



HAL
open science

Embedding reservoirs in industrial models to exploit their flexibility

Thibaut Cuvelier

► **To cite this version:**

Thibaut Cuvelier. Embedding reservoirs in industrial models to exploit their flexibility. SN Applied Sciences, 2020, 2, pp.2171. 10.1007/s42452-020-03925-2 . hal-03053322

HAL Id: hal-03053322

<https://hal.science/hal-03053322v1>

Submitted on 11 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright



Embedding reservoirs in industrial models to exploit their flexibility

Thibaut Cuvelier¹ 

Received: 6 August 2020 / Accepted: 19 November 2020

© Springer Nature Switzerland AG 2020

Abstract

In the context of energy transition, industrial plants that heavily rely on electricity face more and more price volatility. To continue operating in these conditions, the directors become continually more willing to increase their flexibility, i.e. their ability to react to price fluctuations. This work proposes an intuitive methodology to mathematically model electro-intensive processes in order to assess their flexibility potential. To this end, we introduce the notion of *reservoir*, a storage of either material or energy, that allows models based on this paradigm to have interpretations close to the physics of the processes. The design of the reservoir methodology has three distinct goals: (1) to be easy and quick to build by an energy-sector consultant; (2) to be effortlessly converted into mixed-integer linear or nonlinear programs; (3) to be straightforward to understand by nontechnical people, thanks to their graphic nature. We apply this methodology to two industrial case studies, namely an induction furnace (linear model) and an industrial cooling installation (nonlinear model), where we can achieve significant cost savings. In both cases, the models can be quickly written using our method and solved by appropriate solver technologies.

Keywords Electro-intensive industrial process · Electricity-price volatility · Mixed-integer linear programming · Global optimisation

1 Introduction

Current industries tend to consume large quantities of energy, often electricity, during several stages of their processes: in aluminium production, to extract the material from alumina; in cement making, to crush the limestone and heat the kiln; in paper fabrication, to pulp the ground wood. An industrial site is said to be electro-intensive when the costs of energy compose a large part of the retail price. In this case, the dependency on electricity prices is very high in order to remain competitive. The literature mostly focuses on electricity in this context, as its price can be very volatile, as opposed to other fuels like natural gas.

Nevertheless, these industrial sites can often tune their processes in order to decrease their electricity consumption during the most expensive periods; they could also benefit from selling flexibility services to the electrical

network, by reducing their consumption when the operator asks for it [3]. To achieve this cost reduction, they can make use of decision-support systems based on mathematical modelling of their processes. For example, the heating of a cement kiln can be reduced when the fuel price is too high or completely stopped if the revenue from the cement is not sufficient to cover the production costs. However, the achieved consumption drop is frequently the result of a trade-off:

- in some cases, the system is operated differently to use another fuel (*fuel switching*), and this has little to no impact on the process. For instance, a cement kiln can switch from natural gas to flare gas or other industrial by-products [8]; in ethylene production, production plants can typically use several fuels including naphtha or diesel oil [15, 27];

✉ Thibaut Cuvelier, thibaut@tcuvelier.be | ¹L2S (laboratoire des signaux et systèmes), CentraleSupélec, Gif-sur-Yvette, France.



- in others, the only flexibility is to stop the process and lose some production (*load shedding*), which means that some orders cannot be fulfilled. For example, in aluminium smelting, any reduction of the electrical current fed into the electrolysis potlines implies a loss of production [33];
- other flexibility levers can be available, such as *load shifting* or *scheduling*, depending on the exact process, with varying degrees of impact on the production and the consumption. The paper industry falls into this category, with production plan reordering [10].

Indeed, the plant directors would welcome global models, describing the whole set of processes, that allow them to make the best decisions for their energy consumption based on price forecast. To this end, each industrial process in the factory should be modelled to provide a consumption–production mathematical formulation. These models do not need to offer highly detailed insights into the machines: the goal is not to have a real-time control of the physics thereof. In fact, our aim is rather to get an *estimation* of the consumption of the process depending on how it is operated. Moreover, the staff might not have the required knowledge of operational research to build such models: they have to rely on external consultants to do so.

This article proposes a generic paradigm to help conceive such approximate formulations, so that they can be used to characterise the flexibility of a given process by the means of mathematical optimisation. The low-level basic block of this paradigm is the *reservoir*, which yields simple optimisation models while having a great expressive power.

For example, an oven might be modelled as two such reservoirs: a material tank (expressed in tonnes) and an energy storage (in joules). Over time, the oven loses some thermal energy, which can be modelled as losses from the energy reservoir. The oven is mainly controlled by its temperature, constrained to some operating range; this temperature is determined as a function of both the material and the energy reservoirs. The expressive power of the proposed paradigm comes from the fact that a reservoir may represent an actual physical storage or be more abstract (such as an energy storage or any quantity useful for modelling).

The complete oven model must also contain other blocks, named *processes*: these can impact the reservoirs, such as heating the oven. Only those processes might consume energy, such as electricity or natural gas for oven burners. A set of reservoirs only represents the state of the industrial machine, while processes interact with this state.

Using reservoir-based models for industrial processes has several advantages over typical ad hoc models, especially in the context of building many such models in a short amount of time:

- Reservoir-based models are built from a small number of intuitive components (described in Sect. 2). Consultants can quickly grasp the main ideas.
- Reasoning about such a model is easier than with pure mathematical notations, as this paradigm suggests representing the processes as legible diagrams. This graphical representation helps communicate with technical experts of the industrial process and nontechnical managers.
- The resulting optimisation models are often linear mixed-integer (MILPs).
- Conversion into computer code is straightforward once the building blocks are implemented. A software implementation may even propose a graphical interface to build the complete optimisation model, similar to MATLAB/Simulink [31]. Almost no adaptations are then required to fit the parameters of such a model to known data (as highlighted in Sect. 6.1).

This article first details the developed reservoir taxonomy in Sect. 2, and some usage examples in Sect. 3, including diagrams. Even though the resulting problems are generally linear, more complex (i.e. nonlinear) behaviours can be implemented, as explained in Sect. 4. Based on this experience, we build a typology of processes in Sect. 5. Our numerical results are presented in Sect. 6, both for fitting a reservoir model to industrial data and for assessing the flexibility within an industrial site. We conclude in Sect. 7.

2 Reservoir taxonomy

To develop the aforementioned simplified models, four kinds of building blocks are needed. Each of them is directly associated with a mathematical formulation. In the following notations, **boldface** indicates optimisation variables, as opposed to constants, typeset in roman.

- The *reservoir*, which is as close to the intuition of storage as possible: the level s_t is only impacted by the inflow \mathbf{in}_t and the outflow \mathbf{out}_t .

$$s_{t+1} = s_t + \mathbf{in}_t - \mathbf{out}_t. \quad (1)$$

These inflow and outflow variables link the reservoir to other parts of a global plant model. In the paper industry, for instance, the excess of pulp production may be stored before it is used by the paper machines,

which indicates some decoupling between the production and the consumption—in other words, a source of flexibility.

- The *decaying reservoir*, whose distinct feature is to have leaks **decay**_{*t*} in its content.

$$s_{t+1} = s_t + \mathbf{in}_t - \mathbf{out}_t - \mathbf{decay}_t. \quad (2)$$

Those leaks can be modelled with any kind of mathematical relationship: they can depend on the state of the reservoir s_t , but also on the state of *other* reservoirs. For example, an industrial oven naturally loses some thermal energy over time.

- The *observer*, which is not a reservoir per se. Its state is a function of other reservoirs' state, denoted by the variables $\mathbf{u}_t, \mathbf{v}_t \dots$:

$$s_t = f(\mathbf{u}_t, \mathbf{v}_t \dots). \quad (3)$$

It is mainly useful to implement bounds based on the state of a reservoir, such as the temperature in an oven, when it is modelled with both an energy reservoir and a mass one.

- On top of these, *external processes* must be added to impact the state of the reservoirs, by imposing some value to the inflow and outflow variables. Two examples are heaters (they increase the level of a given heat reservoir) and chemical reactions (they transform some products into some others).

Each of these reservoirs may impose constraints on its level variable s_t , such as bounds (like minimum and maximum temperatures), ramping constraints (e.g. to limit the temperature variations) or process-dependent constraints (an oven cannot be tapped before its content is melted, for instance).

The flows between reservoirs may need to be coupled, especially when a given process is modelled as multiple reservoirs. This kind of constraint is common for heat transfer, for example: the quantity of heat transferred is proportional to the flow of heat-transfer fluid.

Aggregating those building blocks constructs models for the whole process. The final step to mathematically formulate the complete plant is to assemble the various systems in one model.

3 Reservoir models

This section shows a series of models that can be obtained with the methodology described in Sect. 2. All the examples are taken from the industry, with HVAC (Sect. 3.1) being present in many sites; ovens (Sect. 3.2) are mostly present in metallurgy, whereas electrolysis (Sect. 3.3) is a

major process in chemistry. These examples are all electro-intensive processes: a typical metallurgical EAF is around 30 MW, while it is not uncommon to have electrolysis pot-lines of 50 MW.

Our goal is to produce *simple and approximate* models: they are used to detect the flexibility potential of a plant, not to perform complex real-time control on the processes. This is why linear models are usually good enough for our purposes. Would they fail in giving a good enough estimation, nonlinearities would have to be introduced (as done in Sect. 4).

One of the main advantages of the proposed methodology is that the models can be summarised by drawings, while retaining many details. This allows easy designing of models, while the actual formulation for the links between the reservoirs must be dealt with in more details outside the drawings.

3.1 HVAC

Heating, ventilation and air conditioning (HVAC) is a kind of system often found in the industry (food processing, supermarket warehouses, pharmaceutical plants, etc.), but also in offices. This kind of process is usually not the most electro-intensive that can be found, but the industrial partners often agree with its flexibilisation (as opposed to their main business).


The proposed model works on a single *energy reservoir*, corresponding to the room whose temperature is controlled (or a set of rooms, or an industrial shed), as shown in Fig. 1. Exploiting the thermal inertia of the area, the HVAC system can be turned off to save on energy while keeping an acceptable temperature.

A series of *external processes* heats or cools down the room by adding or removing energy (heating and air conditioning): they are the inputs and outputs of the reservoir. Moreover, the reservoir level decreases due to thermal losses, which depend on the temperature difference between the inside and the outside of the room. An *observer reservoir* is used to control directly the room temperature.

The corresponding mathematical formulation is the following.

- The standard *decaying reservoir* model is used for the energy. It defines the main decision variable of this system: the energy within the room, denoted by \mathbf{E}_t at time t . Its state equation takes into account the heater (**heating**_{*t*}), the chiller (**ac**_{*t*}) and thermal losses (**losses**_{*t*}; these may very well be negative, if the outside temperature is higher than that of the room).

Fig. 1 An HVAC installation modelled with reservoirs



$$E_{t+1} = E_t + \text{heating}_t - \text{ac}_t - \text{losses}_t. \tag{4}$$

The decay quantity losses_t can be approximated as directly proportional to the temperature difference between the room (denoted by T_t) and the outside (T_t^{out}) [9], the coefficient being denoted by k .

$$\text{losses}_t = k \times (T_t - T_t^{\text{out}}). \tag{5}$$

- Using an *observer*, the temperature is linked with the energy through the specific heat of the air C_p and the mass of air in the room m (which is considered constant):

$$E_t = C_p \times m \times T_t. \tag{6}$$

This temperature has strict bounds (T_{\min} and T_{\max}) that are imposed inside the observer:

$$T_{\min} \leq T_t \leq T_{\max}. \tag{7}$$

Remark 1 As the mass of air m is considered fixed, those constraints are linear. This approximation works well, because the volume of air to heat or cool down sees almost no variations in typical HVAC conditions. Section 3.1.1 considers a use case where this hypothesis no more holds.

External processes relate the heater and the chiller to their energy consumption (for example, electricity, electricity_t , and natural gas, gas_t ; a superscript h denotes the chiller, while ac is used for the air conditioning). The exact relationships f and g are not made explicit here:

$$\text{heating}_t = f(\text{electricity}_t^h, \text{gas}_t^h), \tag{8}$$

$$\text{ac}_t = g(\text{electricity}_t^{ac}, \text{gas}_t^{ac}). \tag{9}$$

These relationships may be either exact or approximated, depending on the need for precision. A basic model for these relations could use an efficiency or a coefficient of performance to relate the energy consumption and the produced effect (either heating or air conditioning).

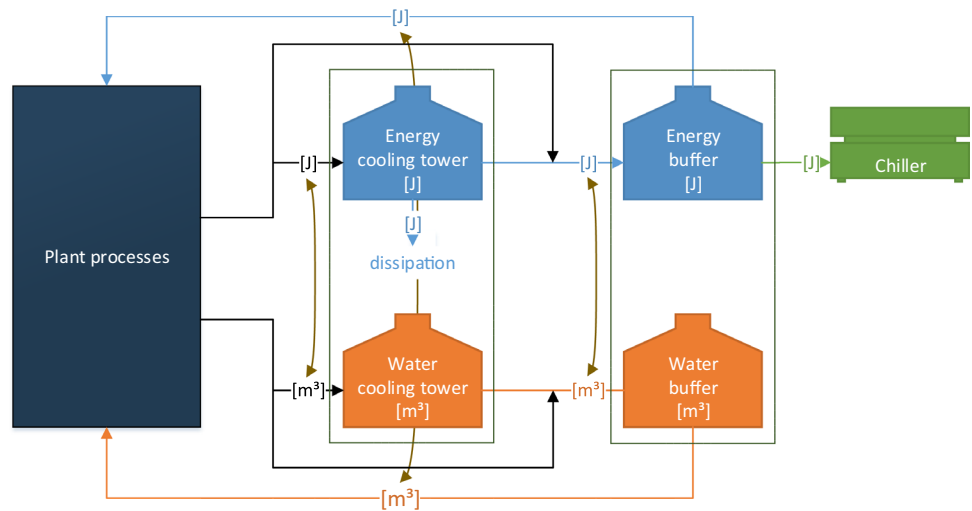
The initial conditions indicate the state of the room to model at the beginning of the optimisation horizon, which gives the initial energy (E_0) based on the initial temperature (T_0):

$$E_0 = C_p \times m \times T_0. \tag{10}$$

In a practical application, to fit this kind of model on experimental data, three important variables stand out: the energy that is pushed into or extracted from the room, the measured temperature and the outside temperature. These variables are sufficient to get the values for all the needed parameters (k for the losses and the product $C_p m$ for the temperature). With linear regressions, the parameter $C_p m$ is determined as the proportion coefficient between the energy and the temperature variations. With the same technique, the loss coefficient k may be determined from (4).

This model is simpler than existing ones in the literature [18, 19], but similar to models used in the context of flexibility [29].

Fig. 2 Industrial cooling modelled with reservoirs



Remark 2 This formulation neglects the effects rooms may have on each other. Those may be included in such a model by using an energy reservoir per room. The losses would then depend on the connected rooms.

Also, the average specific heat variations in the room are neglected, as well as the temperature uniformity hypothesis is made, i.e. the air is considered homogeneous.

3.1.1 Industrial cooling

A very similar model can be built for industrial cooling. The main differences are that no heating is performed, and that cooling can be done with two processes: either a cooling tower (whose efficiency depends on the outside temperature) or a chiller.

This model is more complicated due to the heat exchanges between the three components (the processes to cool and the two cooling mechanisms): the mass of heat-transfer fluid cannot be assumed to be constant. Indeed, the flows of hot and cooled water can be variable, but also the split between the two cooling processes. Each part of the model has two reservoirs: one for heat and the other for water, as shown in Fig. 2. The flows between those reservoirs must be coupled by the means of temperature and the water’s heat capacity C :

$$\mathbf{flow}_{\text{heat}} = C \times \mathbf{T} \times \mathbf{flow}_{\text{water}}. \tag{11}$$

The model is thus the following. Subscript CT denotes the cooling tower, p the processes that must be cooled, and B the buffer to which the chiller is connected. The arrow \rightarrow denotes a flow between two components.

- Neither the water nor the heat reservoirs are decaying, as this decay is usually very small in these applications.

$$\mathbf{water}_{CT,t+1} = \mathbf{water}_{CT,t} + \mathbf{water}_{p \rightarrow CT,t} - \mathbf{water}_{CT \rightarrow B,t}. \tag{12}$$

$$\mathbf{water}_{B,t+1} = \mathbf{water}_{B,t} + \mathbf{water}_{p \rightarrow B,t} - \mathbf{water}_{CT \rightarrow B,t}. \tag{13}$$

$$\mathbf{E}_{CT,t+1} = \mathbf{E}_{CT,t} + C_{\text{water}} \times T_{p,t} \times \mathbf{water}_{p \rightarrow CT,t} - \mathbf{E}_{CT \rightarrow B,t} - f(\mathbf{P}_{CT,t}). \tag{14}$$

$$\mathbf{E}_{B,t+1} = \mathbf{E}_{B,t} + C_{\text{water}} \times T_{p,t} \times \mathbf{water}_{p \rightarrow B,t} + \mathbf{E}_{CT \rightarrow B,t} - C_{\text{water}} \times T_{B,t} \times \mathbf{water}_{B \rightarrow p,t} - g(\mathbf{P}_{\text{chiller},t}). \tag{15}$$

- The total water and energy from the processes must correspond to the input scenario.

$$\mathbf{water}_{p \rightarrow CT,t} + \mathbf{water}_{p \rightarrow B,t} = \mathbf{water}_{p,t}. \tag{16}$$

- The flows of water and energy are linked.

$$\mathbf{E}_{CT,t} = C_{\text{water}} \times \mathbf{water}_{CT,t} \times \mathbf{T}_{CT,t}. \tag{17}$$

$$\mathbf{E}_{B,t} = C_{\text{water}} \times \mathbf{water}_{B,t} \times \mathbf{T}_{B,t}. \tag{18}$$

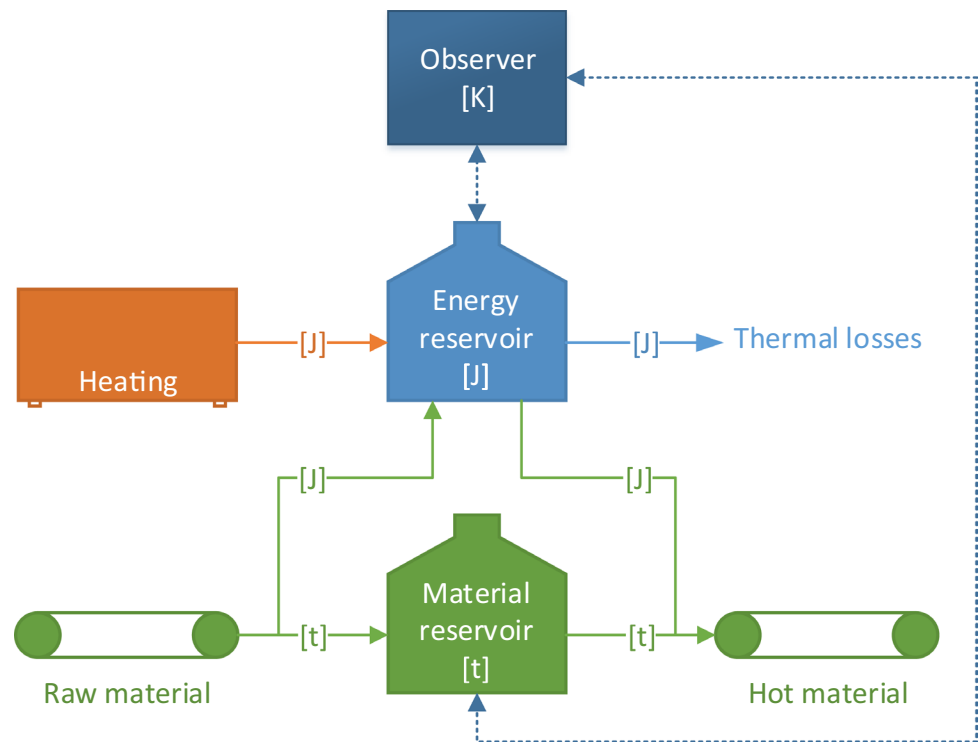
$$\mathbf{E}_{CT \rightarrow B,t} = C_{\text{water}} \times \mathbf{water}_{CT \rightarrow B,t} \times \mathbf{T}_{CT,t}. \tag{19}$$

- The buffer’s temperature must be within acceptable bounds for the processes, as its water is directly sent back to the processes.

$$T_{\min} \leq T_{B,t} \leq T_{\max}. \tag{20}$$

This model is simpler than existing ones [30], even for similar applications, and has fewer parameters to fit [26].

Fig. 3 Industrial oven as reservoirs



Actually, in Sect. 6.2, we only use manufacturer-provided data to perform flexibility estimation. The real performance may significantly differ from the specifications [25], but such a precision in the numerical results is usually not required for our use case.

A highly similar model can be used for petrochemical chains, like ethylene production [15]. More complex models can still be suited to mixed-integer linear formulations, but cannot be expressed as reservoirs; they also come with higher computational costs with traditional optimisation tools [13].

3.2 Oven

A more complex example of reservoir modelling is the industrial oven (see Fig. 3), which heats large quantities of material (several tonnes) to high temperatures (around 1000 °C), usually with a high thermal inertia (which allows for turning off or reducing the heating from time to time). Different kinds of mechanisms can be used to heat the materials: gas or oil burners (for relatively low temperatures), electric arcs (for conductive materials), etc.

An oven model is similar to HVAC in that it has an energy reservoir. However, the main difference is that the quantity of material to heat may vary significantly and thus cannot be neglected. This material is modelled as a second reservoir, coupled with the first one.

The relationships in this model are more intricate than in the previous ones: quantity of material, energy and

temperature are tightly and nonlinearly intertwined. For instance, when heating, the impact on temperature is not direct: for the same quantity of heating, the impact on temperature is less when the oven contains a large quantity of material than when it is almost empty.

In order for the oven to meet its operational goals, it must provide the needed quantities of material at the right temperature. When material is removed from the oven, it has lost an associated quantity of energy—but not temperature. However, when new material is inserted into the oven, this matter is at the outside temperature; it brings some energy into the oven, but lowers the overall temperature.

All in all, the obtained mathematical formulation is the following.

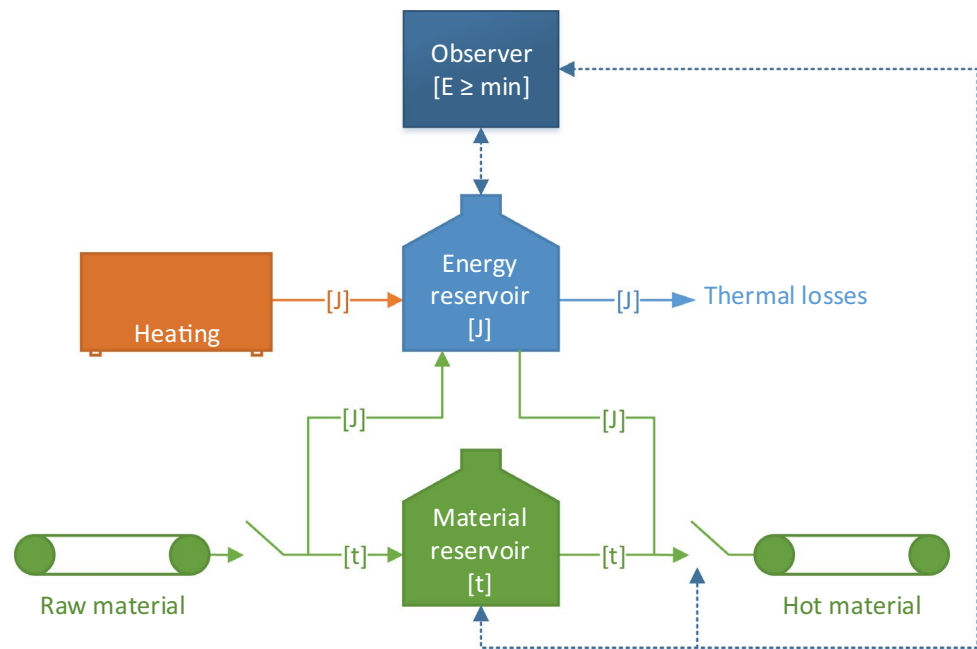
- A first *reservoir* considers the material, \mathbf{m}_t : some may be added ($\Delta\mathbf{m}_t^+$) or removed ($\Delta\mathbf{m}_t^-$) at any time step.

$$\mathbf{m}_{t+1} = \mathbf{m}_t + \Delta\mathbf{m}_t^+ - \Delta\mathbf{m}_t^- \tag{21}$$

- A *decaying reservoir* is used for the energy. Adding or removing material has an impact on the energy content (with the specific heat C_p , as previously). The withdrawn matter is at the oven temperature; however, the added material has a lower temperature, namely T_t^{out} :

$$E_0 = 0, \tag{22}$$

Fig. 4 Induction furnaces modelled with reservoirs



$$\begin{aligned}
 \mathbf{E}_{t+1} = & \mathbf{E}_t + \text{heating}_t - \text{losses}_t \\
 & + C_p \times \Delta \mathbf{m}_t^+ \times T_t^{\text{out}} \\
 & - C_p \times \Delta \mathbf{m}_t^- \times T_t
 \end{aligned}
 \tag{23}$$

The losses take the same form as previously, being linked to the temperature difference with the exterior:

$$\text{losses}_t = k \times (T_t - T_t^{\text{out}}).
 \tag{24}$$

- Likewise, an *observer* gives the temperature, but this time *nonlinearly*:

$$\mathbf{E}_t = C_p \times \mathbf{m}_t \times T_t.
 \tag{25}$$

This nonlinearity is not a problem for the temperature bounds, as constraint (7) can be rewritten in joules-kilograms *without* using constraint (25):

$$C_p \times T_{\min} \times \mathbf{m}_t \leq \mathbf{E}_t \leq C_p \times T_{\max} \times \mathbf{m}_t,
 \tag{26}$$

- The heater is still an *external process* that consumes energy:

$$\text{heating}_t = f(\text{electricity}_t, \text{gas}_t).
 \tag{27}$$

As opposed to the HVAC model, this formulation cannot be linear due to the losses, as they involve temperature (24), i.e. the ratio between the energy and the mass. These issues are discussed in Sect. 4.

Remark 3 This model exploits the hypothesis that the specific heat C_p remains constant with temperature. Also, it does not consider phase change, i.e. it only allows to heat

material, not to melt it. A large quantity of energy being required for the material to change phase, adding this possibility in the model would require another reservoir that specifically deals with the molten part.

3.2.1 Induction furnace

Induction furnaces are a kind of industrial oven. They work by magnetic induction: the bucket containing the metal to melt (such as cast iron) is surrounded by a coil through which high-voltage alternative currents are sent. The created magnetic field induces eddy currents in the metal, which in turn heats it. The power dissipated in the metal is directly proportional to the electrical power fed into the circuit P_t^{elec} , with a constant ratio α , as shown in Appendix 1.

A reservoir model that fits this kind of furnace would be made up of two reservoirs: a mass reservoir that evolves at discrete time steps and an energy decaying reservoir (filled and emptied at the same time as the mass reservoir). There is no need for a temperature observer, but only a “binary” observer that indicates whether the required total energy has been transferred to the metal at the end of the batch. The model is shown in Fig. 4.

The corresponding model in mathematical form is the following:

- The mass reservoir is fixed for a given batch:

$$\mathbf{m}_t = m_0.
 \tag{28}$$
- The energy reservoir is filled by electrical currents:

$$E_{t+1} = E_t + \alpha P_t^{elec}. \tag{29}$$

- The observer reservoir imposes that, at the end of the batch, the total energy in the bucket is at least at the required level:

$$E_T \geq f(m_0). \tag{30}$$

- Initially, the energy reservoir indicates that the material is at the exterior temperature:

$$E_0 = C_p \times m_0 \times T_0^{out}. \tag{31}$$

$$E_{t+1} = E_t + \text{heating}_t - \text{losses}_t + C_p^{Al_2O_3} \times \Delta m_{Al_2O_3t}^+ \times T_t^{out} - C_p^{Al} \times \Delta m_{Alt}^- \times T_t + \Delta H \times \Delta m_{Alt}^+ \times (T_t - T_t^{out}) \tag{32}$$

Typical models for such furnaces deal with the details of the electromagnetic field required to melt the metal [1, 23], which is too cumbersome for our application.

3.3 Electrolysis

The chemical industry and metallurgy often use electrolysis, with aluminium smelting being a prime example. A potline is typically made of hundreds of individual baths, where the electrochemical reactions take place. Each of them can be controlled independently and turned off to save energy, but with a loss of production [33].

In front of the potline, there is a single transformer from AC to DC current; its output has a relatively low voltage (around 5 V, usually), but very high amperage (multiple hundreds of kiloamperes are not rare) [2, 4, 33]. After the transformer, all pots resemble ovens (as in Sect. 3.2): they are heated by the current that flows through them, so that their bath remains in a given temperature range where the electrochemical reaction can happen.

The main difference with ovens is that the contents of the pots have smaller fluctuations: inputs are continuous, so the mass increases at a predictable and constant rate; the product is siphoned off periodically, typically once per day [2, 4, 33]. However, the contents of the bath continuously evolve: the alumina and the carbon anode are transformed into aluminium and gas, the actual reactions being $Al_2O_3 + 3C \rightarrow 2Al + 3CO$ and $2 Al_2O_3 + 3C \rightarrow 4Al + 3CO_2$. As such, 415 kg of carbon anode (not fed continuously, the anode being replaced after a few weeks) is consumed to produce 1 t of metallic aluminium, and this carbon is ejected as exhaust gas (carbon mono- and dioxide) [32]. Nevertheless, these exhausts are neglected in our model.

The flexibility impacts the production in a slightly more complicated way than in the aforementioned processes: it depends (approximately linearly) on the average current

through the pot, as long as the temperature is in the right range (which is controlled by voltage) [22].

As a consequence, the electrolytic bath model contains four reservoirs:

1. The energy is included as a *decaying reservoir*, exactly like in previous models (4) and (23), with some energy contribution due to the alumina input $\Delta m_{Al_2O_3t}^+$ and to the metallic aluminium output Δm_{Alt}^- , but also due to the fact that the reaction is exothermic (with ΔH being the enthalpy change due to the reaction):

Again, the coefficient C_p^X is the specific heat of material X. The thermal losses can be expressed linearly with respect to the temperature difference with the outside:

$$\text{losses}_t = k \times (T_t - T_t^{out}). \tag{33}$$

2. The carbon anode *reservoir* can only be consumed by the electrochemical reaction (it is replaced during maintenance, which we do not mean to optimise):

$$m_{Ct+1} = m_{Ct} - \Delta m_{Ct}^-. \tag{34}$$

3. The alumina *reservoir* is constantly fed and consumed by the electrochemical reaction:

$$m_{Al_2O_3t+1} = m_{Al_2O_3t} + \Delta m_{Al_2O_3t}^+ - \Delta m_{Al_2O_3t}^-. \tag{35}$$

4. The aluminium *reservoir* is filled by the electrochemical reaction and periodically emptied:

$$m_{Alt+1} = m_{Alt} + \Delta m_{Alt}^+ - \Delta m_{Alt}^-. \tag{36}$$

5. The temperature *observer* depends on the complete mass within the electrolytic bath. A simplifying assumption is to consider that the temperature is uniform within the bath:

$$E_t = \left(C_p^{Al_2O_3} \times m_{Al_2O_3t} + C_p^C \times m_{Ct} + C_p^{Al} \times m_{Alt} \right) T_t. \tag{37}$$

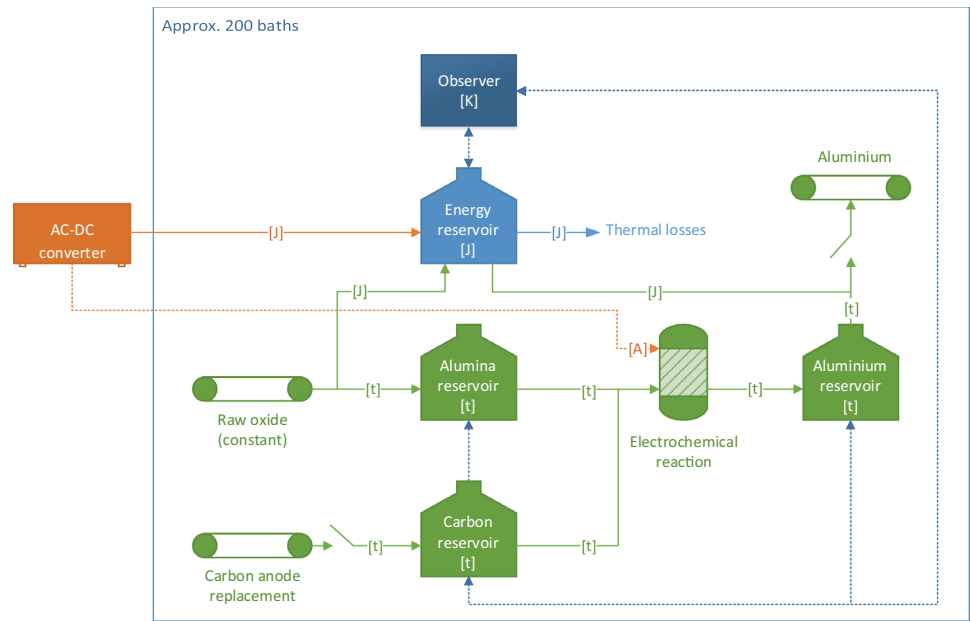
As for oven (26), the temperature bounds can be written linearly based on this expression:

$$T_{min} \times \left(C_p^{Al_2O_3} \times m_{Al_2O_3t} + C_p^C \times m_{Ct} + C_p^{Al} \times m_{Alt} \right) \leq E_t, \tag{38}$$

$$E_t \leq T_{max} \times \left(C_p^{Al_2O_3} \times m_{Al_2O_3t} + C_p^C \times m_{Ct} + C_p^{Al} \times m_{Alt} \right). \tag{39}$$

6. The last block within the process, the electrochemical reactions, is modelled as *processes*, whose conversion

Fig. 5 Electrolysis modelled with reservoirs



rates k_X'' depend on the average current through the bath (\bar{I}_t):

$$\Delta m_{Al^+} = k_{Al}'' \times \bar{I}_t, \quad \Delta m_{C^-} = k_C'' \times \bar{I}_t, \quad (40)$$

$$\Delta m_{Al_2O_3^-} = k_{Al_2O_3}'' \times \bar{I}_t. \quad (41)$$

- The AC-DC converter must be modelled as an *external process* to provide the required power to heat the bath; moreover, the actual amperage must be explicitly represented in the model for the electrochemical reaction rate:

$$\text{heating}_t = f(\text{electricity}_t, \bar{I}_t). \quad (42)$$

- Finally, the initial conditions are set according to the way the process is managed: the process is rarely started from scratch (with a temperature equal to T_0^{out}), but rather continuously operated.

The model is shown in Fig. 5. Existing models tend to use electrochemical multiphysics techniques [16, 35], which are not well-suited for light applications like flexibility estimation.

Remark 4 Observer Eq. (37) might consider directly the sum of all three material reservoirs (alumina, carbon and metallic aluminium), without a distinction between the various materials: their specific heats are of the same order of magnitude (between 700 and 900 J/K kg). This approximation is justified in the context of flexibility, as the model

does not need to be very precise. This simplification is useful when fitting the model to actual data, as fewer parameters must be estimated.

Remark 5 If the electrical current is not considered for flexibilisation, then the model can be simplified, as the electrochemical reaction then becomes constant: the mass reservoir levels only change when the bath is powered on (and thus when the electrochemical reaction takes place). The only possible values are then naturally discrete.

4 Nonlinearity

Another great advantage of the proposed methodology is that most models are linear. However, these can be limiting for some behaviours that cannot be completely represented with linear equations. For those cases, such as ovens (for which a conceptual model is presented in Sect. 3.2), nonlinearities can be introduced in the reservoir formalism.

Those nonlinearities can be dealt with in different ways.

- Use of nonlinear nonconvex solvers** Nonlinear equalities such as (37) make standard optimisation solvers unusable because of their nonconvexity. Therefore, the solvers must be able to deal with nonconvex constraints and (often) mixed-integer variables. Programs like Couenne [5], POD [24] (global solvers, not relying on the convexity assumption), or Bonmin [7] (approximating the problem as convex) must then be used. The most recent versions of standard optimisation solvers

like CPLEX [17] and Gurobi [14] now allow for certain types of nonconvexity (CPLEX 12.7 accepts nonconvex quadratic objective functions, whereas Gurobi 9.0 also tackles quadratic nonconvex constraints).

- **Linear reformulations** Another option is to formulate the problem linearly. The main technique we used for this is discretisation: instead of having continuous variables, some of them are only allowed to take discrete levels. The choice of variables for discretisation is made so that, when these variables take a fixed value, the constraint becomes linear. For example, constraint (37) is linear once the mass is known. Binary variables are then used to choose among the various discrete values, which keeps the overall linearity.

4.1 Discretisation

Discretisation is an approximation of the previous models. It can however be used in some cases without being a rough estimate, due to the actual operational conditions: for instance, industrial ovens are often loaded by batches, whose sizes can be used to define the discretisation levels.

An oven like in Sect. 3.2 can be reformulated with this discretisation approach, more specifically by rewriting the temperature definition (25). Instead of having a continuous variable \mathbf{m}_t , it takes its value in a discrete set $\mathcal{M} = \{\mathcal{M}_i | i \in [1, M]\}$. For example, instead of having an oven whose mass may freely vary between 500 kg and 5 t, its discretised version may only take the values in $\mathcal{M} = \{500, 1500, 2500, 3500, 5000\}$ kg, or $\mathcal{M} = \{500, 1000, 1500, 2000, 3500, 5000\}$ kg if more precision is needed for low masses.

In order to linearise the quotient $\mathbf{T}_t = E_t / C_p \mathbf{m}_t$, the following expression can be used:

$$\mathbf{T}_t = \frac{E_t}{C_p \mathbf{m}_t}, \quad \forall t \in \mathcal{T}$$

$$= \begin{cases} E_t / C_p \mathcal{M}_1 & \text{if } \mathbf{m}_t = \mathcal{M}_1 \\ E_t / C_p \mathcal{M}_2 & \text{if } \mathbf{m}_t = \mathcal{M}_2 \\ \vdots & \\ E_t / C_p \mathcal{M}_M & \text{if } \mathbf{m}_t = \mathcal{M}_{|\mathcal{M}|} \end{cases}, \quad \forall t \in \mathcal{T}$$

Implementing this in a mathematical optimisation model can be done in a classical way [20, 21, 34]. New variables are introduced:

- The binary variables $\mathbf{m}_t^{(i)}$ indicate whether the material contents of the oven \mathbf{m}_t take the discrete value \mathcal{M}_i .
- The continuous variable \mathbf{T}_t is defined as the quotient $E_t / C_p \mathbf{m}_t$.
- The continuous variable $\mathbf{T}_t^{(i)}$ is defined as the quotient $E_t / C_p \mathbf{m}_t$ if $\mathbf{m}_t = \mathcal{M}_i$, and zero otherwise. These

variables could be called “partial temperatures”, as their sum yields the temperature.

Then, constraints are added to impose those semantics.

- Exactly one discrete mass value is possible at any time step:

$$\sum_{i \in \mathcal{M}} \mathbf{m}_t^{(i)} = 1, \quad \forall t \in \mathcal{T}. \tag{43}$$

- The mass is then defined as a linear combination of those binary choices:

$$\mathbf{m}_t = \sum_{i \in \mathcal{M}} \mathbf{m}_t^{(i)} \mathcal{M}_i, \quad \forall t \in \mathcal{T}. \tag{44}$$

- Similarly, the temperature is given by the sum over all partial temperatures:

$$\mathbf{T}_t = \sum_{i \in \mathcal{M}} \mathbf{T}_t^{(i)}, \quad \forall t \in \mathcal{T}. \tag{45}$$

- The partial temperatures $\mathbf{T}_t^{(i)}$ are linked to the binary choices by the following upper and lower bounds:

$$\mathbf{T}_t^{(i)} \geq \frac{E_t}{C_p \mathcal{M}_i} - \frac{E_{\max}}{C_p m_{\min}} (1 - \mathbf{m}_t^{(i)}), \quad \forall t \in \mathcal{T}, \quad \forall i \in \mathcal{M}, \tag{46}$$

$$\mathbf{T}_t^{(i)} \leq \frac{E_t}{C_p \mathcal{M}_i}, \quad \forall t \in \mathcal{T}, \quad \forall i \in \mathcal{M}, \tag{47}$$

$$\mathbf{T}_t^{(i)} \leq \frac{E_{\max}}{C_p m_{\min}} \mathbf{m}_t^{(i)}, \quad \forall t \in \mathcal{T}, \quad \forall i \in \mathcal{M}. \tag{48}$$

The two last bounds are required to be specified separately, to ensure that a partial temperature is forced to be zero if the corresponding mass level is chosen, and takes its expected value otherwise.

Thanks to this technique, the temperature definition becomes linear. It could be extended to more general expressions.

This formulation can be strengthened in order to improve solving times. The bounds on the partial temperatures $\mathbf{T}_t^{(i)}$ should be as tight as possible to keep good solving times; this issue motivated the choice of temperature $\mathbf{T}_t = E_t / C_p \mathbf{m}_t$ over the raw quotient of variables E_t / \mathbf{m}_t , as the factor $1 / C_p$ can reduce the values that are considered by several orders of magnitude.

- The bounds on the partial temperatures also have lower bounds, using the same binary variables:

Table 1 Solving time for the various nonlinear models

Model	Underlying solver	Time (s)	Solution cost (€)	Gap reported by the solver (%)	Cost improvement with respect to reference (%)
Reference: reach target temperature as soon as possible	CPLEX 12.7.1	0.02	10,793.25	0.00	–
Reservoir model: linearisation by discretisation (16 levels)	CPLEX 12.7.1	3610.37	8208.23	2.68	23.95
	Gurobi 7.5.0	3600.02	8034.72	2.12	25.56
Reservoir model: nonlinear (continuous)	Bonmin 1.8.4	3600.59	7259.46	–	32.74
	Couenne 0.5.4	3601.39	7260.20	100	32.73

All solvers were stopped after approximately 1 h of computations, even when they did not find the optimum solution (nonzero gap). Those tests were run on a machine with two Intel Xeon E5-2650v4 (2.2 GHz) and 128 GB of RAM

Table 2 Typology of industrial processes to help reservoir modelling

Process	Mass	Inputs	Outputs
HVAC	Constant	(Constant mass: no inputs nor outputs)	
Kiln	Variable	Continuous	Continuous
Oven (including electric arc furnace, induction furnace)	Batch	Discrete (batches: when starting)	Discrete (batches: when done, e.g. molten metal)
Electrolysis	Variable	Continuous	Discrete (periodical: e.g. every day)

$$\mathbf{T}_t^{(i)} \geq \frac{E_{\min}}{C_p m_{\max}} \mathbf{m}_t^{(i)}, \quad \forall t \in \mathcal{T}, \quad \forall i \in \mathcal{M} \quad (49)$$

- The mutual exclusion of the $m_t^{(i)}$ can be further imposed with clique constraints:

$$\sum_{i \in \mathcal{V}} \mathbf{m}_t^{(i)} \leq 1, \quad \forall t \in \mathcal{T}, \quad \forall \mathcal{V} \subseteq \mathcal{M} : |\mathcal{V}| \geq 2 \quad (50)$$

4.2 Benchmark

The two nonlinear models are benchmarked against each other, in order to compare their performance when solving the same problem. We use an oven-like model to heat a given quantity of metal that may be added at any rate and time (which gives a nonlinear model). Fifty time steps are considered. Heating happens with temperature ramping constraints. All models have been written using JuMP [11] in Julia [6].

Those models are all compared to a base line, which heats the material as soon as possible (while respecting the same ramping constraints) and keeps it at the right temperature until the end of the horizon.

The results are shown in Table 1. When a high precision is needed in the discretised variables, both CPLEX [17] and Gurobi [14], two state-of-the-art mixed-integer linear solvers, have troubles to reach a very low gap, albeit their solutions are more than satisfying for industrial applications. In the same time budget, the nonlinear formulation with

open-source nonlinear solvers achieves a better solution, without being hindered by discretised variables.

5 Process typology

Based on these models, we can derive a typology of industrial processes based on a few characteristics. Table 2 does so according to two criteria, focusing on the *flows of material to process*:

- the *inputs* to the process: are they continuous or periodical?
- the *outputs* from the process: are they continuous or periodical?

Those two questions can be equivalently formulated as follows. Is the mass currently processed constant? How can it vary (continuously, periodically)?

Combined, these three parameters indicate how the process could be modelled. A mass that is constant (HVAC) or highly predictable (electrolysis, batches; kilns to a lesser degree) can result in model simplifications. Completely variable masses often imply nonlinearity (as most ovens, see Sect. 4).

The diagrams shown in Sect. 3 also help build another typology, this time based on the *potential flexibility* levers for each process. Indeed, the presence of a decaying heat energy reservoir indicates that the process is amenable to load shifting: even if heating is stopped, the process might

Table 3 Analysis of flexibility levers available in industrial processes

Process	Potential flexibility levers
HVAC	Thermal inertia: <i>load shifting</i>
Kiln	Thermal inertia: <i>load shifting</i> (production delayed)
	Starting time: <i>load scheduling</i> <i>Fuel switching</i>
Continuous oven	Thermal inertia: <i>load shifting</i> (production delayed)
	<i>Fuel switching</i>
EAF	Thermal inertia: <i>load shifting</i> (production delayed)
	Starting time: <i>load scheduling</i>
Electrolysis	Thermal inertia: <i>load shifting</i>
	For a long period of time, becomes <i>load shedding!</i>

still continue to produce, albeit probably at a lessened rate. If multiple fuels are possible for heating, then fuel switching can be used. For all noncontinuous processes, load scheduling can be applied; for continuous processes whose production can be tuned (like electrolysis), exploiting flexibility may lead to load shedding. The processes that are studied in Sect. 3 are included in Table 3.

6 Flexibility potential of industrial processes

Some models developed in Sect. 3 are now fit to industrial data, and the flexibility potential of the processes is estimated based on historical scenarios. A major hypothesis is that exploiting the flexibility for these industrial sites has no impact on the electricity market, as they do not consume enough electricity.

All optimisation models have been written using JuMP [11] in Julia [6]. The source code for the simulations is available online at the following address:

<https://github.com/dourouc05/IndustrialProcessFlexibilisation.jl>

6.1 Induction furnace

The induction furnace of Sect. 3.2.1 needs three parameters. The first one, α , characterises the energy level depending on the initial temperature of the material. The second one, β , is the efficiency of converting the electrical power into heat (see Appendix 1). The last one, δ , indicates the energy losses.

Each sample $s \in \mathcal{S}$, taken from historical measurements, corresponds to one use of the furnace. It mainly

consists in a power curve P_s^t , indicating the average power injected through the circuits each hour (a melting cycle lasts 12 h). To have a better fit, each sample may have its own value of β , denoted by β_s . The values of the β_s are brought closer together by a constraint limiting the variance of the β_s to 0.001. All in all, fitting the parameters is done through the following optimisation program (a convex QCQP), based on the reservoir model of Sect. 3.2.1:

$$\begin{aligned}
 & \min \sum_{s \in \mathcal{S}} (E_s^t - E_s)^2 \\
 & \text{s.t. } E_s^0 = \alpha \times m_s \times T_0 \quad \forall s \in \mathcal{S} \\
 & \quad E_s^t = E_s^{t-1} + \beta_s \times P_s^t - \delta \quad \forall s \in \mathcal{S}, \forall t \in \mathcal{T} \\
 & \quad |\mathcal{S}| \times \bar{\beta} = \sum_{s \in \mathcal{S}} \beta_s \\
 & \quad \sum_{s \in \mathcal{S}} (\beta_s - \bar{\beta})^2 \leq 0.001 \\
 & \quad \alpha \geq 0 \\
 & \quad \beta_s \geq 0 \quad \forall s \in \mathcal{S} \\
 & \quad \bar{\beta} \geq 0 \\
 & \quad \delta \geq 0
 \end{aligned} \tag{51}$$


To evaluate the results of this model, we use a leave-one-out procedure [12]: for each sample $s \in \mathcal{S}$, we solve the previous optimisation program over the samples $\mathcal{S} \setminus \{s\}$ to get a value $\beta_{\mathcal{S} \setminus \{s\}}$, and we average the obtained square errors to predict the energy of s based on $\beta_{\mathcal{S} \setminus \{s\}}$. It results in a root-mean-square error on our data set of 65.194 kWh (the ratio of this error to the average energy is 1.78%).

Once a reservoir model is fit, it can be used to estimate the flexibility potential of the process. To this end, we compare our reservoir-based optimisation to current fixed consumption profiles on a historical price scenario (an average day of January 2016 on the Belgian day-ahead market). No optimisation is currently performed on the schedule: the smelting process always starts at the same hour and lasts for 12 h.

Related constraints must be added to the formulation in order to implement real-world constraints. Mostly, the peak power is limited, and the heating cannot abruptly change. Three constraints are thus added for each time step: a minimum and a maximum power (P_{\min} and P_{\max} , respectively), and a ramping constraint (ρ_{\min} and ρ_{\max} are, respectively, the minimum and maximum ratios by which the electrical consumption is allowed to change from hour to hour).

The complete model is therefore the following, where p_h indicates the price of electricity for the hour h (in €/MWh), c_h the electrical consumption (in MWh) and E_h the energy of the metal to melt:

Fig. 6 Comparison of consumption profiles between the existing fixed solution and the result of the optimisation



$$\begin{aligned}
 & \min \sum_{h=1}^H p_h c_h \\
 & \text{s.t. } E_0 = \alpha m T_0, \\
 & E_H = \alpha m T_H, \\
 & E_h = E_{h-1} + \beta c_h - \delta \quad \forall h \in \{1, 2 \dots H-1\} \\
 & P_{\min} \leq c_h \leq P_{\max} \quad \forall h \in \{1, 2 \dots H\} \\
 & \rho_{\min} c_h \leq c_{h+1} \leq \rho_{\max} c_h \quad \forall h \in \{1, 2 \dots H\}
 \end{aligned} \quad (52)$$

Per heat, using a reservoir model to decide the heating power, hour per hour, could decrease the costs by 8.35% per heat, from €239.11 to €219.15; over a year, this corresponds to more than €15,000 of savings. The main difference in the planned power consumption is that its peak is shifted to exploit the lowest prices during the production period (as shown in Fig. 6).

Computationally speaking, this model allows to optimise the required power to a given price scenario in a fraction of a second: a 95% confidence interval is 0.03 ± 0.01 s (with either CPLEX 12.7.1 or Gurobi 7.5.0). Fitting the parameters is also very quick, as the leave-one-out validation phase takes less than 2 min for the available data. The from-scratch implementation in Julia takes 150 lines, but the process could be automatised with a GUI to build the reservoir model. These characteristics are very appealing for energy-sector consultants, who may want to estimate the flexibility potential of an industrial site in very little time: most of the effort can be spent on actual discussions on the flexibility solutions that can be implemented.

6.2 Industrial cooling

Based on industrial data from a polypropylene-film plant, a model similar to that of Sect. 3.1.1 can be tuned. All the needed physical constants are known and match the data set. The model is then fed with scenarios of water volumes to cool down. In practice, the price scenarios can be estimated in advance, while the heat and water volumes are known with a high precision.

The processes in Fig. 2 have to be specified in more details. A linear model is deemed sufficient for our needs, using coefficients of performance, and the functions f and g can be written as:

$$f(\mathbf{P}_{CT,t}) = \text{cop}_{CT} \times \mathbf{P}_{CT,t} \quad (53)$$

$$f(\mathbf{P}_{\text{chiller},t}) = \text{cop}_{\text{chiller}} \times \mathbf{P}_{\text{chiller},t} \quad (54)$$

Thus, two energy budget constraints (14) and (15) become:

$$\begin{aligned}
 \mathbf{E}_{CT,t+1} &= \mathbf{E}_{CT,t} + C_{\text{water}} \times T_{p,t} \times \mathbf{water}_{p \rightarrow CT,t} \\
 &\quad - \mathbf{E}_{CT \rightarrow B,t} - \text{cop}_{CT} \times \mathbf{P}_{CT,t} \\
 \mathbf{E}_{B,t+1} &= \mathbf{E}_{B,t} + C_{\text{water}} \times T_{p,t} \times \mathbf{water}_{p \rightarrow B,t} + \mathbf{E}_{CT \rightarrow B,t} \\
 &\quad - C_{\text{water}} \times T_{B,t} \times \mathbf{water}_{B \rightarrow p,t} - \text{cop}_{\text{chiller}} \times \mathbf{P}_{\text{chiller},t}
 \end{aligned}$$

Two other constraints must be added for the cooling processes, as they have a limited maximum power:

$$\mathbf{P}_{CT,t} \leq P_{CT,\max} \quad (55)$$

$$\mathbf{P}_{\text{chiller},t} \leq P_{\text{chiller},\max} \quad (56)$$

Remark 6 Variable coefficients of performance are not considered in this case in order to keep the model simple. The increased complexity is not justified in a context of *approximate* models.

In order to compare the impact of flexibility on the cooling behaviour, the temperature bounds are set in two different ways:

- either $T_{\text{target}} \pm 0.5^\circ\text{C}$ (*low-flexibility scenario*), as currently implemented in the studied use case
- or $T_{\text{target}} \pm 3^\circ\text{C}$ (*high-flexibility scenario*), the maximum temperature variation that the production equipment may tolerate, based on a deeper analysis of their data sheets

Table 4 Comparison between the two formulations, for a time horizon of 8 h, with a time step of 1 h

	Mixed-integer linear	Nonconvex	
	CPLEX 12.7.1	Couenne 0.5.4	Gurobi 9.0.0
Number of constraints	1120	100	
Number of variables	480 (378 integers)	102 (0 integers)	
Average number of explored nodes	1,193,050,714	2,210,153	662,159
Average solving time (s)	62,742.96	3694.52	46.95

25 discretisation steps are used for the temperature. The programs are solved within a rolling-horizon algorithm; averages are based on its iteration. The mixed-integer linear model could not be solved over the 1-week horizon due to memory problems (the formulation is too large). Those tests were run with CPLEX 12.7.1 (MILP formulation), Couenne 0.5.4 and Gurobi 9.0.0 (nonconvex formulation), on a machine with two Intel Xeon E5-2650 (2.2 GHz) and 128GB of RAM

The model includes three nonlinear constraints: (17), (18), and (19). They are implemented as nonlinear constraints; we also compare this formulation to a finely discretised linear version. In practice, the nonconvex formulation can be solved to optimality faster than the discretised version with a time horizon of 8 h (Table 4).

In order to compare these two flexibility scenarios, a rolling-horizon algorithm is implemented. This choice helps keep the running times low: the optimisation program is run for 8 h; then, the result for the first time step is used as the initial condition for the next program, whose horizon is shifted by one time step. Even though it is closer to the plant operating conditions and therefore better estimates the actual flexibility potential, it no more guarantees a global optimality over the complete time horizon.

When using this methodology on a synthetic heat inflow to cool down¹, the energy costs can be lowered by 24% when going from the low- to the high-flexibility scenario (Fig. 7): it goes from €123,569 down to €93,909 for 1 week (with a price scenario corresponding to the first week of January 2016 on the Belgian day-ahead market). Results on historical data are highly similar. What is more, the obtained solution uses the extra flexibility to lower the temperature before an increase in both the electricity price and the heat to eliminate: this is exactly the expected kind of solution.

Similarly to the induction furnace (Sect. 6.1), computation times are very encouraging: a complex, real-sized, nonconvex model can be solved quickly to optimality (albeit only using Gurobi 9.0's new nonconvex functionalities). Model-building times are again reduced to a very low amount: each model corresponds to 100 lines of Julia when implemented from scratch, with very little thought required to build the reservoir model, even for someone who is not a specialist of these processes.

¹ No historical scenario could be shown in article due to the data being proprietary.

7 Conclusion

The reservoir framework helps build models of actual industrial processes. It does so by following a physically based approach, around the concept of storage (either material or energy). This principle corresponds to the industrial and physical reality, as these models can correctly approximate the behaviour of processes. Such reservoir models can lend themselves to interpretation about flexibility (as performed in Sect. 5).

This framework defines a series of blocks that can be used to think about processes. However, it does not specify anything about the actual energy consumption. For now, they are supposed to be easily represented by so-called processes. They may hide different kinds of mathematical expressions, from a constant efficiency that translates fuel consumption into energy (which is probably the best fit for this framework, due to its simplicity) to more complex models.

Simple extensions to the proposed building blocks allow modelling more processes. For example, a crusher can correspond to a reservoir that transforms one product (such as rock) into another one (like gravel), with a conversion rate that depends on the energy consumption. The major problem of this kind of formulation is that linearity is forgone. More work is needed to look into these kinds of generalisations, and the way to produce linear models from slightly different building blocks.

Nevertheless, without needing many alterations, the reservoirs can already be used to study the impact of flexibility on some industrial processes, as done in Sect. 6. More generally, reservoir-based models share many similarities and are still able to deal with many different processes, while being lightweight.

Acknowledgements We would like to thank Mr. Nicolas Descoumont (N-SIDE at the time of the collaboration, now with OMP), Dr. Olivier Devolder (N-SIDE), and Pr. Quentin Louveaux (university of Liège) for their help in building and assessing the accuracy of our models.

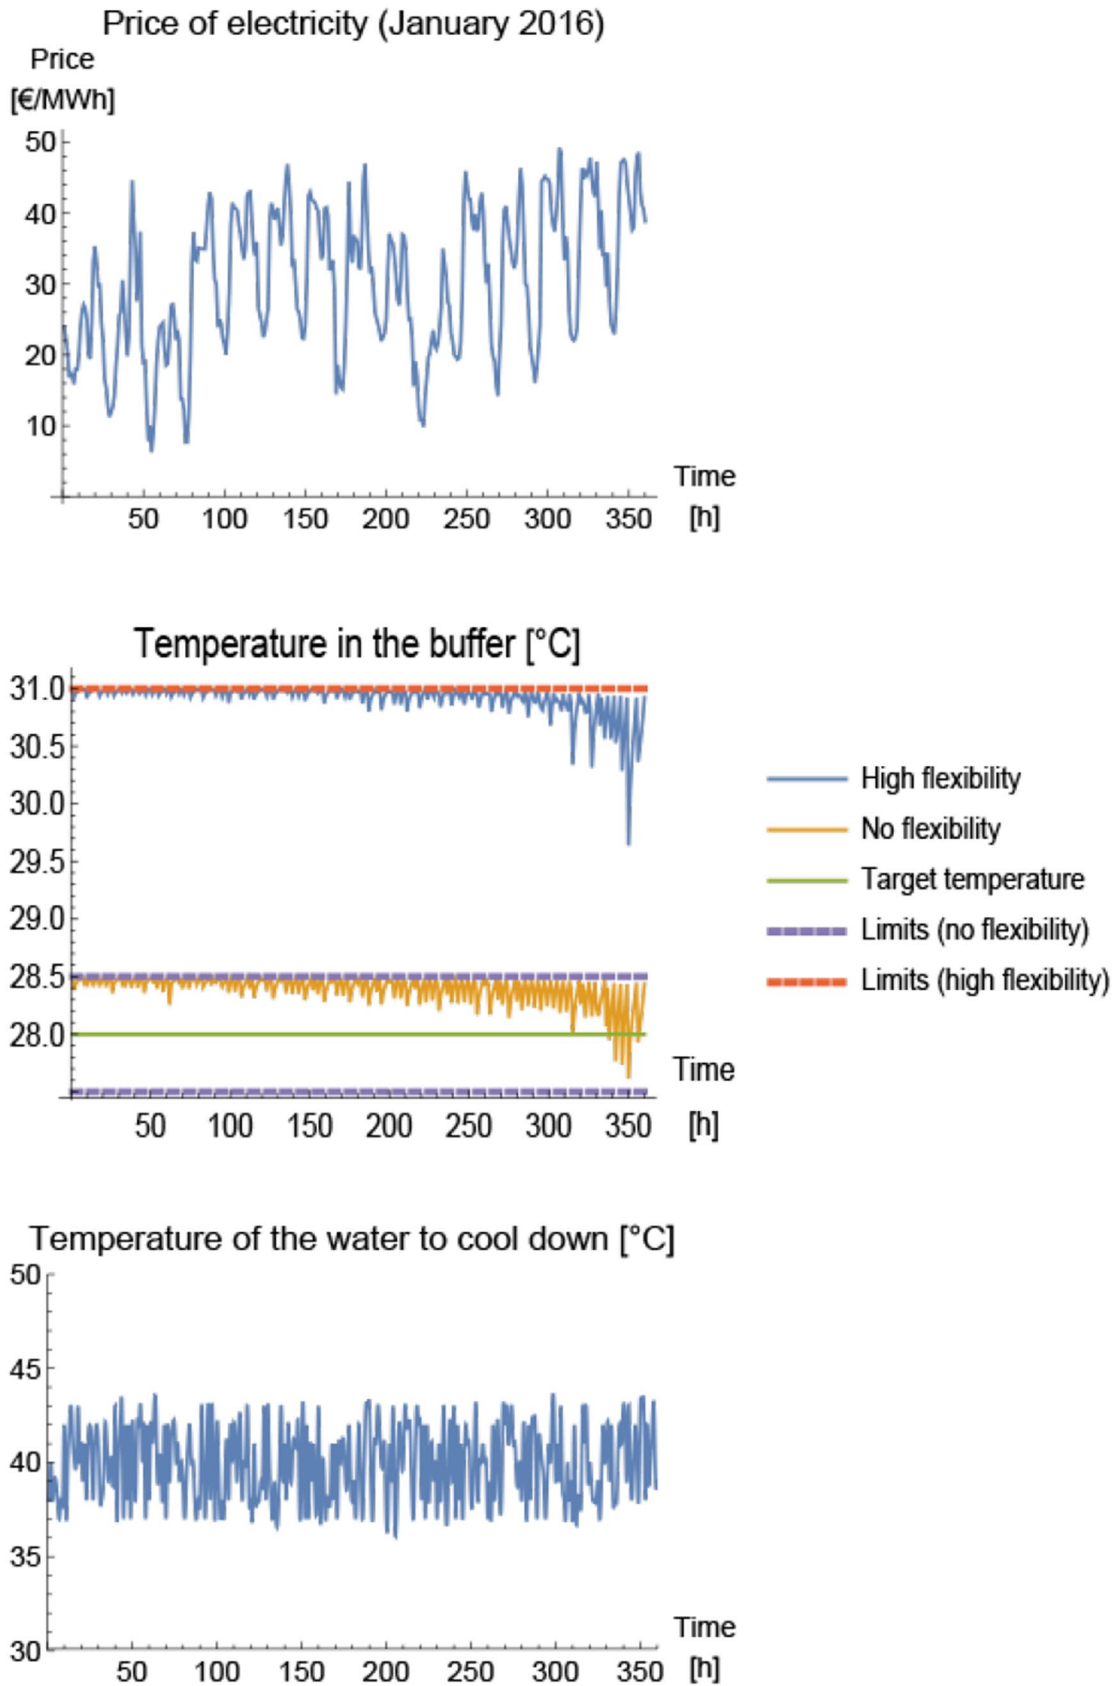


Fig. 7 Obtained experimental result for one scenario: increasing flexibility lowers the energy costs by 24%

Funding This research was supported by the public service of Wallonia, within the framework of the InduStore project (Grant 1450300). This work was carried out, while the author was a research engineer at the University of Liège, Belgium.

Compliance with ethical standards

Conflict of interest The author declares that they have no conflict of interest.

Appendix: Mathematical derivation of the induction furnace's electrical efficiency

The power dissipated by eddy currents can be expressed as a linear function of the electrical current in the coil. Indeed, the dissipation is related to the maximum magnetic field B_{\max} (if it varies sinusoidally in time) [28]:

$$P_{\text{eddy}} = \frac{\pi^2}{6} \frac{e^2 B_{\max}^2 f^2}{\rho}, \quad (57)$$

where e is thickness, f current frequency, ρ material resistivity. In an induction furnace, f is fixed, ρ depends on material, the other parameters are a function of the furnace's geometry. The only variable is thus the maximum magnetic field, which can be given by Ampère's law, as the coil corresponds to a solenoid with N wire turns and a total wire length of ℓ :

$$B = \mu \frac{N}{\ell} I. \quad (58)$$

μ is the permeability of the metal to heat and the bucket, and I is the intensity of the electrical current. This formula can be rewritten to exhibit the electrical power P in the coil instead of the current I , using Ohm's law $P = RI^2$:

$$B = \mu \frac{N}{\ell} \frac{\sqrt{P}}{R}. \quad (59)$$

Finally, the dissipated power in the metal is linked to the injected electrical power by:

$$\begin{aligned} P_{\text{eddy}} &= \frac{\pi^2}{6} \frac{e^2 f^2}{\rho} \left(\mu \frac{N}{\ell} \frac{\sqrt{P}}{R} \right)^2 \\ &= \frac{\pi^2}{6} \frac{e^2 f^2}{\rho} \mu^2 \frac{N^2}{\ell^2} \frac{P}{R^2} \\ &= \underbrace{\frac{\pi^2 e^2 f^2 \mu^2 N^2}{6 \rho \ell^2 R^2}}_{\text{constant}} P. \end{aligned}$$

In other words, the power dissipated in the metal is directly proportional to the electrical power fed into the circuit, with a constant ratio.

References

1. Abubakre OK, Muriana RA et al (2009) Mathematical model for optimizing charge and heel levels in steel remelting induction furnace for foundry shop. *J Miner Mater Charact Eng* 8:425
2. Alcoa Inc (2004) Aluminium smelting. https://www.alcoa.com/global/en/about%7B_%7Dalcoa/pdf/Smeltingpaper.pdf. Accessed June 2020
3. Asadinejad A, Tomsovic K (2017) Optimal use of incentive and price based demand response to reduce costs and price volatility. *Electr Power Syst Res* 144:215–223
4. BALCO (2013) Aluminium production technology. <http://www.balcoindia.com/operation/pdf/aluminium-production-process.pdf>. Accessed June 2020
5. Belotti P (2009) Couenne: a user's manual. <https://projects.coin-or.org/Couenne/export/914/html/couenne-user-manual.pdf>. Accessed June 2020
6. Bezanson J, Edelman A, Karpinski S, Shah VB (2017) Julia: a fresh approach to numerical computing. *SIAM Rev* 59(1):65–98
7. Bonami P, Lee J (2013) BONMIN users' manual. Technical report. <https://www.coin-or.org/Bonmin/>. Accessed June 2020
8. Bosoaga A, Masek O, Oakey JE (2009) CO₂ capture technologies for cement industry. *Energy Procedia* 1(1):133–140
9. Crombie M (2006) Calculating heat loss. <http://www.process-heating.com/articles/87988-calculating-heat-loss>. Accessed June 2020
10. Dejemeppe C, Devolder O, Lecomte V, Schaus P (2016) Forward-checking filtering for nested cardinality constraints: application to an energy cost-aware production planning problem for tissue manufacturing. In: Quimper CG (ed) *Integration of AI and OR techniques in constraint programming*. Springer, Berlin, pp 108–124
11. Dunning I, Huchette J, Lubin M (2017) JuMP: a modeling language for mathematical optimization. *SIAM Rev* 59(2):295–320
12. Efron B (1982) The jackknife, the bootstrap and other resampling plans. SIAM, Philadelphia
13. Geng Z, Zhang Y, Li C, Han Y, Cui Y, Yu B (2020) Energy optimization and prediction modeling of petrochemical industries: an improved convolutional neural network based on cross-feature. *Energy* 194:116851
14. Gurobi Optimization LLC (2020) Gurobi optimizer reference manual. <http://www.gurobi.com>. Accessed June 2020
15. Han Y, Zhou R, Geng Z, Bai J, Ma B, Fan J (2020) A novel data envelopment analysis cross-model integrating interpretative structural model and analytic hierarchy process for energy efficiency evaluation and optimization modeling: Application to ethylene industries. *J Clean Prod* 246:118965
16. Hofer T (2011) Numerical simulation and optimization of the alumina distribution in an aluminium electrolysis pot. Technical report, EPFL
17. IBM (2020) IBM ILOG CPLEX 12.10 user's manual. <https://www.ibm.com/analytics/cplex-optimizer>. Accessed June 2020
18. Kusiak A, Li M, Tang F (2010) Modeling and optimization of HVAC energy consumption. *Appl Energy* 87(10):3092–3102
19. Lee YM, Horesh R, Liberti L (2015) Simulation and optimization of energy efficient operation of HVAC system as demand response with distributed energy resources. In: 2015 Winter simulation conference (WSC). IEEE, pp 991–999

20. Liberti L (2009) Reformulations in mathematical programming: definitions and systematics. *RAIRO Oper Res* 43(1):55–85
21. Liberti L, Maculan N (2009) Reformulation techniques in mathematical programming. *Discrete Appl Math* 157(6):1165–1166
22. Molina-Garcia A, Kessler M, Bueso MC, Fuentes JA, Gomez-Lazaro E, Faura F (2011) Modeling aluminum smelter plants using sliced inverse regression with a view towards load flexibility. *IEEE Transact Power Syst* 26(1):282–293. <https://doi.org/10.1109/TPWRS.2010.2051566>
23. Naar R, Bay F (2013) Numerical optimisation for induction heat treatment processes. *Appl Math Model* 37(4):2074–2085
24. Nagarajan H, Lu M, Yamangil E, Bent R (2016) Tightening McCormick relaxations for nonlinear programs via dynamic multivariate partitioning. In: Rueher M (ed) *Principles and practice of constraint programming*. Springer, Cham, pp 369–387
25. Peesel RH, Schlosser F, Schaumburg C, Meschede H (2017) Prädiktive simulationsgestützte Optimierung von Kältemaschinen im Verbund. *Simul Prod Logist*. http://www.asim-fachtagung-spl.de/asim2017/papers/Proof_108_Peesel.pdf
26. Peesel RH, Schlosser F, Meschede H, Dunkelberg H, Walmsley TG (2019) Optimization of cooling utility system with continuous self-learning performance models. *Energies* 12(10):1926
27. Petracci NC, Hoch PM, Eliceche AM (1996) Flexibility analysis of an ethylene plant. *Comput Chem Eng* 20:S443–S448
28. Serway RA, Jewett JJ (2010) *Physics for scientists and engineers*, 9th edn. Cengage Learning, Boston
29. Short M, Rodriguez S, Charlesworth R, Crosbie T, Dawood N (2019) Optimal dispatch of aggregated HVAC units for demand response: an industry 4.0 approach. *Energies* 12(22):4320
30. Söderman J, Ahtila P (2010) Optimisation model for integration of cooling and heating systems in large industrial plants. *Appl Therm Eng* 30(1):15–22
31. The MathWorks I (2020) Simulation and model-based design. <https://www.mathworks.com/products/simulink.html>. Accessed June 2020
32. The New Zealand Institute of Chemistry (2008) The production of aluminium. <http://nzic.org.nz/ChemProcesses/metals/8D.pdf>. Accessed June 2020
33. Todd D, Caufield M, Helms B, Starke M, Kirby B, Kueck J (2008) Providing reliability services through demand response: a preliminary evaluation of the demand response capabilities of Alcoa Inc. Technical report
34. Vielma JP (2015) Mixed integer linear programming formulation techniques. *SIAM Rev* 57(1):3–57
35. Zhao QW, Liu CL, Sun Z, Yu JG (2020) Analysing and optimizing the electrolysis efficiency of a lithium cell based on the electrochemical and multiphase model. *R Soc Open Sci* 7(1):191124

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.