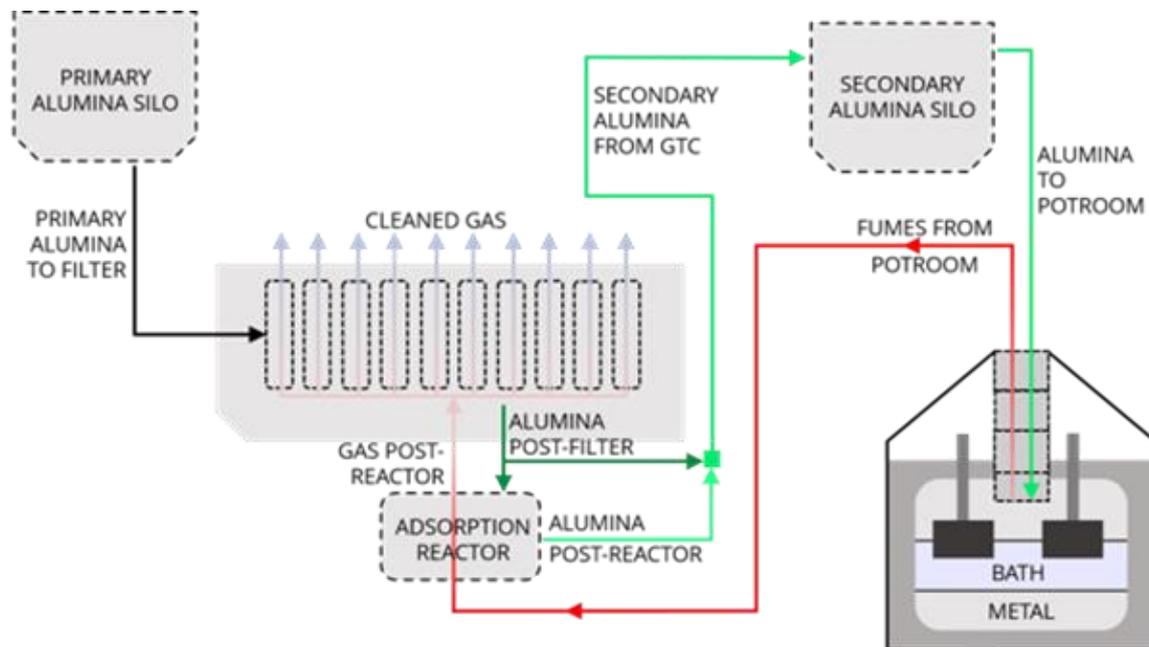


Figure 21: Diagram of dynamic GTC model control volumes and process flows. Alumina flows are shown in black (low fluoride) and green (high fluoride). Gas flows are shown in red (high fluoride) and blue (low fluoride). Control volumes are marked with dashed lines.

The model currently contains 51 dynamic states and is configured to run with the CVODE solver for ordinary differential equations. These combined improvements from the previous version (86 dynamic states with Forward Euler integration) lead to a 50X+ model speed-up. Three dynamic model parameters have also been included (discussed in section 1.4.3)

1.4.3 Hybrid digital twin - model adaptivity

Deliverable 1.3 introduced the estimation strategy used by the hybrid digital twin to achieve self-adaptive behavior and best follow process data. The following text describes the current status and highlights updates to the augmented Kalman filter used for state and parameter estimation. Three measurements are used by the Kalman filter in order to adapt model parameters to best follow process data.

The parameter with the most direct importance to model results is related to the Loss-on-Ignition (LOI – structural water in alumina). Currently, the LOI information fed to the digital twin is static as the alumina certificate has not been fully integrated with the online data flow. Figure 22 shows digital twin results when the first-principles model is configured to run with the Kalman filter enabled – the change in the anonymised LOI-related estimation parameter is shown in the bottom panel. Specific logic has been coded into the model software in order to enable the Kalman filter to handle the missing/anomalous data described above. Significant improvement is seen in the model's prediction

of the most relevant process measurement (HF concentration in gas leaving pot fumes). Tested and compared over the same period, the digital twin with enabled Kalman filter achieved a prediction-measurement correlation for this time period of 0.797 compared to 0.605 achieved by the ballistic digital twin (model run without adaptation).

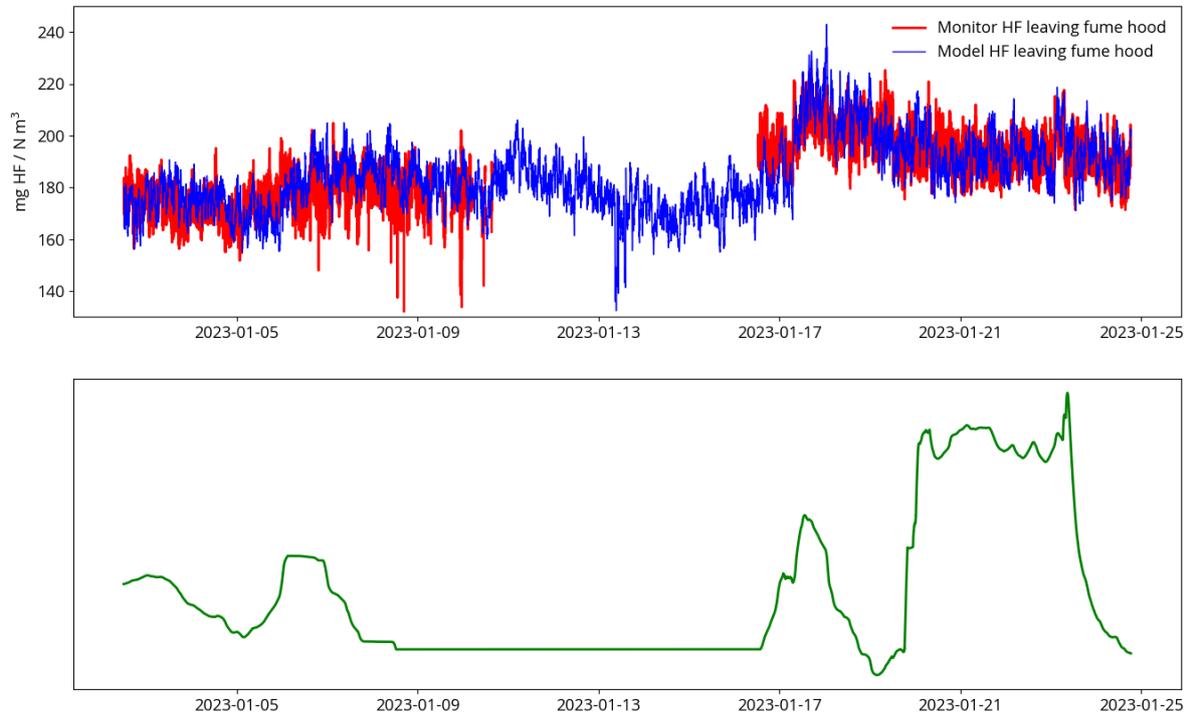


Figure 22: Dynamic model results for several weeks in January 2023 with adaptive LOI-parameter estimation.

The augmented Kalman filter also includes the following measurement/parameter pairs:

- (measurement) HF emissions from the GTC / (parameter) saturation concentration of HF in alumina
- (measurement) secondary silo height / (parameter) additive bias parameter for primary alumina feed to the GTC

Model predictions vs. online data for these two measurements as well as adaptive parameter estimation to best capture their behaviour are shown in Figure 23 and Figure 24, respectively. Model adaption to capture HF emissions levels from the GTC is relevant for ensuring the correct fraction of HF uptake. Model adaption to silo level measurements is important for capturing the correct residence time of HF and moisture through the GTC. Since the introduction of model adaption in Deliverable 1.3, the silo parameter has been modified to represent an additive bias to the primary alumina feed instead of the feed itself. This change improves the predictivity of the hybrid digital twin and was made possible by new GTC data from the primary alumina feeders.

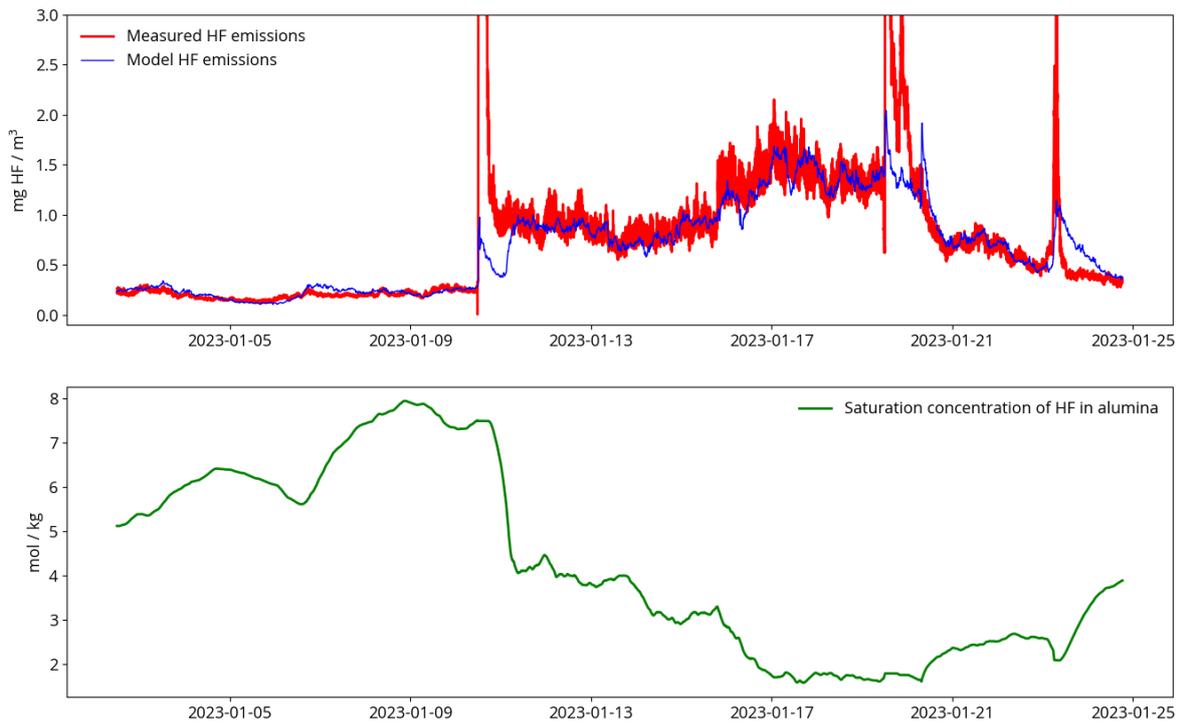


Figure 23: Model predictions vs. measurements of HF emissions from the GTC with adaptive parameter (saturation concentration of HF in alumina) estimation.

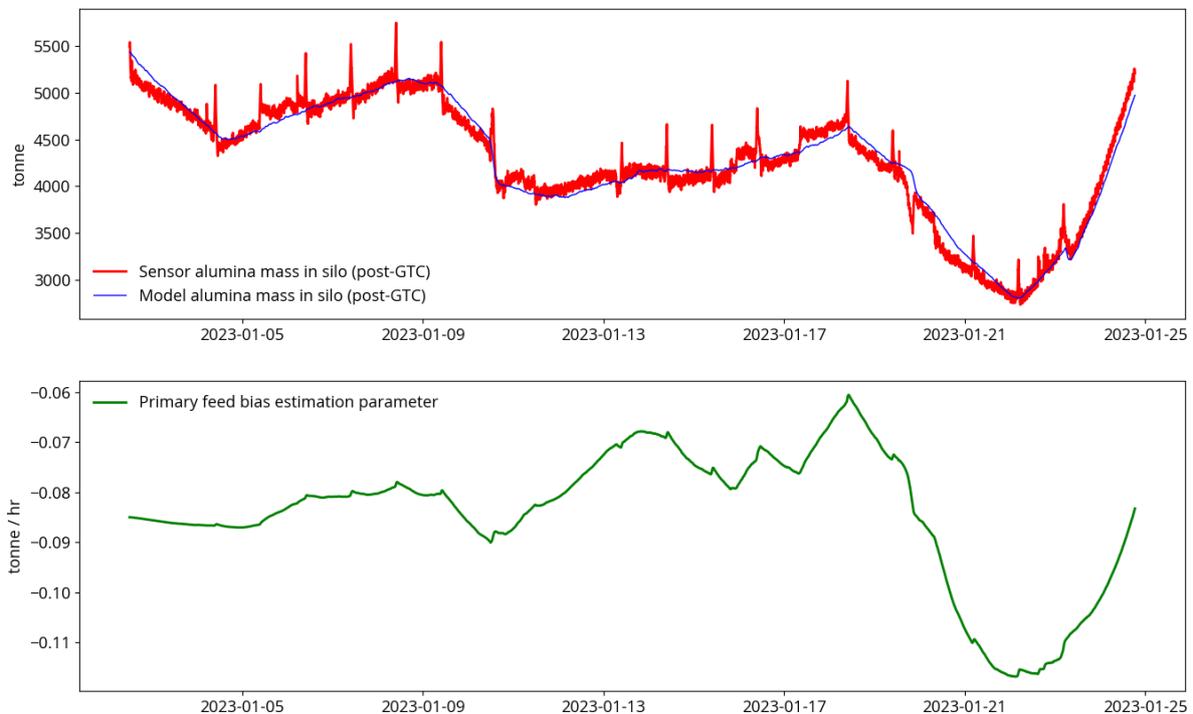


Figure 24: Model vs. sensor predictions of mass in the secondary alumina silo with adaptive parameter (primary alumina input to the GTC) estimation.

All three measurements used by the Kalman filter are prone to anomalous and/or invalid data readings. Anomalous readings are problematic when used by the Kalman filter to influence model states and

parameters. The following measures have been taken in order to minimise the impact of invalid data readings on the application performance:

- Upper and lower bounds for data readings have been defined for all three measurements
- If data readings are missing or out of bounds, a **warning** status for the affected measurement is triggered
- If a measurement remains in **warning** status for a specified number of samples (defined per measurement), an **invalid** status is triggered
- If model adaptivity based on a particular measurement is defined as critical and this measurement has **invalid** status, an **invalid** flag for the overall application is triggered, signaling for manual control and intervention to fix the problematic measurement.

1.5 Cognitive digital twin

1.5.1 Optimisation to match primary alumina feed to HF-content in alumina

Nonlinear model predictive control (NMPC) has been implemented to stabilise the HF-content in secondary alumina by optimising the primary alumina feed. The NMPC control scheme has been implemented in the Cybernetica CENIT toolbox component, where the optimisation takes the form of minimising the objective function J evaluated at specified intervals over a predictions horizon:

$$J = \frac{1}{2}(Z - Z_{ref})^T Q(Z - Z_{ref}) + \frac{1}{2}\Delta U^T S \Delta U + R^T \varepsilon$$

The first term in the objective function represents deviation of the controlled variables (CVs) Z from their set points Z_{ref} penalized by weights found in the diagonal of matrix Q . The elements of the CV-vector Z are:

- Z_1 or CV_1 : HF-content in secondary alumina (mol/kg)
 - Evaluated multiple times during the predictions horizon
 - Set point based on operator experience
- Z_2 or CV_2 : Cumulative variance of HF-content in secondary alumina (mol / kg)²
 - Evaluated at the end of the predictions horizon
 - Set point is zero (minimisation)

The result of tuned optimisation with these two CVs is to achieve a balance between reaching the set point for HF in secondary alumina and keeping the HF-content stable over the evaluated predictions horizon.

The second term in the objective function represents changes to the manipulated variables (MVs) U penalized by weights found in the diagonal of matrix S . The elements of the MV-vector U are:

- U_1 or MV_1 : Primary alumina feed to the GTC (tonne/hr)

The result of penalizing the MV changes by weights S is to ensure that the NMPC proposes balanced, stable step changes to the primary alumina feed in order to achieve the CV-set points.

The third term in the objective function represents maximal constraint violations ϵ penalized by linear weights R . The elements of the constraint vector ϵ are:

- ϵ_1 : Static maximum/minimum allowed secondary alumina silo levels (mm)
- ϵ_2 : Cyclical maximum/minimum allowed secondary alumina silo levels (mm)

The result of implementing the above constraints is to ensure that the NMPC proposes a solution that always adheres to operation requirements. The cyclical silo level constraints will be discussed more in 1.5.2.

The NMPC uses the following pieces of information to evaluate the predictions horizon:

- Current process operating conditions
- Anticipated future operating conditions (weather forecast, other known process disturbances)
- First-principles model predictions (past and future) of GTC states

The online GTC application is currently configured with a 3.5-day long predictions horizon.

1.5.2 Cognitive silo level constraint for avoidance of segregation effects

Experience-based manual GTC operation involves taking measures to minimize negative consequences of alumina segregation in the secondary silo. This is in practice achieved by periodically increasing and decreasing the silo alumina level. Because the flow of alumina to the electrolysis process is relatively constant, the need to vary silo levels introduces an inherent constraint on the primary alumina feed. By defining cyclical constraints for the alumina silo levels, the experience-based operation profile can be introduced to the digital twin optimisation solution.

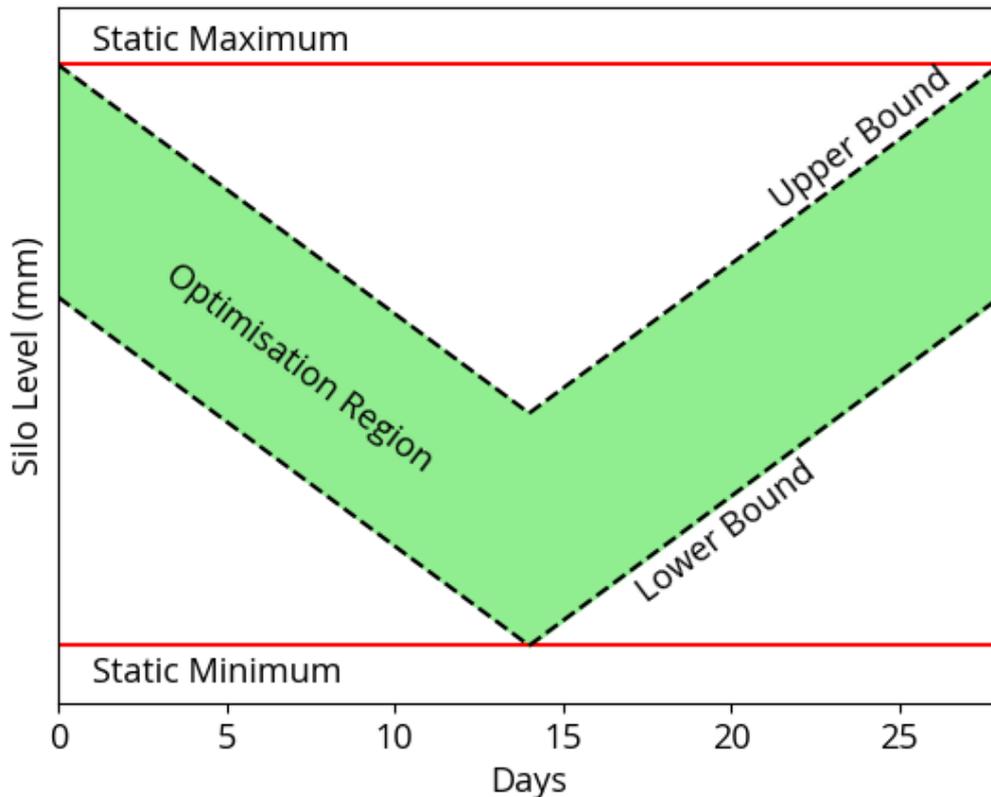


Figure 25: Schematic of cognitive (cyclical) silo level constraint. The total period for a silo cycle is configured to last for 28 days.

Figure 25 shows a schematic of the cyclical silo level constraint. Static maximum and minimum values for the silo level are defined and are always active during the optimisation in addition to the cyclical constraints. Cyclical upper and lower bounds are then calculated based on the static maximum/minimum and an allowable bound. The optimisation is then allowed to propose primary alumina feed solutions that result in silo levels between the upper and lower bounds. A larger optimisation region gives the optimisation more degrees of freedom when attempting to match the primary alumina feed to the HF-content in alumina.

The cyclical silo constraint is related to operator experience and process knowledge beyond the scope of the first-principles model. As such, it can be considered as a cognitive element of the digital twin.

1.5.3 User interface / dashboard

GTC HF Monitoring & Control

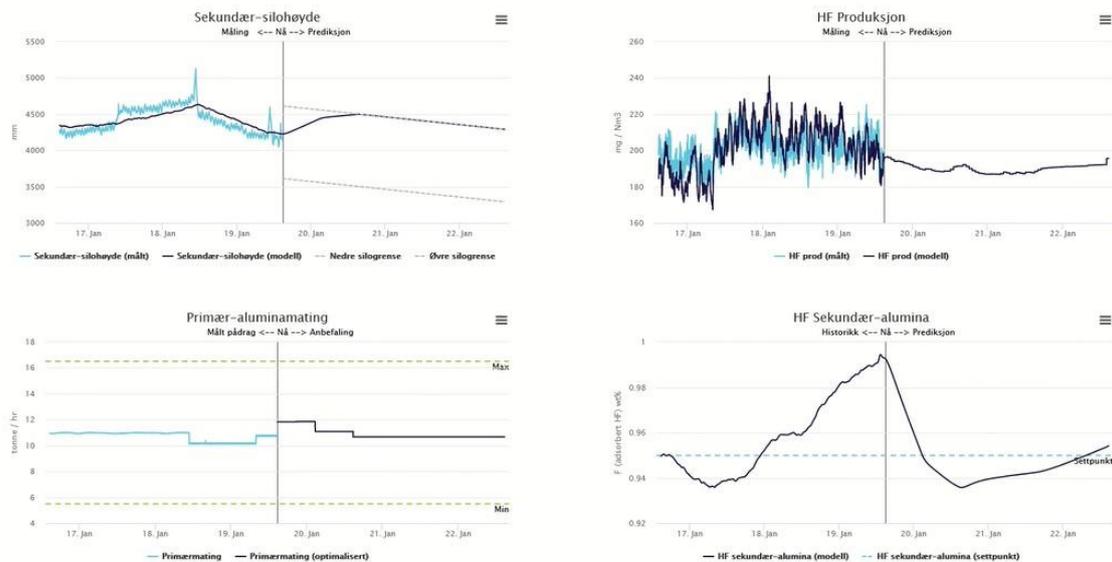


Figure 26: Preview (screenshot) of web-based dashboard for the Digital Twin operator interface.

In order to facilitate oversight and utilisation of the digital twin application by operators, a custom graphical user interface (digital twin dashboard) has been developed. Figure 26 shows screenshots of the main page of the dashboard (with Norwegian language text). This page is accessible for operators via web server.

Descriptions of the data displayed in the operator dashboard:

- Top left: Model and sensor silo levels, as well as future upper and lower limits for silo levels (grey dashed lines) for comparison
- Top right: Model and monitor results of the main pilot measurement (HF concentration in raw gas into GTC)
- Bottom left: Recorded and optimised primary alumina feed input
- Bottom right: Model results of the variable to be stabilised/controlled (HF content in secondary alumina) and the variable set point (blue dashed line) for comparison

In all four plots, the black vertical line divides measured and modelled *historical* data from anticipated *future* model predictions. Variations in *future* data is based on a combination of the weather forecast and the model's prediction for dynamic GTC states based on current and optimised process inputs.

In addition to displaying historical and future data application data, the digital twin dashboard also includes functionality for the operator to change controlled variable setpoints and enable/disable digital twin optimisation.

1.5.4 Autoencoder for anomaly detection

Toward the end of the COGNITWIN project, SINTEF developed a prototype model for anomaly detection. The basis of this work has been the data analysis carried out on the HF-gas measurements, together with Hydro, and resulted in an *autoencoder* structure that can detect anomalies in the sensor data coming from the HF laser.

Autoencoders are a particular type of architecture used in neural network-based machine learning characterised by a strong dimensionality reduction in the middle of the network; an example is shown in Figure 27. As such, an autoencoder is constructed by several layers of varying size (number of neurons or cells in each layer, i.e. the width of the layer), with the characteristic that at least one layer in the middle sports a width much smaller than that of the input or output layers, forcing the data streaming in the network through this “bottleneck”. Autoencoders are trained to replicate in the output layer exactly the same data that was provided as input (hence the *auto-* part of the name). This training, coupled with the bottleneck structure in the middle, forces the network to learn a meaningful representation of the data, which must be expressive enough to allow the correct reconstruction of the data toward the output layer, but compact enough to fit into the smaller size of the bottleneck layer. Typically, the part of the network before and including the bottleneck layer is called the *encoder*, and the part after the bottleneck is called the *decoder*. Since the decoder needs to perform the inverse operation of the encoder, it is also customary that the two parts have symmetrical structure about the bottleneck.

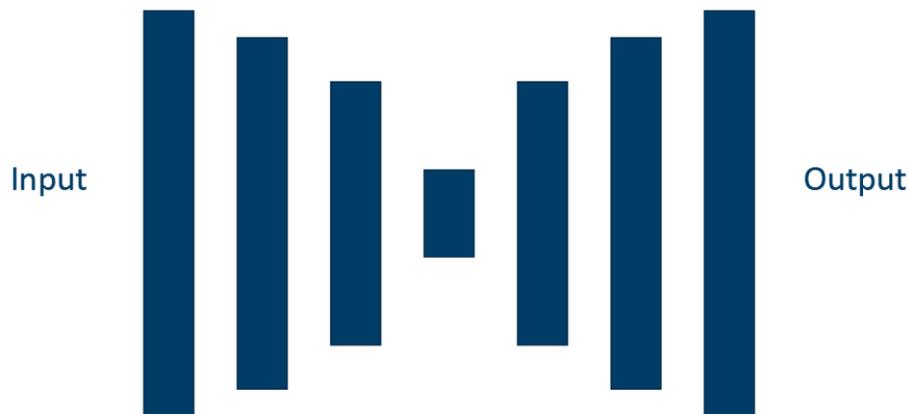


Figure 27. Example of the layers' structure in an autoencoder.

This particular structure and training makes autoencoders particularly suited for anomaly detection tasks: By training it only on data corresponding to normal operation, the network will have poor reconstructive performance when encountering data that deviates from the normal. The higher reconstruction error can then be used to flag anomalous data inputs. Further, use of autoencoders on time series permits tracking the reconstruction error over time, which can provide an indication of when the conditions are starting to deviate from the normal, thereby enabling a "warning" to be issued before an anomaly is actually encountered.

SINTEF developed and trained an autoencoder on the HF measurements from the laser, and tested it on data not used for training, as is normal in ML operations. An application of the autoencoder in a period of the data is shown in Figure 28, where we see some anomalies being detected (shaded red areas). Indeed, in those periods we notice not only abrupt changes in the monitor values, but also the difference between the monitor and the model developed by Cybernetica (indicated as *Model* in the figure), which is a further confirmation. Note that the values from Cybernetica's model are not used by the autoencoder; they are included in the figure just as a "second opinion" to evaluate the autoencoder's performance.

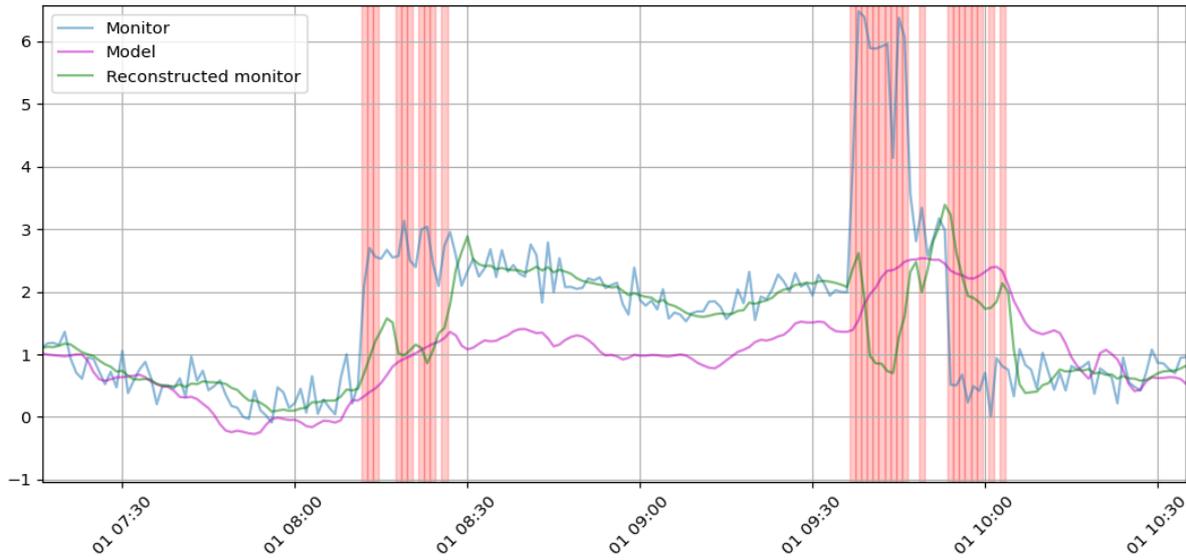


Figure 28. Detection of anomalies in the HF laser measurements (Monitor). The shaded red areas are the ones flagged by the autoencoder as anomalies. The reconstructed monitor data is also shown, together with the model data from Cybernetica's model (here displayed only for comparison, the Model data is not used by the autoencoder). We see that the autoencoder flags areas where the Monitor data behaves unexpectedly, and significantly differs from the model data.

Later, estimations of the HF-laser signal's status became available. This quantity, based on the transmission's strength of the laser, can assume three discrete values: 2 = trustable signal, 1 = unsure signal, and 0 = not trustable signal. The autoencoder was tested in this new situation, again by training it to reconstruct the Monitor signal, but without using any information from the signal's quality estimation. Results over a period of the data where the HF-laser lost connection are show in Figure 29.

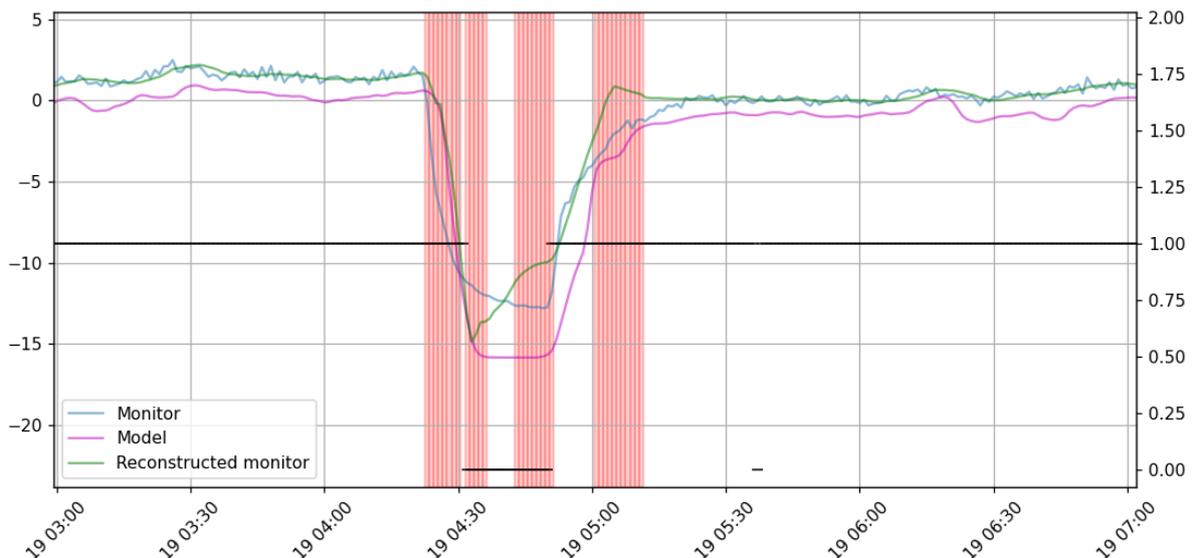


Figure 29. Detection of anomalies in the HF laser measurements (Monitor). The shaded red areas are the ones flagged by the autoencoder as anomalies. We notice that the autoencoder correctly flags the moment the HF-laser loses connection.

Here we see how the autoencoder correctly flags the moment where the laser loses connection, which is also confirmed by a signal status of 0 (black line, right vertical axis). Cybernetica's model also does a perfect job of adjusting for the lost connection.

1.6 Future work (post-COGNITWIN)

1.6.1 Testing of automatic control

The GTC digital twin has been using the Cybernetica CENIT platform to run in advisory optimization mode since the summer of 2022: operators can follow the primary alumina feed suggestions and can choose to implement them. This man-in-the-loop control strategy is valuable for testing the stability and general behavior of the digital twin, but loses the precision of closed-loop control.

The next step for advancing the contributions of the digital twin will be to test closed-loop control by allowing the digital twin to update the primary alumina feed at regular time intervals. This step is planned to continue during 2023 and 2024.

1.6.2 Container solutions

Container solutions are becoming increasingly popular in the deployment of digital applications. By standardizing the application's runtime environment, containers offer a number of important advantages, including:

- Faster deployment of new applications/rapid migration of existing applications to new servers
- Simplification of application maintenance, especially in the case where many similar applications are running online
- Overseasable control over data streams to/from the container workspace

Hydro and Cybernetica are underway with developing a workflow for deployment of Cybernetica CENIT applications via container solutions (Docker/Kubernetes).

1.6.3 Integration of Autoencoder

The autoencoder for anomaly detection in the HF-laser signal as described in Section 1.5.4 has been developed by SINTEF and tested on the available data. However, due to time constraints within the project, it has not yet been integrated in Hydro's systems or put in online functioning.

In order to achieve online operability, the development of appropriate software connections layers is required; this in order to manage an online stream of data, as opposed to offline deployment as of present. In addition, testing on more recent data (the one used in the development is from early 2022) would be appropriate.

With the above done, which is not expected to require long time, the autoencoder could be put into online service, integrated with Cybernetica's model and Hydro's data systems. This would enable the following:

1. The autoencoder would work as a secondary process control tool in parallel to Cybernetica's model, being able to detect when the HF-laser measurements could diverge from the model's predictions due to anomalies in the sensor;

2. Warnings to the operators could be issued when the autoencoder detects an anomaly in the HF-laser sensors, informing the engineers of possible sensors malfunctions or that the measurements are not to be trusted. This type of information is important both with respect to deciding on manual process steering vs automatic process steering, and maintenance of the sensor.
3. By integrating the autoencoder model into both Cybernetica's model and Hydro's data systems, it would add an additional element of *cognition* to the digital twin developed in this project, enhancing data-driven decision support.

1.6.4 Follow-up on KPIs

By the end of the project we see that the following effects can be possible:

Table 1 Summary of KPI for the Hydro pilot case

Hydro - Gas Treatment Centre (GTC)		
	Target KPIs	Achieved/achievable result
KPIs	<ul style="list-style-type: none"> • Reduce suction rate overall by 10%, i.e. for the pilot in question, 1500 MWh/y saved fan work, and increased available recovered thermal energy of 13500 MWh/y 	<ul style="list-style-type: none"> • Achieved targets
	<ul style="list-style-type: none"> • Reduce energy consumption in GTC by 15% 	<ul style="list-style-type: none"> • Achieved 10% reduction
	<ul style="list-style-type: none"> • Decrease process disturbance by preventive maintenance by 5% 	<ul style="list-style-type: none"> • Achieved 4% reduction
	<ul style="list-style-type: none"> • Maintain balanced flow distribution to different filter compartments within $\pm 5\%$ 	<ul style="list-style-type: none"> • Achieved target

2 Elkem Pilot

Some of the information in this chapter was introduced in the confidential deliverables D1.1, D1.2. and D1.3. It is also included here for the readability and context of this D1.4 public deliverable.

2.1 Introduction and background

2.1.1 Motivation

The goal for this project is to develop a dynamic, on-line model for a complete mass/energy balance for the post tap hole processing of liquid ferrosilicon, utilizing measured process data in real-time and

continuously updating the optimal process route. The model should consider cost of raw materials, product pricing as well as energy cost and optimize with respect to maximizing the profit. During the latest reporting period, the main focus for the pilot has been to install an infrared camera for the tapping station, develop algorithms for temperature and slag measurements from the refining station IR camera as well as further development of the digital twin.

2.1.2 Process description

The production of ferrosilicon (FeSi) requires multiple process steps that will be described briefly in the following. A schematic of the production route is given in Figure 30.

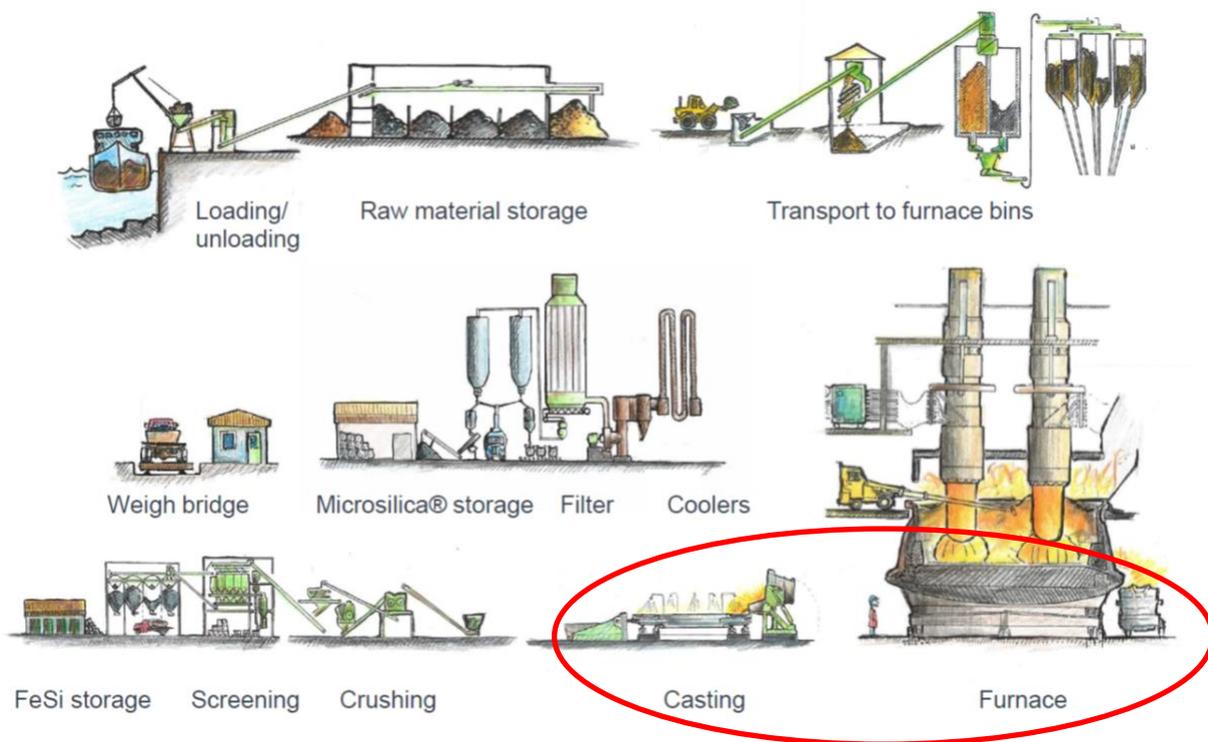


Figure 30 - Production of ferrosilicon. The focus of this project is the treatment and control of liquid FeSi after tapping of the submerged arc furnace.

The scope of the COGNITWIN project is to develop a digital twin with cognitive elements for the post taphole process (tapping-refining-alloying-casting), see Figure 31.

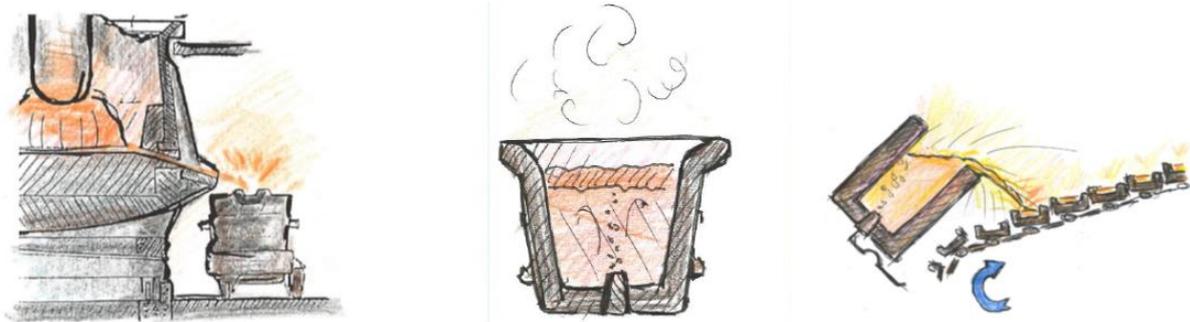


Figure 31 - Tapping, refining/alloying and casting of ferrosilicon. The ladle may (as here) or may not be equipped with a bottom plug for addition of oxygen/air.

The output of the digital twin is a dynamic description of the status of the batch process. The minimum requirement is a continuous evaluation of chemistry and temperature. Based on this evaluation, the model should suggest future actions by the operator. The goal is to maximize the post tap hole yield (tonnage cast/tonnage tapped) and meet chemical specifications for every batch. In addition, slag that follows the metal during casting should be automatically detected so that the batch can be downgraded if necessary and not shipped to customer if the slag content is too high. Most customers have an upper tolerance for slag contamination of the product. Preferably we should be able to stop the casting process if slag is detected and make necessary adjustments, for example by checking and modifying the temperature. High temperature during casting will result in a less viscous slag that will have a higher tendency to follow the metal.

2.2 Use case and challenges

Table 2 – Use case and challenges.

Use Case Template	Description
Use Case Name	Temperature in tapping stream
Use Case ID	Elkem-UC-1
User story expression of use case	<i>As process responsible for furnace operation, I want to be able to measure the temperature of the tapped ferrosilicon to establish a link between output and furnace operation and adjust operation as necessary when the temperature gets too low.</i>
Goal	Automated temperature measurement of tapping stream
Measurable KPIs for the goal (if any)	Post-taphole yield
Actors and stakeholders involved	Furnace operators, process owners
Input data	Infrared images

Output data / actions	Temperature of tapping stream
Summary description – Main success scenario	Installation of IR camera – programming – recording of data – development of algorithm to determine temperature – visual output to operator/process owner
Extensions, exceptions, variations	
Possible generalisation of use case	Setup and methods can be applied to all plants
Use case analysis – related to which Digital Twin pipeline steps	This use case will be supported through the following Digital Twin pipeline steps: Digital Twin Data Acquisition, Digital Twin Representation, Hybrid/Cognitive Digital Twin Generation,
Use Case Template	Description
Use Case Template	Description
Use Case Name	Slag in refining ladle
Use Case ID	Elkem-UC-2
User story expression of use case	As process responsible for refining and alloying, I want to know the amount of slag that accompanies the liquid metal at the start of the refining process so that proper adjustments to the operator decision support can be made
Goal	Quantify the amount of slag per ladle
Measurable KPIs for the goal (if any)	Product quality, post taphole yield
Actors and stakeholders involved	Refining operators, post taphole process owner
Input data	Infrared images of liquid metal surface in refining ladle
Output data / actions	Kilogram slag per tap – modification of total metal weight
Summary description – Main success scenario	Installation of IR camera – programming – recording of data – development of algorithm to determine slag coverage and thickness – visual output to operator/process owner
Possible generalisation of use case	Setup and methods can be applied to all plants where slag amount in ladle is important for correct liquid metal processing

Use case analysis – related to which Digital Twin pipeline steps	This use case will be supported through the following Digital Twin pipeline steps: Digital Twin Data Acquisition, Digital Twin Representation, Hybrid/Cognitive Digital Twin Generation,
--	--

Use Case Template	Description
Use Case Name	Slag in refining ladle
Use Case ID	Elkem-UC-2
User story expression of use case	As process responsible for refining and alloying, I want to know the amount of slag that accompanies the liquid metal at the start of the refining process so that proper adjustments to the operator decision support can be made
Goal	Quantify the amount of slag per ladle
Measurable KPIs for the goal (if any)	Product quality, post taphole yield
Actors and stakeholders involved	Refining operators, post taphole process owner
Input data	Infrared images of liquid metal surface in refining ladle
Output data / actions	Kilogram slag per tap – modification of total metal weight
Summary description – Main success scenario	Installation of IR camera – programming – recording of data – development of algorithm to determine slag coverage and thickness – visual output to operator/process owner
Possible generalisation of use case	Setup and methods can be applied to all plants where slag amount in ladle is important for correct liquid metal processing
Use case analysis – related to which Digital Twin pipeline steps	This use case will be supported through the following Digital Twin pipeline steps: Digital Twin Data Acquisition, Digital Twin Representation, Hybrid/Cognitive Digital Twin Generation,

2.3 Installations of IR cameras in production line

All planned installations of IR cameras are now completed. A brief description of the process, typical process conditions and integration with process control are given below for the three main sub-process related to post taphole processing of liquid ferrosilicon.

2.3.1 IR-camera – refining

After completion of the tapping process, the ladle is transported with truck to the refining station. Depending on the target specification, several elements can potentially be added to the ladle by an automatic silo-and weighing system. Because of this automatic system, there has been added a logic so that the IR camera does not interfere with automatic adding of the alloying and refining elements. There is limited amount of space and therefore the IR camera with its housing and air cooling is mounted on a pneumatic arm that works in harmony with automatic silo-and weighing system in order to avoid collision with other moving parts. Currently images are gathered on a virtual machine for offline analysis and live view is also available using manufacturers SDK through Cybernetica InSight, a custom-made software for thermal imaging from Cybernetica.

2.3.2 IR camera – tapping

Liquid ferrosilicon is tapped from the submerged arc furnace at fixed intervals depending on the production rate. The temperature of the metal is 1800-2000 °C when it leaves the furnace, accompanied with significant amounts of hot gases and smoke. These conditions are quite extreme and severely limits opportunities to measure key process parameters such as temperature. For this purpose, Infrared camera have been installed for the tapping station. This has required some infrastructure modifications such as electrical cabling and air supply for cooling. We have also built a custom housing for the camera with a mechanical shutter to protect the camera lens when it is not in use. Currently images are gathered on a virtual machine for offline analysis and live view is also available using Cybernetica InSight.

2.3.3 IR camera – casting

The ladle is transported to a tilting chair which offers both manual and automatic control of the tilting speed. The ladle is slowly emptied into a runner which leads the liquid metal to a caster which consists of multiple rectangular molds where the typical cast amount is 20 kg/mold. The molds are fixed to a moving belt. In this last part of the production process, the main emphasis is stable quality, which is achieved by keeping the casting thickness constant in all molds. In addition, slag contamination can in some cases render the entire batch unusable for the intended product grade. Downgrading the batch means value loss for the plant.

The installation of the camera was completed during 2022. A live image from the camera is available for the operator in charge of the casting. Remote access is available through the internal network.

2.4 Image analysis and algorithm development

2.4.1 Introduction and background

Sampling and data-acquisition are generally complex for metallurgical processes due to harsh environments for sensors, operators, and equipment in general. IR-cameras represent one of the currently available technologies suitable for these environments, allowing continuous measurements from afar. Though IR-cameras have previously been used for temperature measurements in ELKEM, other critical process information may be attained by analyzing the video-feeds by computer-vision algorithms. This work concerns itself with gaining access to previously unavailable process information, namely the estimation of furnace and refining slag, while additionally attaining a reliable temperature measurement during tapping and refining of liquid Si/FeSi alloys.

Prototype algorithms have been created, distinguishing between different phases by identifying transitions, also known as edges, in various parameters, like observed temperature and color. Sharper transitions create more easily identifiable edges, and as such, the greatest portion of this work dedicated to filtering out enough process noise to enable standard computer-vision methods. A large portion of this noise is attributed to effects which distort the intensity of IR-radiation reaching the camera. The temperature given to any given pixel is calculated from the intensity it is exposed to, with various effects like emissivity, reflectance, surface roughness, curvature, and viewing angle, altering the emitted intensity to appear hotter or colder. This is analogous to how shadows and powerful illumination may change the apparent color of an object, how creases and folds in a garment changes its appearance, and how various metals begin to glow at different temperatures. Thus, waves and ripples in the liquid alloy creates the illusion of massive temperature gradients, while slag and alloy measured at similar temperatures are observed to be hundreds of degrees apart. Attaining reliable temperature measurements is therefore not trivial, but while these distortions create a high degree of noise, they also create opportunities. Differences in emissivity between slag and alloy creates easily detectable transitions between the two phases, a feature exploited heavily when looking for slag in liquid alloy.

Algorithms outputting consistent temperature measurements and slag amounts have been created for both tapping and refining, with their methodology described below.

2.4.2 Refining camera

The objective was to provide an estimate on the amount of slag floating on-top of the liquid alloy, while providing a reliable temperature measurement of the liquid alloy. Algorithms for satisfactorily attaining both have been created and tested on the video-feeds made available. It should also be noted that good agreement between IR-camera measurements and mathematical modeling was found, allowing gaps in the video-feed to be filled in by the model. Refining requires a different set of solutions than tapping, as the camera looks down onto the refining ladle. Even at a slight angle, reflections from the wall necessitates removal of the visible ladle wall, with edge filtering made easier by knowing the rough ladle geometry.

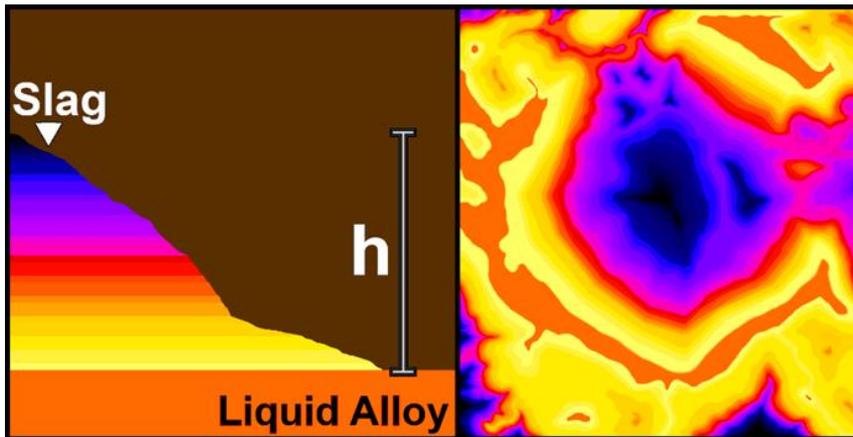


Figure 32. Depiction of how the slag thickness is estimated, with a simplified view of how the system appears to the camera to the right.

To estimate the amount of floating slag, information on the slag thickness is needed. The liquid alloy can be assumed a heat source, as its volume is order of magnitude greater than the floating slag, at constant temperature, changing with time. Heat flows from the liquid alloy to the slag and is transported through it, bleeding off heat-energy to heat the slag as it passes through. After passing through the slag, the remaining heat-energy will be emitted as radiation from the slag surface and picked up by the IR-camera. Therefore, the surface temperature at different parts of floating slag can help approximate the surface slag thickness if the liquid alloy temperature is known, as depicted in Figure 32. Figure 33 shows the algorithm being applied on a frame from an industrial video-feed.

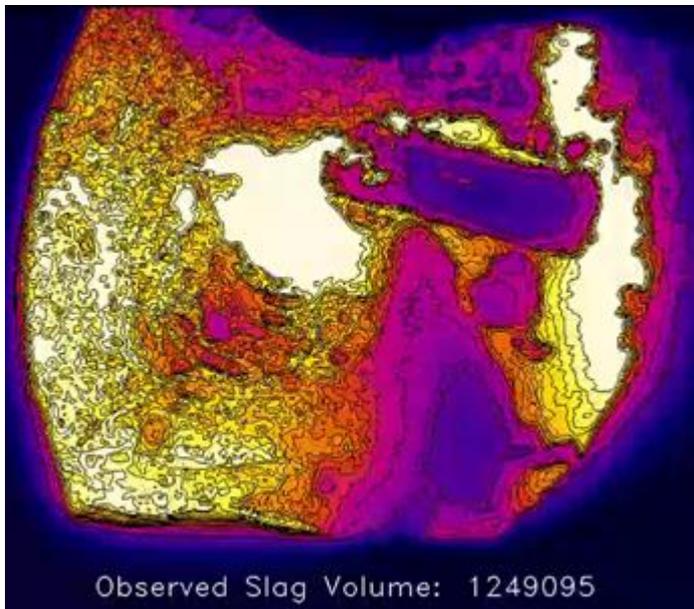


Figure 33. Surface slag divided into different temperature regions, with an uncalibrated volume estimate.

While a wide range of methods have been tried, no consistent method for separating liquid alloy and floating slag has been found in this work. Texture analysis, liquid probability factors, and different edge detection algorithms to name a few, have all been found wanting. The main issue is that slag covers most, sometimes all, of the liquid alloy, and liquid alloy with waves and thin pieces of slag floating in it shares similar intensities and textures with parts of the slag. Dust on the lens, fuming, varying process

conditions, and viewing angle, may additionally cause the liquid alloy to appear different within the same frame, adding far too many parameters to solve the problem. However, during refining floating slag gets swallowed by the liquid alloy due to a stirring current. It was found that a consistent liquid alloy temperature could be found at lower degrees of slag coverage. These measurements agree satisfactorily with the modelled temperature, which opens the possibility of substituting an initial measurement with the modelled liquid alloy temperature. The slag estimate can then be calibrated after the slag coverage is reduced and a temperature measurement is possible with the IR-camera.

2.4.3 Tapping camera

The objective was to locate slag-chunks and estimate their size, while providing a consistent temperature measurement. Algorithms for satisfactorily attaining both have been created and tested on the video-feeds made available. It was found prudent to limit the search space to only the initial tap-jet erupting from the tap-hole for improved performance. The initial tap-jet is thicker, lowering its curvature compared to the final tap-jet, with a viewing angle allowing high contrast between the liquid alloy jet and its surroundings. This highlights the importance of camera placement, where changing the viewing angle may make a feature easier or impossible to locate, much like in normal photography.

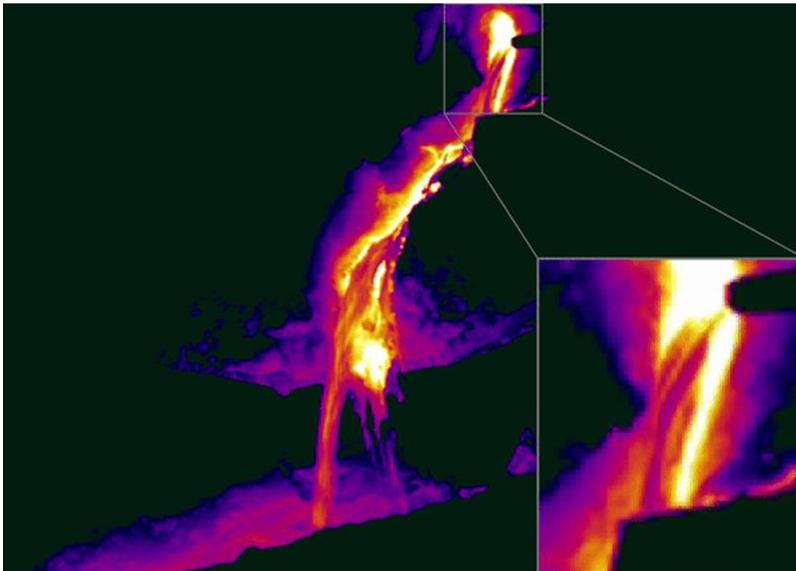


Figure 34 - Infrared image from tapping furnace #2.

While multiple methods were tried, the best results for both slag detection and temperature measurement were obtained by first locating the tap-jet boundary. Locating and updating this boundary is not trivial. Pure computer-vision methods were found lacking, instead locating geometric features of the tap-hole and jet in conjunction with fluid dynamics to estimate said boundaries. Combining images of the tap-hole without liquid alloy together with the IR video-feed was found to simplify this greatly.

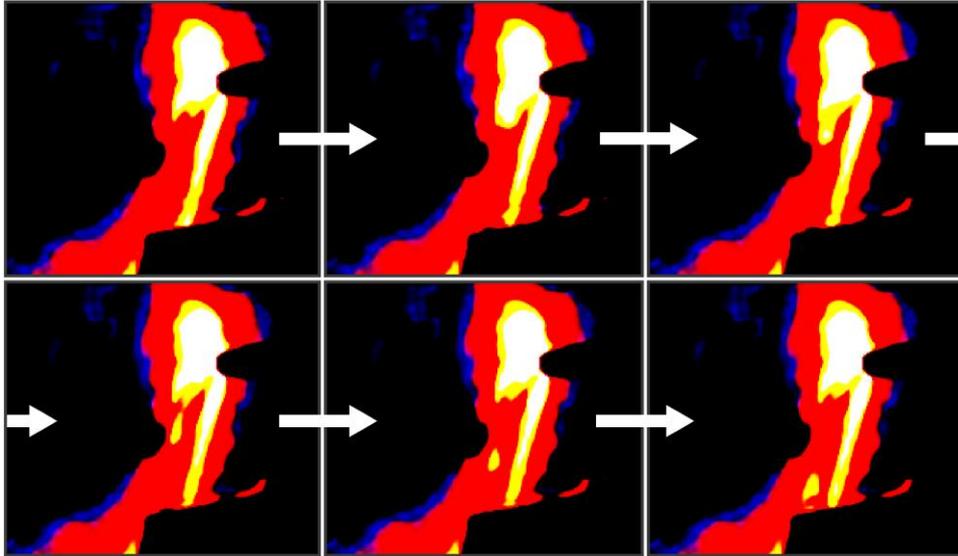


Figure 35. Depiction of a slag-chunk within the liquid alloy flow field. Arrows indicate the image order.

After locating the tap-jet boundaries, static features are masked out to focus on the dynamic parts of the frame. Slag-chunks are detected and measured by looking for high temperature features floating in the tap-jet, moving from the tap-hole to the runner and overlapping from frame to frame. Noise in the observed liquid alloy temperature signal is mostly due to flow dynamics and varying geometric shape. By following consistent streamlines, this variation may be limited as a streamline experiences similar flow dynamics along its path. A temperature profile over a streamline in a low curvature part of the jet, generally found towards its center, was found to be sufficiently consistent to fall within the confidence region of industrial measurements. Naturally, locating streamlines requires an understanding of the jet's flow patterns, which requires knowledge of the tap-hole geometry, justifying the effort spent in locating the tap-jet boundaries.

2.4.4 Casting camera

At the time of writing this report, the IR camera installed at the casting belt is operational but the development of suitable algorithm for image analysis has been delayed due to technical issues. However, as shown in Figure 36, the IR camera gives the operator a much better view of the casting process. It is much easier to detect slag contamination just by observing the IR image.



Figure 36 - Casting process as viewed from the control room (left) and through IR camera (right).

2.4.5 Infrared image analysis

The IR cameras for tapping and refining are each delivered with a GUI application for visualization and performing measurements. The refining camera, delivered by Optris, connects to the PIX-Connect application, and the tapping camera, delivered by DIAS, connects to the PyroView application.

Both of these environments provide several basic functionalities. Temperature may be sensed at single points, or across regions-of-interest (ROIs). From a ROI, the software can extract statistics such as maximum, minimum and mean temperature. Measured values may be logged at user-defined intervals to a text file (typically CSV-format) for uplink to the OPC server.

The above solution, as supplied by the device manufacturers, provides the minimal functionality needed for the digital twin models: temperature measurements of the molten fluid. It is therefore likely that the first sensor data will be provided as described above. However, the solution has several drawbacks that the project should seek to address.

Firstly, the PIX-Connect and PyroView platforms do not provide advanced or customized measurement functionality. Examples of advanced measurements that might be useful include:

- Determining what parts of an image are fluid, slag crust, or ladle walls (segmentation)
- Tracking of advected particles, e.g. for measuring flow velocity
- Image stabilization
- Measuring temporal and spatial gradients, e.g. for heat flux measurements and ladle condition monitoring

The first item is especially important, as it pertains to the quality and reliability of measurements. The current solution puts strain on downstream data QA to determine which time-series values represent fluid temperature and which ones do not. At best, this will lead to incomplete time series; at worst, the digital twin will receive and process unreliable or garbage data.

Secondly, the provided software environments are limited to data collection from one device at a time, and only from the manufacturer's own cameras. As a plant seeks to expand its inventory of thermal

imagers, this will lead to a proliferation of drivers, applications, storage locations, VMs, and runtime environments.

For these reasons, we envision the development of a purpose-built platform, with the ability to collect data from multiple cameras and manufacturers, and process images in a highly customizable machine vision environment. A schematic of this platform is shown in Figure 37 and comprises the “image/video processing” block of the overall system architecture.

The COGNITWIN project will develop DAQ modules for each of the above-mentioned cameras; however, all above data layers will be generic to allow for easy introduction of new camera manufacturers. The DAQ modules will use a combination of available drivers and libraries, and new file parsers for interpreting proprietary files.

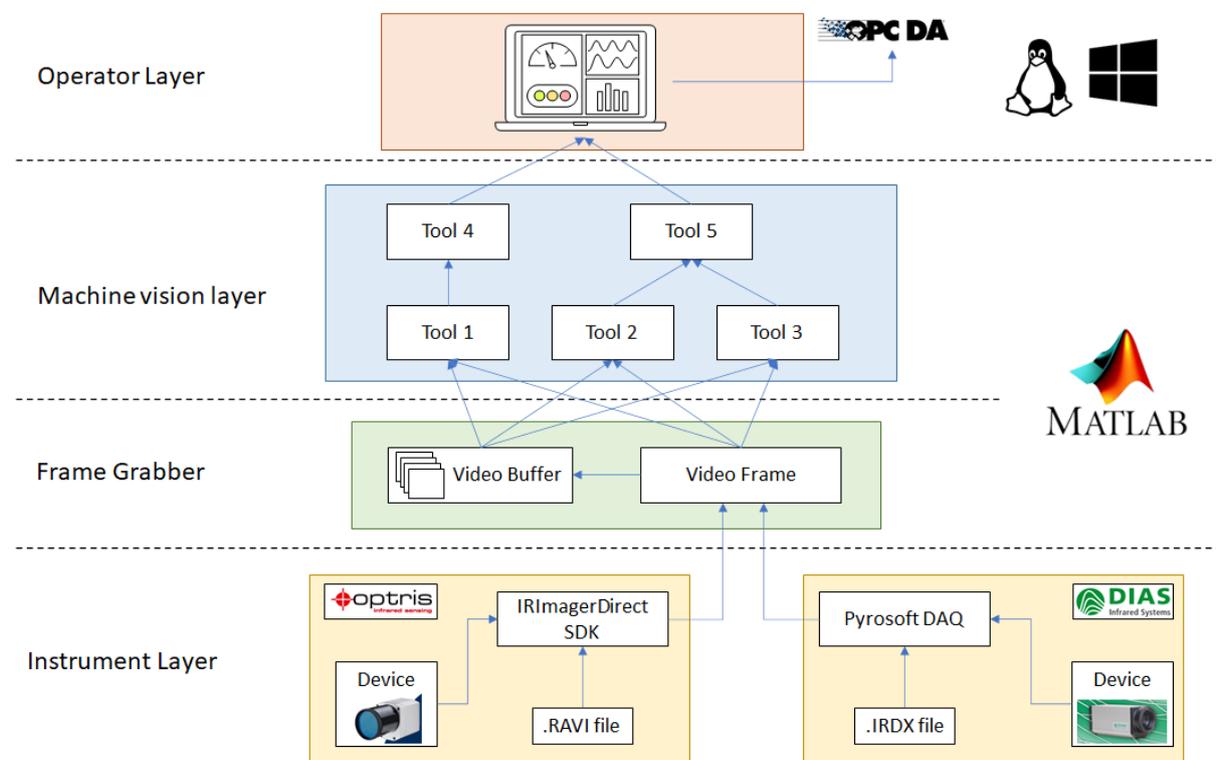


Figure 37. Proposed IR camera architecture.

The machine vision layer will provide at minimum the functionality offered by the manufacturer software: point- and ROI-measurement of temperature and logging to data files. In addition, we will implement at least one of the advanced functions listed above, and additional functions as need dictates. We expect a sub-architecture for this layer, with separate sub-layers for, e.g., resampling, analogue filtering, thresholding and binarization, morphological filtering, and data extraction.

The frame grabber layer is planned to have three functions:

- Fast processing of recorded data, e.g., for building training data sets.
- Simulated real-time processing of recorded data, e.g. for testing processing speed and handling frame dropping.

- Real-time processing of live data from device.

The operator layer will be defined toward the end of the development track, to reflect needed GUI and logging functionality based on the actual measurements developed.

2.5 Online cognitive digital twin

2.5.1 Data connectivity and system architecture

An overview of the system architecture is shown in Figure 38. Manual analyses taken at the plant are not available on OPC, but in a plant database. A bespoke extension to the Cybernetica OPC UA Server has been created to extract data from the plant database and publish it to OPC, making it available for the other components. Images from the infrared cameras are captured, analysed and stored by Cybernetica InSight, and the results of the image analyses are available on OPC. The machine learning slag model described in Deliverable 1.2 will be trained with data from the IR-cameras before it is installed online. Unfortunately, due to various operational challenges as described in section 2.4.2, the images from the refining camera were not reliable enough to train this machine-learning model. The digital twin, i.e., the first-principles refining model is running online at the plant and is used by the nonlinear model predictive controller (NMPC) to calculate optimal recommendations for the operators. The results are stored in a PostgreSQL database and presented to the users by an operator support system (Cybernetica Viewer).

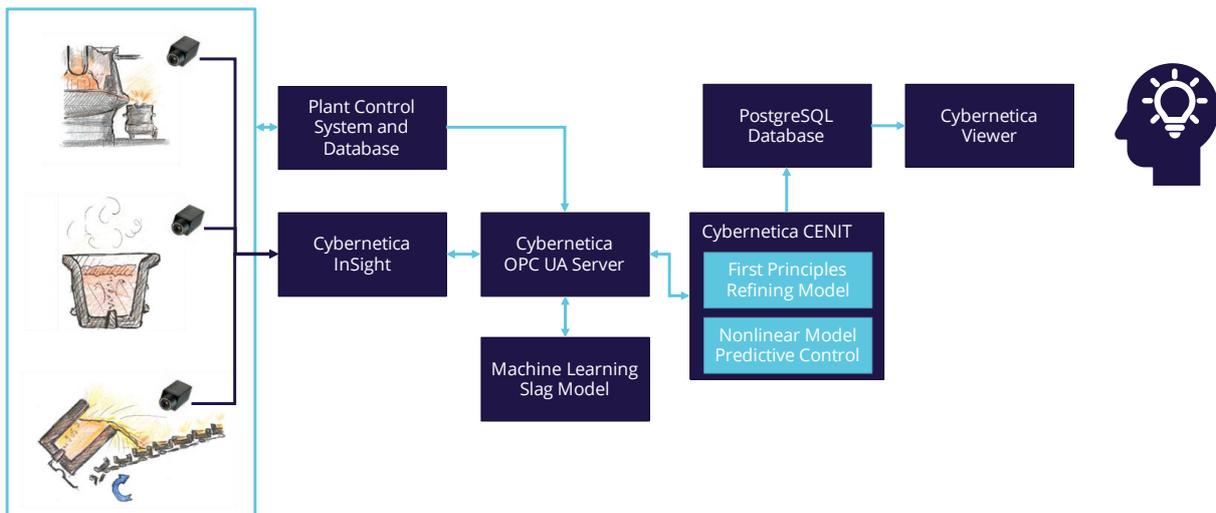


Figure 38: An overview of the system architecture for the Elkem pilot. All components are currently installed at the plant and running online except the machine learning slag model.

The digital twin pipeline for the Elkem pilot is shown in Figure 39.

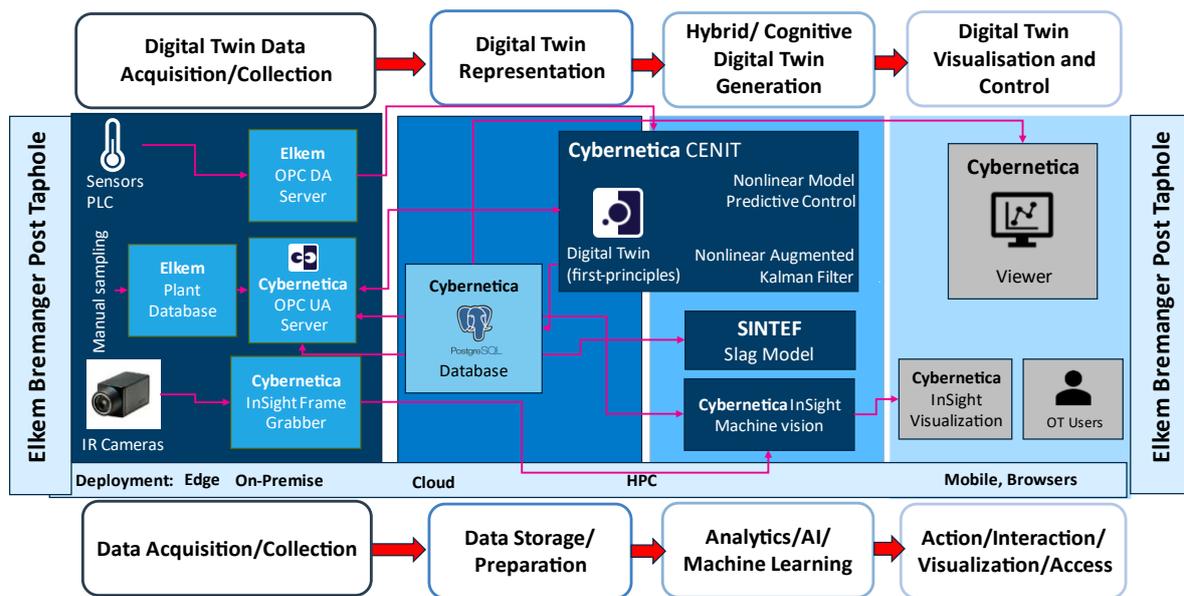


Figure 39: Digital twin pipeline for the Elkem pilot.

2.5.2 Digital twin

The digital twin of the Elkem pilot is a first-principles model of the alloying and refining process. The model is based on an existing offline model of the refining processes at Elkem. The general offline model has been made specific to the Elkem pilot. This is described in more detail in Deliverable 1.2

Considerable work has been done to include logics for initialization and handling signals specific to the pilot, as there are many manual operations for which there is no digital signal the model can utilize.

In Deliverable 1.3 we described how two datasets, one for calibration and one for validation had been utilized to fit the model to plant data and validate results. The resulting error distribution of key process variables were wider than expected, and significant effort went into finding the sources for the discrepancies between the model and the measurements at the plant. The following investigations were carried out:

- The signals for the amount of each addition fed to the ladle were delayed, meaning that they were received later in the real-time model than at the plant. After discussions with the plant's automation personnel, an acceptable solution was found for the most important material addition, the amount of cooling FeSi. For the other materials, there is still a delay, albeit less significant than previously.
- The amount of slag in the ladle after tapping can have significant variations and will affect the heat-loss of the liquid metal below the slag layer. Our goal was to estimate the amount of slag after tapping by combining images from the IR cameras and machine learning. Due to the unreliability of the refining camera, it was not possible to get an estimate of the slag amount. A simulation study was carried out to study the effect the amount of initial slag had on the metal temperature at the end of refining. This study showed that an increase of 60 kg in the amount of slag caused difference in the end temperature of almost 10 °C.

- Errors in the thermodynamic calculations for the added alloys due to unmodelled phenomena such as enthalpy of mixing, enthalpy of dissolution and intermediary phase transitions. A comparison between the model calculations and the thermodynamic calculation tool FactSage was carried out. The thermodynamic properties of certain alloys in the model were adjusted based on this comparison to account for the aforementioned unmodelled effects.
- The pilot usually operates on a “one ladle” strategy, where the same ladle is reused for each tapping and refining. Thereby keeping the metal as warm as possible by avoiding the use of fresh, cold ladles. However, there are occurrences where a ladle has to be switched out due to operational problems or maintenance. A cold ladle will then be used for the tapping and refining, which cools the metal. This effect was previously not modelled, but has now been taken into account in the model.

The result of the investigations carried out is that the standard deviation of the modelled temperature error has been reduced from 50 °C to 20 °C, which is close to the the standard deviation of the measurement itself. A histogram of the temperature error distribution is shown in Figure 40. Similar plots for the aluminium and calcium concentrations are shown in Figure 41 and Figure 42, respectively.

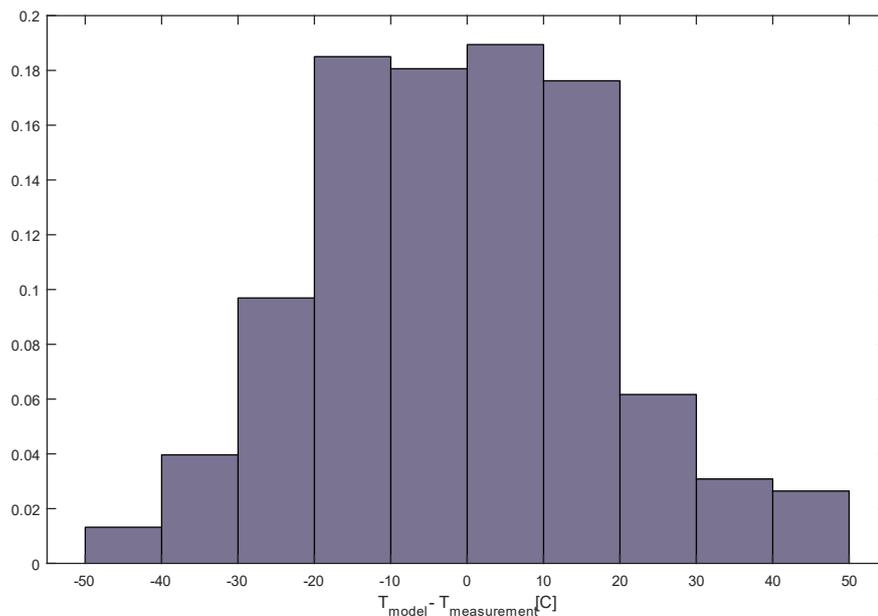


Figure 40: A histogram of the error distribution of the modelled temperature.

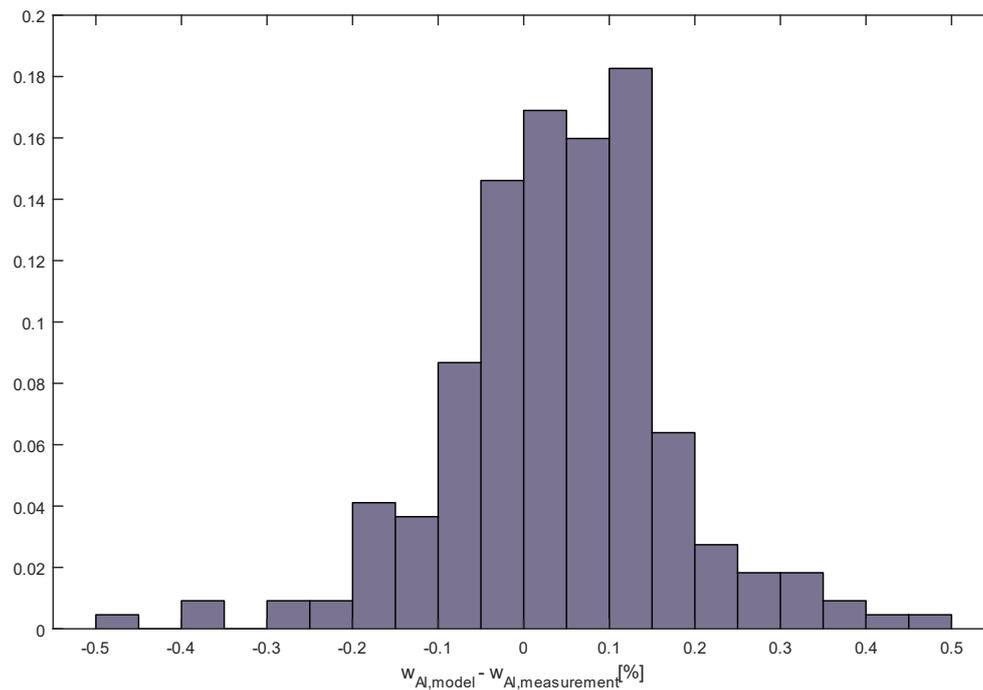


Figure 41: A histogram of the error distribution of the modelled aluminium concentration.

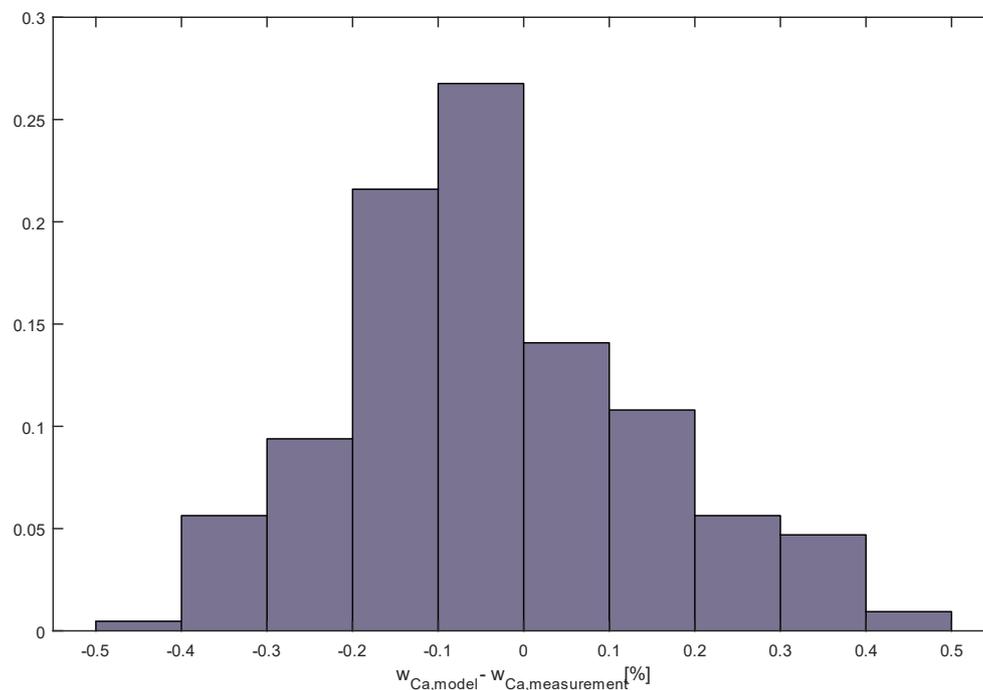


Figure 42: A histogram of the error distribution of the modelled calcium concentration.

The digital twin is installed at the plant, and is running online. The digital twin calculates the current state of the refining, and provides predictions for how the refining will develop if the operator follows the recipe given in the control system. These predictions are shown to the operators by the Operator

Support System presented in section 2.5.5. The digital twin is also used by the nonlinear model predictive controller to calculate optimal recommendations for the operators, as described in section 2.5.3.

2.5.3 Hybrid digital twin

The hybrid digital twin was planned to consist of the following components

- The data-driven slag model
- The image analysis algorithms
- The digital twin extended with an online, recursive estimator, e.g. a Kalman filter or a moving horizon estimator

Due to the IR camera installation delays, a data-driven slag model has not been possible to develop early in the project. Image analysis algorithms have thus not been tested online during periods without problems. This is being handled in the internally funded extension to the project activities.

The issues with the IR cameras have been resolved towards the end of the project, and will then further enable a data-driven slag model to provide an estimate of the initial amount of slag in the ladle to the digital twin. The amount of slag affects the equilibrium reactions in the model in addition to having an insulating effect. Having an estimate of this unmeasured variable will improve the accuracy of the digital twin with respect to the concentration and the temperature of the metal.

The results of the machine vision algorithms will be used to train the data-driven slag model offline, and will provide inputs online. The results from the camera at the refining station will also provide online measurements of the temperature and slag coverage to an online estimator.

By employing an online estimator such as a Kalman filter, the digital twin will be adapted to the conditions at the plant in real time, thereby reducing the discrepancy between the plant and the model. The estimator can be used to both mitigate discrepancies in the current batch and account for slow variations that happen over several batches.

2.5.4 Cognitive digital twin

The cognitive digital twin for the Elkem Bremanger pilot is an extension of the hybrid digital twin, with the possibility to calculate optimal recommendations for the operators. The nonlinear model predictive controller (NMPC) of Cybernetica CENIT will be used to calculate these recommendations.

The NMPC is set up to achieve the specification of the current grade with respect to concentrations and temperature and utilize as much ferrosilicon fines for recycling as possible. The NMPC realises these goals by calculating the optimal amount of each addition to be added by the operator. The NMPC uses the digital twin to predict how additions will affect the temperature and the concentrations. The results from the NMPC are shown to the operators by the Operator Support System.

2.5.5 Operator support system

The results from the Hybrid Digital Twin are presented to the operators by the Operator Support System. The Operator Support System is a graphical user interface developed with the Cybernetica Viewer framework. Figure 43 shows a conceptual mockup of the Operator Support System.

The plots displayed in Figure 43 show the initial measurement taken before refining start as a light blue cross, and the model calculation as a continuous dark blue line. The vertical black line indicates the present, with past history to the left and future predictions to the right. Two predictions are shown together in the same plot, a ballistic trajectory from the digital twin (dark blue dashes), and an optimal trajectory calculated by the NMPC (green dashes). The band near the bottom of each plot indicate where the variable must lie at the end of the batch in order to meet the specification of the grade. By showing two different trajectories together the operator can quickly decide on whether to apply the recipe or to follow the recommendations from the NMPC.

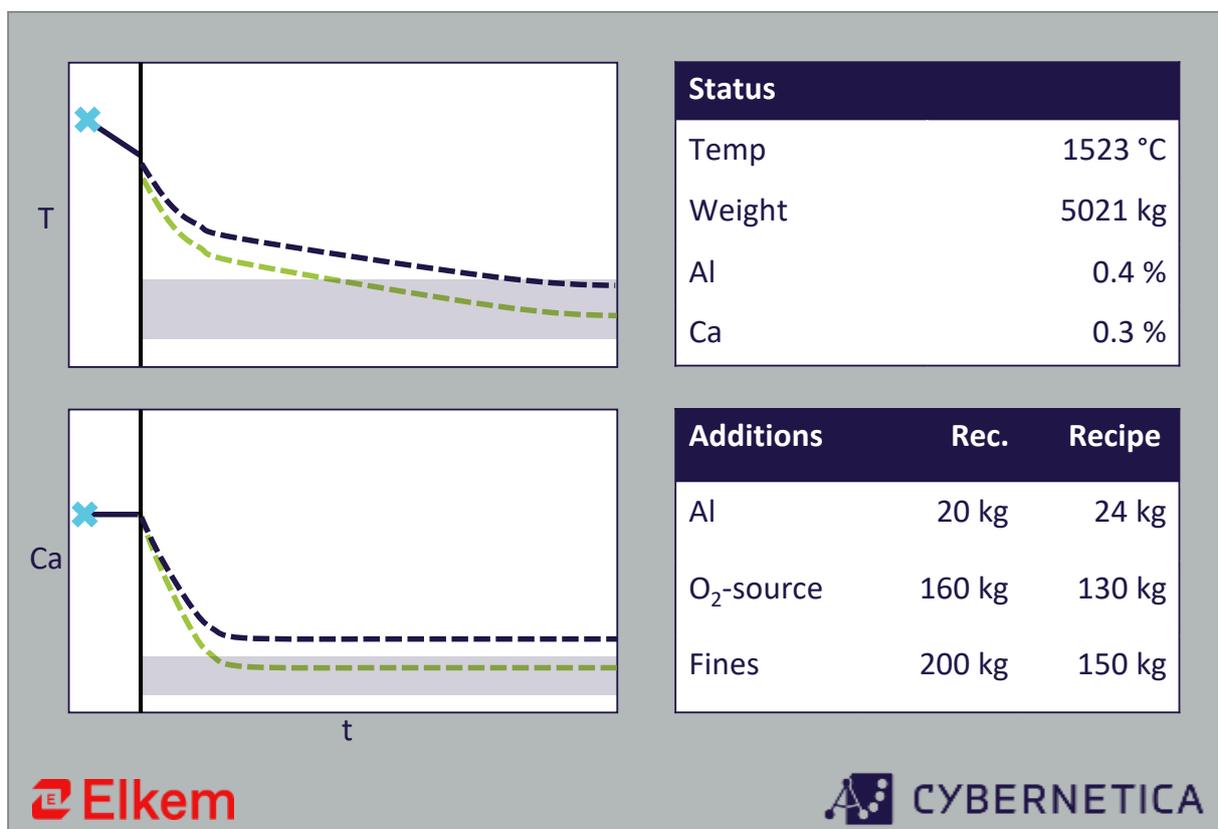


Figure 43: A mockup of the Operator Support System.

2.6 Future activities after project ends

The COGNITWIN project has opened up new opportunities for enhanced process control and production stability. Elkem will continue to build on the learnings from the COGNITWIN project, using state-of-the-art sensor technology and advanced image analysis to record and visualize key process parameters for the production of ferrosilicon alloys.

2.6.1 Camera installation

A second IR camera for direct observation of the furnace’ taphole will be installed during Q2 2023. The installation will be done at another plant in the Elkem system. The image analysis algorithms developed during the COGNITWIN project will be adapted and optimized for the new location.

2.6.2 Hybrid/cognitive twin

The main element of cognition in this project is the process of using an IR camera at the casting belt for slag detection. A high amount of slag means that the batch must be downgraded, for an economic loss. Today, the operator determines whether the amount is acceptable or not by his/her visual observations and experience. Obviously, two operators can have a different opinion regarding what constitutes too much slag. In Cognitwin, the idea was to test the following solution:

1. Determine suitable IR-camera technology
2. Purchase and install camera in casting area
3. Connect camera, optimize image and transfer video stream to local computer.
4. Develop algorithms for detection of slag in cast metal.
5. Perform image analysis to determine the amount of slag per batch.
6. Compare actual decision by operator with slag amount decided by IR camera.
7. Gradually implement cognitive decision making once proven reliable.

Some of the items on the list was completed as planned during the project period. The main reasons for not being able to fully test the solution (i.e. complete item 6) were lack of personell combined with unexpected technical issues related to data transfer needs which required additional investments in the data infrastructure. It was necessary to resolve the technical issues in order to make any headway. As of now, the technical issue is in principle resolved but full integration in the plant’s control system still awaits support from local IT. The R&D department in Elkem will complete this part of the project so that the cognitive element can be added to the hybrid twin. *Elkem has allocated additional internal funding to COGNITWIN project partner Cybernetica to evolve and continue with the deployment for one more year – until summer 2024.*

2.6.3 Follow-up on KPIs

In order to determine whether KPI-targets have been met successfully, it is required that the online model can be run continuously over a longer period of time (6-12 months). As the project comes to an end, the presumed positive effect of the completed camera installations combined with an online model for energy and mass balance will be evaluated further during the internal extended activities of the project result deployment. Individually, both of these have given valuable new information that is used on a daily basis. As of now, the goals (shown in 3) reads:

Table 3 – Summary of KPI for the Elkem pilot case. Numbers are relative, i.e. 100 = current average (%).

Table 3 ELKEM pilot KPIs summary ELKEM – Post taphole silicon yield		
	Target	Achieved

<p>KPIs</p>	<ul style="list-style-type: none"> • Increase post taphole output by up to 1% - for pilot plant this equals 400 MT/year for a gross value of 6.5 MNOK, or equivalent to an industry potential of 1.6 MEuro per 100 000 metric ton produced. • Increase PTH yield from 100 to 102 • Increase hit rate on intended products from 100 to 102 • Reduce energy consumption from 100 to 99 • Increase lifetime of ladles from 100 to 105 	<ol style="list-style-type: none"> 1. Calculated additional output 120 MT, i.e. 30% of the targeted value. 2. PTH yield calculated increase equals 100.4, i.e. 20% of the targeted increase. 3. Will be determined during the continued deployment activities until mid 2024. 4. Calculated energy consumption by using online model is 99.6, i.e 60% of the targeted reduction. 5. Cameras are able to collect data describing initial conditions that will be used to explain ladle lifetime.
-------------	---	--

The collection of data to further evaluate that these goals have been met is continuing in the Elkem internally funded project until mid 2024.

3 Conclusions

3.1 Conclusions Hydro Pilot

The COGNITWIN project seen from Hydro’s point of view has been very helpful in many aspects. It has very much clarified the digital level a process needs to be at to be able to harvest from digital twins and even further cognitive digital twins.

There were 3 cases in with targets for this project:

- I) Case I – Match primary alumina feed to HF content in alumina
 - a) To even out the fluoride into the secondary alumina sent to the electrolysis cells
- II) Temperature control
 - a) To optimise the adsorbtion conditions of the HF into alumina
- III) Main fan control
 - a) To reduce energy consumption in the GTC’s main fans

All the cases have met their targets:

Case I; has resulted in a digital twin predicting HF evolvment, hence able to optimise the feed of primary alumina to the GTC accordingly. The twin utilises data’s from process system, meteorological servers and also from the flow control/measurement in Case III. The twin is based on a first-principles

dynamic model. This model has been extended to be partially data-driven (hybrid), and lastly operators' experience has been incorporated into the optimisation system, i.e. a cognitive element.

Case II; the case of temperature regulation has, for all practical purposes met its objective, i.e. keeping the temperature in the reactor $90 \pm 5^\circ\text{C}$. The achieved range was $91.8 \pm 5,9^\circ\text{C}$ – a huge improvement over non-optimised performance and very much within the preferred range for HF adsorption in GTCs.

Case III; has demonstrated its value in saving energy simply by harvesting the changes in temperature. The savings during testing/measuring were found to be 7%. Moreover, the control of the main fans enables the suction rate to be lowered down towards the leakage limits of the cells, resulting in a further 15% potential reduction of energy use. The case targets have therefore been met.

Hydro consists of several production sites and will be able to bit-by-bit include elements from the pilot in all of our plants after COGNITWIN ends. Moreover, the methodologies developed/refined in COGNITWIN Hydro pilot will also applied to other purposes; for example, meteorological data will be used for pot room ventilation measurements, which is a vital part of emission control.

Lastly, integrating the autoencoder capabilities into the current system will add an additional element of *cognition* to the digital twin developed in COGNITWIN, which will supply the operators and engineers with valuable information regarding the status of the HF-laser sensor.

The final Hydro pilot demonstrator is shown in the COGNITWIN Toolbox [6] with the Hydro digital twin pipeline description [9] and the final Elkem demonstrator video [10]. This is also further described in the final public deliverable D6.4 Best "Digital Twins" practices report [7].

3.2 Conclusions Elkem Pilot

3.2.1 Cognitive twin vs hybrid twin

The main element of cognition in this project is the process of using an IR camera at the casting belt for slag detection. A high amount of slag means that the batch must be downgraded, for an economic loss. Today, the operator determines whether the amount is acceptable or not by his/her visual observations and experience. Obviously, two operators can have a different opinion regarding what constitutes too much slag. In COGNITWIN, the idea was to test the following solution:

8. Determine suitable IR-camera technology
9. Purchase and install camera in casting area
10. Connect camera, optimize image and transfer video stream to local computer.
11. Develop algorithms for detection of slag in cast metal.
12. Perform image analysis to determine the amount of slag per batch.
13. Compare actual decision by operator with slag amount decided by IR camera.
14. Gradually implement cognitive decision making once proven reliable.

Some of the items on the list was completed as planned during the project period. The main reasons for not being able to fully test the solution (i.e. complete item 6) were lack of personell, as described in D1.3 combined with unexpected technical issues related to data transfer needs which required additional investments in the data infrastructure. It was necessary to resolve the technical issues in order to make any headway. As of now, the technical issue is in principle resolved but full integration in the plant's control system still awaits support from local IT. The R&D department in Elkem will complete this part of the project so that the cognitive element can be added to the hybrid twin.

3.2.2 Summary of challenges

In this section a summary will be given to explain the various challenges that came up during the project period. In general, any pilot project will experience unforeseen issues – this is the very nature of a pilot project. To some extent, the project plan did not sufficiently plan for failure but was rather ambitious which is acceptable and proper for research projects. Although unforeseen challenges came up during the project, Elkem is still planning to finish the tasks as stated in the plan – to the extent they were not fully completed and reported in the deliverables.

3.2.2.1 Tapping camera

This IR camera was installed in 2021 and came online, although with a temporary solution later that year. The initial testing showed promising results and the image analysis algorithms were progressing well. Some videos and examples of use have already been shared with the project. Early 2022, we experienced technical issues and lost contact with the camera. The lack of relevant resources at the plant caused ongoing delays with the troubleshooting. In this period, the plant faced several operational challenges as well, meaning that the repair of the camera system dragged on. During fall of 2022 the project engaged with resources from outside the plant to make some progress on this work, and the troubleshooting could be completed. This caused delays in terms of testing the system over an extended period of time – which also affects other tasks in the project.

The project group determined that this technology was promising, but to ensure better access to resources a new camera was purchased and will be installed at a different plant in 2023. One advantage with the new plant is that the electric arc furnace is stationary as opposed to the furnace at Elkem Bremanger. The rotating furnace at Elkem Bremanger caused additional challenges as the taphole moves (albeit slowly).

3.2.2.2 Camera at refining station

This camera is installed and online. During the project, there were various challenges regarding connectivity as well as keeping the camera “alive” in a challenging industrial environment. We had issues with dust buildup on the protection glass in front of the camera, pieces of machinery crashing into the camera as well as data transfer limitations. All of this is fairly normal and to be expected for an industrial pilot project. The main delay caused by these issues was again related to resource availability, particularly for re-programming some of the control system. The project tasks would not be prioritized over operational challenges (these are decisions that are taken locally) which caused delays for the project. During 2022 we combatted this by engaging directly with the external supplier

of programming resources, which helped us resolving the issues. However, also for this camera an extended test period is required to conclude on the merits of the digital twin, and we could not complete sufficient testing within the project time. Therefore, Elkem will continue to test the hybrid model beyond the project. Additional internal funding has been secured to engage with our partner Cybernetica during 2023. The project now has access to a new trainee that will help out locally with IT-related tasks, this will be of great help and we are optimistic that we can make significant progress.

The data transfer issues have been resolved (in principle) by purchasing additional computational power, however, the installation has dragged on. Again, this relates to lack of internal resources and the prioritization of these.

3.2.2.3 Camera at casting

This camera was installed later in the project than originally planned, caused by delays with the other cameras as described above. It was decided that running too many tasks in parallel with the actual resource situation would be counter-productive. It was also recognized, based on the experience with the two other cameras, that a more robust setup for processing of data from the camera needed to be in place before there could be any progress on an online model for slag recognition during casting.

The camera is online at the end of the project and the casting process can be viewed on a screen in the operator room. The operator can use this information to make a more informed decision about the quality of the batch – this is one of the cognitive elements of the Elkem pilot. What currently is missing is processing of the data such that a “slag number” can be attached to every batch. Later, we can compare the operator’s actual decision vs “slag number” thereby allowing for a fully cognitive approach to casting of ferrosilicon.

The next step is to transfer data to the dedicated computer for image analysis, and this is assisted by a new trainee that supports this work. The algorithms are developed and ready for testing with this IR camera. Elkem is moving forward with this project internally as this is seen as an important development project that could influence all other plants in the Elkem family, i.e. the ability to use a cognitive twin for detection of slag in the cast product.

3.2.3 Summary and conclusive remarks

The main goals for Elkem during this project were

- 1) Install suitable infrared cameras for direct observation of the post taphole processing of liquid ferrosilicon
- 2) Establish an online cognitive twin for improved decision support for the operator.

Installation of 3 IR cameras has been completed according to the original plan. Video streaming from the cameras is available via remote access (all three cameras), and for two of them direct streaming is used during normal operation (refining and casting). It is anticipated that direct streaming from the tapping camera will be available during the internally extended project until mid 2024. This work was delayed due to installation of a new control room for tapping, which was not foreseen in the original plans. During the project, one of the main challenges has been the survival of equipment (both IR

cameras and infrastructure supporting the use). The industrial environment with high temperatures, smoke, flames and heavy machinery pose a continuous threat to sensitive equipment. Another challenge has been that the local data network does not have sufficient capacity to accommodate the large amount of data generated by the three IR cameras. In order to reach the full potential of the online cognitive twin, both of these issues must be properly resolved.

The online model is using all available data from the existing process control system and runs the model in parallel. The information from the IR cameras will be made available for the online model.

Elkem has taken the following planning steps for the further operational deployment of the ir cameras and COGNITWIN technologies: Allocated additional internal funding from Elkem to COGNITWIN project partner Cybernetica to evolve and continue with the deployment for one more year – until mid 2024.

The final Elkem pilot demonstrator is shown in the COGNITWIN Toolbox [6] with the Elkem digital twin pipeline description [11] and the final Elkem demonstrator video [12]. This is also further described in the final public deliverable D6.4 Best "Digital Twins" practices report [7].

4 References

- [1] Gordon Elinam Kofi Agbenyegah;” Mechanism and Kinetics of Hydrogen Fluoride Capture with Smelter Grade Alumina”, PhD Thesis, University of Auckland, June 2019
- [2] Yang, Y. J.; Hyland, M.; Wang, Z. W.; Seal, C. Modelling HF Generation in Aluminium Reduction Cell. *Canadian Metallurgical Quarterly*. 2015, 54, 149-160
- [3] Hyland, M.; Patterson, E.; Welch, B. Alumina Structural Hydroxyl as a Continuous Source of HF. *Light Metals*. 2016, 936-941
- [4] Pereander, L. M.; Stam, M. A.; Hyland, M. M.; Metson, J. B. Towards Redefining the Alumina Specifications Sheet – The Case of HF Emissions. *Light Metals*, 2011, 285-290
- [5] <https://frost.met.no/index.html>
- [6] COGNITWIN Toolbox: <https://cognitwin.github.io/toolbox/>
- [7] COGNITWIN Deliverable D6.4 Best "Digital Twin" practices report (Public)
- [8] COGNITWIN Deliverable D1.4 "A complete digital twin enabled with cognitive elements for the Non-Ferrous pilots" (Public)
- [9] Hydro Digital Twin Toolbox pipeline: <https://cognitwin.github.io/toolbox/pipelines/hydro.html>
- [10] Hydro Demonstrator: (2) [Hydro Pilot \(Final COGNITWIN Demonstrator - D1.4\) - YouTube](#)
- [11] Elkem Digital Twin Toolbox pipeline: <https://cognitwin.github.io/toolbox/pipelines/elkem.html>
- [12] Elkem Demonstrator: (2) [Elkem Pilot \(Final COGNITWIN Demonstrator - D1.4\) - YouTube](#)