



HAL
open science

Economic optimisation of cold production: a matheuristic with artificial neural network approach

Alnour Ribault, Samuel Vercraene Vercraene, Sébastien Henry, Yacine Ouzrout

► To cite this version:

Alnour Ribault, Samuel Vercraene Vercraene, Sébastien Henry, Yacine Ouzrout. Economic optimisation of cold production: a matheuristic with artificial neural network approach. *International Journal of Production Research*, 2021, 59 (22), pp.6941-6962. 10.1080/00207543.2020.1831705 . hal-03003595

HAL Id: hal-03003595

<https://hal.science/hal-03003595v1>

Submitted on 13 Nov 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.

Economic optimisation of cold production: a matheuristic with artificial neural network approach

Alnour Ribault^{a,b}, Samuel Vercaene^{b*}, Sébastien Henry^b, Yacine Ouzrout^b

^aEnergy Pool, Le Bourget-du-Lac, France; ^bDISP laboratory, University of Lyon, France

ARTICLE HISTORY

Compiled October 14, 2019

ABSTRACT

In this paper, the economic optimisation of cold stores with a dynamic electricity price is studied. The main cost in cold storage is energy, and the combination of three factors make the optimisation of cold-room costs difficult to achieve. First, the temperature inside the cold rooms must be maintained within the allowed range, and complex thermodynamic processes make predicting temperatures a hard task. Second, production schedules that induce higher risks of compressor wear, and thus high maintenance costs, have to be avoided. Third, the price of electricity, which is known in advance, varies during the day. In this paper a cold store model is presented, which includes a fixed cost payment when the compressor production level changes, as well as an Artificial Neural Network model for the temperature. A Dynamic Programming solution is presented to solve simple cases and a matheuristic algorithm is designed to solve the general case. The matheuristic is compared to a linear solver, and performs better overall. Finally, managerial insights are derived from this study. The impact of storage capacity on operating costs is evaluated, and the influence of the fixed cost value on the electricity cost value of the optimised production schedules is discussed.

KEYWORDS

Cold storage; smart grid; neural network; cost optimisation; matheuristic

1. Introduction

The cold storage industry consumes considerable amounts of energy. In 2002, the International Institute of Refrigeration estimated that cold stores used between 30 and 50 kWh/m³/year, while the total cold storage volume in Europe was estimated between 60 and 70 million m³. The ICE-E (Improving Cold storage Equipment in Europe) project¹, funded by the European Union and run from 2010 to 2012 concluded that the potential energy savings in European cold stores range from 30% to 40%, going up to 72% in some cases, with one of the possible means of energy saving being improvements in the production and distribution of cold. In this section, the general contexts of electricity production and cold storage in Europe are first presented. The potential of cold storage for economic optimisation is then discussed, and, finally, the structure of this paper is presented.

*Corresponding author. Email: samuel.vercaene@insa-lyon.fr

¹Improving Cold storage Equipment in Europe, <https://ec.europa.eu/energy/intelligent/projects/en/projects/ice-e>.

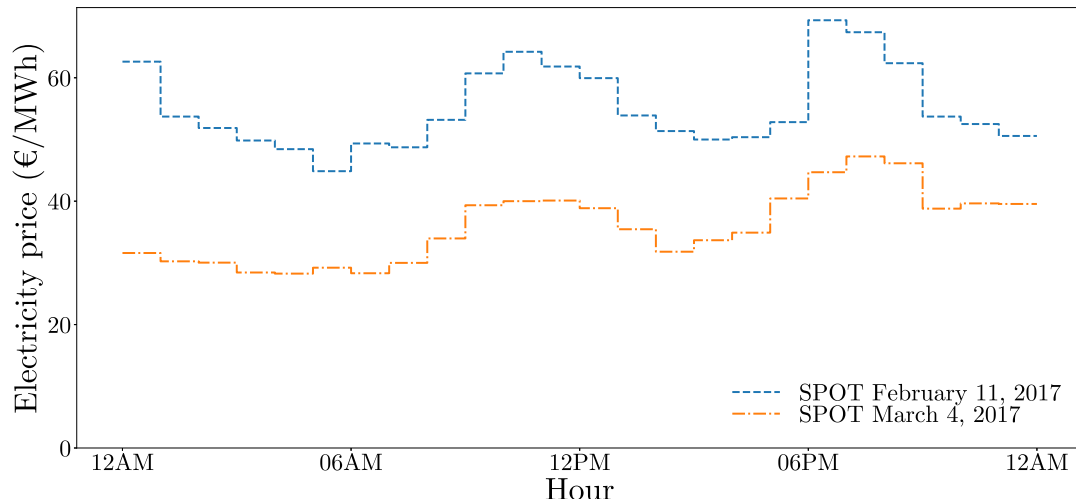


Figure 1. Instances of SPOT price over two different days.

1.1. Context of electricity production

A critical issue with regard to the synchronous electrical grid of Continental Europe is the balance between electricity supply and demand. Since production resources can be subject to operating constraints (e.g. scheduled maintenance on a power plant) or unplanned exogenous events (e.g. meteorological events for solar or wind energy), energy consumption must be controlled. This is called *demand side management*. One of the incentive mechanisms to encourage consumers to modulate their electricity consumption according to the production and demand is the modulation of the price of electricity. This mechanism is common in smart grid systems.

In Europe, an increasing number of industries are offered the possibility to pay the *SPOT* price for their electricity. The *SPOT* price is an hourly price published around 12 pm for the next day based on supply and demand forecasts by the European Power Exchange ². Therefore, the *SPOT* price is known and deterministic for at least the next 12 hours and at most the next 36 hours. This tariff can be financially advantageous in industries where production can be rescheduled easily in order to take full advantage of the low prices opportunities. Figure 1 presents two instances of *SPOT* price scenarios.

1.2. Context of cold storage

The aim of cold storage is to maintain stored goods in a given temperature range. Temperature control is paramount for cold storage since a failure might cause damage to the stored goods, and in the case of edible products, health hazards.

Cold stores are highly complex, nonlinear systems. The macro-organisation of a cold store can generally be described as follows.

First, one or several compressors, condensers, and an expansion valve convert a coolant from a gas state to a liquid state, and in so doing consume electricity.

Since the compressors are the main electricity consumer of the system, with approximately 80% of the total electricity consumption, in the rest of this paper we group the electric consumption of the condensers together with that of the compressors.

²EPEX SPOT SE, <https://www.epexspot.com>.

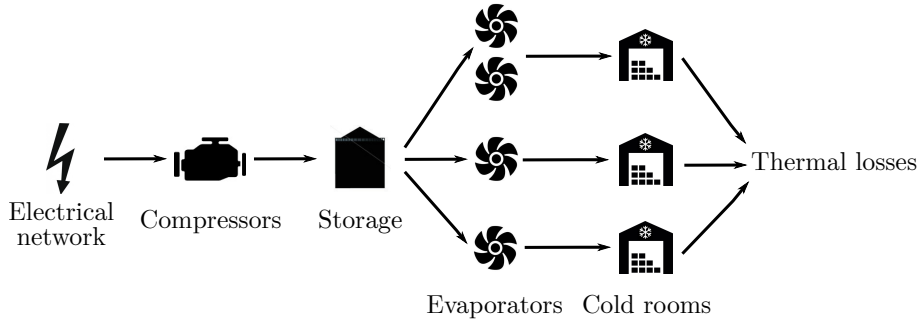


Figure 2. Diagram of a cold store.

The coolant in a liquid state is stored in a tank. Several evaporators physically present in the cold rooms pump liquid coolant from the tank and transfer the heat from the cold room to the coolant, which decreases the temperature in the cold room. Doing so converts some of the liquid back into gas. Meanwhile, several complex exchanges of energy with the exterior of the cold room, the stored goods, the staff, and so on increase the temperature in the cold room.

From an energy management point of view, this can be expressed more simply as several compressors that convert electricity to cold stored in a tank. This cold is taken from the tank and sent to the cold rooms by the evaporators, and the cold is then consumed by way of several mechanisms (see Figure 2). Aside from the efficiency of each element (compressor and evaporator), the thermal losses are functions of the gradient between the internal temperatures (tank, cold room) and the outdoor temperature which has strong variations depending on the time and the season. Since the energy loss at a given time depends on the gradient between internal and outdoor temperature, moving a production block forward or backwards in time may impact the total energy production required on the complete time horizon.

Due to thermal phenomena that generate premature wear, compressors are subject to a risk of failure when they start-up. This can lead to costly maintenance or even the replacement of the machine. The efficiency of a compressor depends on its production level. Thus, the thermal power output as a non-linear function of the electric power input. In this paper the production levels of the compressors are considered discrete (which is realistic based on professional practices) and the wear is modeled as a fixed cost when starting a compressor.

Evaporators are the second main energy consumer of the system. Generally, several on/off controlled evaporators are present in each cold room. The set of evaporators in a cold room can thus be seen as a single evaporator with several discrete production levels.

The temperature forecast in a cold room as a function of the evaporator production level based on a physical model is complex and specific to each installation. Several nonlinear exchanges of energy and nonlinear behavior of the cold production system itself make classical thermodynamic models inefficient or intractable for this context. As described by [Ruano et al., 2006] it is possible to get around this problem with Artificial Neural Networks (ANN), which allow one to model nonlinear functions of multiple variables, provided one has enough training data.

1.3. *optimisation opportunity*

The thermal inertia of the cold rooms offers a potential for economic optimisation, by taking advantage of low price opportunities to lower the internal temperature of the cold room. Given that the tanks of modern cold stores enable one to cover around an hour of maximal evaporator consumption of frost energy, the tank can also be designed and used for economic optimisation.

The temperature control of cold rooms generally involves a constant hysteresis mechanism without energy cost optimisation. Given that the cost of electricity is known in advance and that the temperature can be forecast efficiently with an ANN, temperature control can be optimised from an economic point of view.

In this paper we focus on finding the optimal control of frost energy production and distribution of frozen stores. The objective is to minimise the total operating cost, including the time-varying production cost and the maintenance cost, while respecting the store's operating constraint, i.e. maintaining the cold rooms' temperature between its bounds.

The paper is organised as follows. In Section 2, existing works on optimising temperature control problems are discussed. In Section 3, we propose a mathematical model of cold stores. In Section 4, a dynamic programming algorithm is presented to optimise a simple case of the model, and in Section 5, a matheuristic algorithm is presented which allows us to optimise the general model. In Section 6, temperature forecast methods are discussed. A numerical validation of the matheuristic is presented in Section 7. In Section 8, the impact of the storage size on operating costs, as well as the fixed cost influence on the optimised schedules, are discussed. Finally, Section 9 presents a conclusion and some perspectives.

2. Literature Review

Energy efficiency is a growing concern in industry. The factory of the future where energy consumption or gas emissions are adequately controlled has been called an "ecofactory" [May et al., 2016] and has been the subject of several optimisation studies (see e.g. Absi et al. [2016]).

In this paper we focus on minimising the functioning costs of a cold store. Section 2.1 presents some similar nonlinear lot-sizing problems found in the literature. Section 2.2 presents the more specific problem of temperature control, which has been studied in the industrial sector as well as the residential sector. Section 2.3 presents existing work using matheuristics in lot-sizing context. Then, as a conclusion of the literature review, Section 2.4 presents the contributions of this paper.

2.1. *Lot-sizing*

An analogy can be made between our problem and a special case of the single-stage, single-item Capacitated Lot-sizing and Scheduling Problem (CLSP). The frost production facility can be viewed as a single production server with a fixed setup cost. Frost losses in the cold room can be seen as demands, and the cold room itself as a capacitated storage. A review on the modeling of industrial extensions of lot sizing problems is proposed by Jans and Degraeve [2008]. Another review, on the single-item version of the dynamic lot-sizing problem and many of its extensions, can be found in Brahimi et al. [2017]. The considered problem is also reminiscent of the One-Warehouse Multi-Retailer

(OWMR) [Gayon et al., 2017], where a centralised warehouse must satisfy the demand of several retailers at each time step, while minimising its production and holding costs.

In our case the thermal losses are a function of the internal temperature of the cold room. By analogy with the CLSP, the demand is a function of the stock level. The demand is also a function of the production level and the outdoor temperature.

Examples of stock-dependent demand in inventory management can be found for example in Jolai et al. [2006] and Pal et al. [2006] in the discrete case, and in Teng and Chang [2005] in the continuous case. Panda et al. [2010] give an example of stock-dependent demand for the specific case of the OWMR. In these works, the demand is considered to be linear in the stock level, which results in an exponential evolution of the stock level.

2.2. *Optimal Temperature Control*

The optimal temperature control problem in buildings has mainly been studied in the context of buildings occupied by humans (in this respect, see [Shaikh et al., 2014] for a review).

Being able to forecast the temperature is necessary in order to control its evolution. Different types of models can be found in the literature with various complexities. Those models can be classified into two categories: physical models and data based models.

For the former, some authors use a simpler physical model called first order model [Sonderegger, 1977; Ha et al., 2006], that describes the heat exchanges between the cold room and the outside air using a first order differential equation

$$C \frac{d\theta(t)}{dt} = P_t - \frac{1}{R} (\theta(t) - \Theta_t^{ext})$$

The small number of parameters makes them easy to estimate based on historical data, by solving a constrained ordinary least-squares problem. The differential equation can then be solved and the resulting model discretised for use in a linear program. This first-order model is analog to the linear stock-dependent demand used in lot-sizing problems [Jolai et al., 2006; Pal et al., 2006; Teng and Chang, 2005; Panda et al., 2010] resulting in an exponential stock level evolution.

Alternatively, others use a complete physical model that aims at describing the whole frost production system [Hovgaard et al., 2013]. However, those models require knowledge or estimation of many physical parameters, for example the enthalpy of the coolant before and after compression, the thermal capacity of the goods stored in the cold rooms, or the thermal resistance of the walls of the cold rooms.

An other way to obtain a temperature forecast model is to use a data-based approach. Artificial neural networks (ANNs) have been shown to be efficient tools for the forecast of parameters in manufacturing process [Cook and Shannon, 1992]. In the same way Ruano et al. [2006] use ANNs and argue that they are a suitable tool to produce adequate forecasts for temperature control, whereas physical models are more costly to develop and therefore better suited to the conception phase of a building. With our historical temperature data, Figure 3 shows a comparison between the predictions from a first-order model [Sonderegger, 1977; Ha et al., 2006], and an ANN model. Clearly, the ANN provides a superior forecast in our case.

While some authors [Lu et al., 2005; Ferreira et al., 2012] have chosen to minimise the environmental impact of the system by using energy consumption as the objective




Figure 3. Discretised ANN prediction compared to real temperature and to the linear model on a given day.

function, others [Ha et al., 2006; Hovgaard et al., 2013] focus on economic optimisation by minimising the total cost of energy consumption. Other optimisation criteria have been considered as well, such as the discomfort felt by the inhabitants of the building [Ferreira et al., 2012; Ha et al., 2006] because of the temperature changes.

As far as we know, the work closest to our concern in the literature on cost optimisation, in a context of temperature control of multiple cold rooms, is Hovgaard et al. [2013].

They model a three cold-room facility with one compressor. The objective is to minimise the cost to maintain the temperature of the three cold rooms within a certain range.

All the frost energy produced by the compressor is directly transferred to the cold rooms, so there is no active energy storage (like a tank of coolant).

They use a first-order model modeling heat exchanges between the air inside the cold room, the outside air, the coolant, and the stored food. This requires estimating several physical parameters based on historical data. Hovgaard et al. [2013] add a penalty to the objective function for considering the compressor wear due to by changes in its consumption level. The considered penalty is quadratic in the difference in compressor power over subsequent time steps and also serves to improve the convergence of their non-convex solving method.

2.3. *Matheuristics*

Matheuristics are algorithms that combine elements of metaheuristics and exact solving such as mathematical programming.

As shown in Ball [2011], matheuristics have been successfully applied to a wide variety of problems. In particular, capacitated lot-sizing and scheduling problems have been addressed by Guimarães et al. [2013] and Bayley et al. [2018] by applying a column-generation based matheuristics and a Benders decomposition matheuristics respectively.

Ball [2011] provides a survey on matheuristics and proposes a classification into three classes. 1) decomposition approaches, where the problem is decomposed into subproblems which are solved to optimality; 2) improvement heuristics, where mathematical programming is used to improve an initial solution found by some heuristic method;

and 3) approaches using mathematical programming to generate approximate solutions, where a relaxation of the problem is solved to optimality. The approach proposed in our paper, which is a column-generation based matheuristic, arguably falls into both categories 2) and 3) of this classification.

2.4. Contributions with respect to the literature

In this paper the temperature control problem in a cold store with several cold rooms and active energy storage is considered. This problem is modeled as a capacitated nonlinear lot-sizing problem. The objective is to minimise both the cost of the energy consumption and the fixed cost payments at each compressor start-up. The latter allow one to account for compressor wear, a concern which was reported to us by a cold-store manager.

The work closest to our concern is Hovgaard et al. [2013]. We extend their work in several directions: 1) the possibility to use the tank as an active storage is considered; 2) several compressors may be used simultaneously; 3) an ANN temperature forecast, which is simpler to establish, is used; and 4) according to the people of the profession the fixed cost paid at each compressor start-up is more realistic than the quadratic cost modeled in Hovgaard et al. [2013].

Our problem is solved with a column-generation-based matheuristic. The subproblem of the column-generation-based matheuristic is a dynamic programming that uses an artificial neural network for the temperature forecast. The performance of the solution method is validated by comparing it with the performance of a MILP solver applied to a linear approximation of our problem. Finally, we derive managerial insights such as the impact of the storage capacity on the annual operation costs of a cold store, as well as the impact of a fixed-cost payment at each compressor start-up on the number of start-ups, on the time between consecutive start-ups, and on the electricity cost.

3. Model

3.1. Parameters

External parameters Let $\mathcal{T} = \{0, 1, \dots, T\}$ be the set of time periods. For each period $t \in \mathcal{T}$, the price of electricity is denoted C_t and the outdoor temperature is denoted Θ_t^{ext} .

Compressors Let $\hat{\mathcal{L}}$ be the set of production levels for the compressors. During one period in level $l \in \hat{\mathcal{L}}$, the frost power injected into the tank is denoted \hat{P}_l and the electric power consumed is denoted \hat{E}_l . Note that the relationship between \hat{P}_l and \hat{E}_l is not linear in our case. The initial compressors' production level (during period $t = 0$) is \hat{L}_0 . Switching the compressors from a production level l to a production level l' leads to a maintenance cost $M_{l,l'}$. This switch can be done at the beginning of each period, except for period $t = 0$.

Tank The maximal quantity of frost energy stored in the Tank is denoted \bar{S} . Let S_0 be the initial storage level of the tank. The quantity in the tank cannot be negative.

Evaporators Let $\mathcal{K} = \{1, \dots, K\}$ be the set of cold rooms and \mathcal{L}_k be the set of production levels for evaporators in cold room $k \in \mathcal{K}$. During one period at level $l \in$

\mathcal{L}_k , the frost power injected in cold room k is denoted $P_{k,l}$ and the electric power consumption is denoted $E_{k,l}$. Let \mathcal{P}_k be the set of frost powers that can be injected in the room k such that $\mathcal{P}_k = \{P_{k,l}\}_{l \in \mathcal{L}_k}$. The initial evaporators' production level (during period $t = 0$) for cold room k is $L_{k,0}$. The production level can be changed at the beginning of each period, except for period $t = 0$.

Cold rooms For each cold room $k \in \mathcal{K}$ let \mathcal{E}_k be the set of possible temperatures for room k . Given an evaporator production level $l \in \mathcal{L}_k$, an outdoor temperature Θ_t^{ext} and a current temperature $\theta_{k,t}$ at the beginning of period $t \in \mathcal{T}$, the temperature forecast function F_k yields the temperature in period $t + 1$ in cold room $k \in \mathcal{K}$.

$$\theta_{k,t+1} = F_k(P_{k,l}, \Theta_t^{ext}, \theta_{k,t}).$$

For cold room $k \in \mathcal{K}$, let $\bar{\Theta}_k$ and $\underline{\Theta}_k$ be the maximal and minimal allowed temperature inside the room, with $\bar{\Theta}_k, \underline{\Theta}_k \in \mathcal{E}_k$. The initial temperature for cold room k is $\Theta_{k,0}$.

3.2. Mathematical program

To simplify the formulation we define $\mathcal{T}^{-T} = \mathcal{T} \setminus \{T\}$.

Decision variables

- $s_t \in [0, \bar{S}]$ is the frost energy stored at the beginning of period $t \in \mathcal{T}$.
- $\theta_{k,t} \in [\underline{\Theta}_k, \bar{\Theta}_k]$ is the temperature inside cold room $k \in \mathcal{K}$ at the beginning of period $t \in \mathcal{T}$.
- $x_{l,t} \in \{0, 1\}$ indicates if the compressors' production level is set to $l \in \hat{\mathcal{L}}$ at the beginning of period $t \in \mathcal{T}$.
- $y_{k,l,t} \in \{0, 1\}$ indicates if the evaporators' production level is set to $l \in \mathcal{L}_k$ in cold room $k \in \mathcal{K}$ at the beginning of period $t \in \mathcal{T}$.
- $z_{l,l',t} \in \{0, 1\}$ indicates if the compressor production level changes from level $l \in \hat{\mathcal{L}}$ to level $l' \in \hat{\mathcal{L}}$ at the beginning of period $t \in \mathcal{T}$.

Objective The objective is to minimise the total cost including maintenance costs, the energy consumption costs for compressors, and the energy consumption costs for evaporators.

$$\min f = \sum_{t \in \mathcal{T}} \left[\sum_{l \in \hat{\mathcal{L}}} \sum_{l' \in \hat{\mathcal{L}}} z_{l,l',t} M_{l,l'} + C_t \left(\sum_{l \in \hat{\mathcal{L}}} x_{l,t} \hat{E}_l + \sum_{k \in \mathcal{K}} \sum_{l \in \mathcal{L}_k} y_{k,l,t} E_{k,l} \right) \right] \quad (1)$$

Constraints

$$s_t + \sum_{l \in \hat{\mathcal{L}}} x_{l,t} \hat{P}_l = s_{t+1} + \sum_{k \in \mathcal{K}} \sum_{l \in \mathcal{L}_k} y_{k,l,t} P_{k,l} \quad \forall t \in \mathcal{T}^{-T} \quad (2)$$

$$\theta_{k,t+1} = F_k \left(\sum_{l \in \mathcal{L}_k} y_{k,l,t} P_{k,l}, \Theta_t^{ext}, \theta_{k,t} \right) \quad \forall k \in \mathcal{K}, t \in \mathcal{T}^{-T} \quad (3)$$

$$\sum_{l \in \hat{\mathcal{L}}} x_{l,t} = 1 \quad \forall t \in \mathcal{T} \quad (4)$$

$$\sum_{l \in \mathcal{L}_k} y_{k,l,t} = 1 \quad \forall k \in \mathcal{K}, t \in \mathcal{T} \quad (5)$$

$$x_{l,t} + x_{l',t+1} - 1 \leq z_{l,l',t+1} \quad \forall l, l' \in \hat{\mathcal{L}}, t \in \mathcal{T}^{-T} \quad (6)$$

$$x_{\hat{L}_{0,0}} = 1 \quad (7)$$

$$s_0 = S_0 \quad (8)$$

$$y_{k,L_{k,0},0} = 1 \quad \forall k \in \mathcal{K} \quad (9)$$

$$\theta_{k,0} = \Theta_{k,0} \quad \forall k \in \mathcal{K} \quad (10)$$

$$0 \leq s_t \leq \bar{S} \quad \forall t \in \mathcal{T} \quad (11)$$

$$\underline{\Theta}_k \leq \theta_{k,t} \leq \bar{\Theta}_k \quad \forall k \in \mathcal{K}, t \in \mathcal{T} \quad (12)$$

$$x_{l,t}, y_{k,l,t}, z_{l,l',t} \in \{0, 1\} \quad \forall k \in \mathcal{K}, l, l' \in \mathcal{L}, t \in \mathcal{T} \quad (13)$$

Flow constraints (2) ensure the conservation of the frost energy stored in the tank. Constraints (3) specify the temperature evolution in all rooms. Constraints (4) and (5) impose the unicity of the production level for compressors and evaporators in each cold room, respectively. Constraints (6) ensure that the fixed maintenance cost is paid when the production level of the compressors changes. Constraints (7) to (10) set the initial state. Constraints (11) and (12) limit the capacity of the tank and the temperature in the rooms, respectively. Then, Constraints (13) specify the binary decision variables.

The model, outside of constraints (3), is linear. In our case, functions F_k are implemented by nonlinear ANNs. The ANNs make a good temperature forecast, but linearising their sigmoid activation function finely enough to preserve forecasting performance induces too many binary variables for efficient use in a MILP solver.

4. Dynamic programming

The minimisation problem presented in the previous section, including an ANN, can be solved with dynamic programming.

4.1. Problem reformulation

Let us first define the stock evolution function. With s the stock level at the beginning of period t , $l \in \hat{\mathcal{L}}$ the production level of the compressors, and $l_k \in \mathcal{L}_k$ the production level of the evaporators in room k , we define $S(s, l, l_1, \dots, l_{|K|})$ as the stock level at the beginning of period $t + 1$ such that

$$S(s, l, l_1, \dots, l_K) = s + \hat{P}_l - \sum_{k \in \mathcal{K}} P_{k,l_k} \quad (14)$$

Let $\mathcal{V} = \{v_t | t \in \mathcal{T}\}$ be the set of vertices with $v_t = \langle t, l_t, s_t, \theta_{1,t}, \dots, \theta_{K,t} \rangle$, $t \in \mathcal{T}$, $l_t \in \hat{\mathcal{L}}$, $s_t \in [0, \bar{S}]$, and $\theta_{k,t} \in \mathcal{E}_k \forall k \in \mathcal{K}$. With the same notation, the set of arcs is denoted $\mathcal{A} = \{(v_t, v_{t+1}) | t \in \mathcal{T}^{-T}\}$ with $s_{t+1} = S(s_t, l_t, l_{1,t}, \dots, l_{K,t})$, $l_{k,t} \in \mathcal{L}_k$, and $\theta_{k,t+1} = F_k(P_{k,l_{k,t}}, \Theta_t^{ext}, \theta_{k,t}) \forall k \in \mathcal{K}$. Then, the weight $H_{v_t, v_{t+1}}$ on the arc $(v_t, v_{t+1}) \in \mathcal{A}$ is defined as

$$H_{v_t, v_{t+1}} = C_t \left(\hat{E}_{l_t} + \sum_{k \in \mathcal{K}} E_{k, l_{k,t}} \right) + M_{l_t, l_{t+1}} \quad (15)$$

Finding the cost-optimal production schedule is equivalent to finding the shortest path from the initial state $v_0 = \langle 0, \hat{L}_0, S_0, \Theta_{1,0}, \dots, \Theta_{K,0} \rangle$ to a final state v_T (that could be any vertex of the last period) in the graph $\{\mathcal{V}, \mathcal{A}\}$.

4.2. discretisation

The time complexity of this dynamic program is $|\mathcal{V}| \cdot \prod_{k \in \mathcal{K}} |\mathcal{P}_k|$. Because s_t and $\theta_{k,t}$ are continuous, the vertex set \mathcal{V} is infinite so it is not possible to solve this problem in the general case with the dynamic program. Therefore, the first approximation is to make \mathcal{V} finite.

With

$$\epsilon_k^\theta = \frac{\bar{\Theta}_k - \underline{\Theta}_k}{N_k^\theta}$$

the temperature step and N_k^θ the number of temperature steps, let

$$\mathcal{E}'_k = \{[\underline{\Theta}_k, \underline{\Theta}_k + \epsilon_k^\theta[, \dots, [\bar{\Theta}_k - \epsilon_k^\theta, \bar{\Theta}_k]\}$$

be the discrete temperature set for cold room $k \in \mathcal{K}$.

In the same way, with

$$\epsilon^s = \frac{\bar{S}}{N^s}$$

the storage step and N^s the number of storage steps, let

$$\mathcal{S}' = \{[0, \epsilon^s[, \dots, [(N^s - 1)\epsilon^s, \bar{S}]\}$$

be the discrete storage set.

With a naive implementation, where the temperatures and storage are rounded at each time step, the discretisation accumulated errors are $T \cdot \epsilon_k^\theta$ and $T \cdot \epsilon^s$ at the end of the time horizon. In order to reduce the discretisation error, the dynamic programming is computed forward and the vertices are labelled with variables denoting *real temperatures* (one per cold room) and *real storage* values. For a given vertex, these variables correspond to the values of temperatures and storage that would be found, without discretisation, following the best path found so far from the initial vertex v_0 to the current vertex. During the solving, for a given combination of frost powers (for compressors and evaporators), the successor of a vertex at period t is computed based on its real temperatures and real storage. First, the real temperatures and real storage at period

$t + 1$ are computed based on real temperatures and real storage, then the successor (at period $t + 1$) is found such that its range of temperatures and storage contains the real temperatures and real storage found for period $t + 1$. This procedure guarantees that for any vertex, its real temperatures and real storage values are in its range of temperatures and storage.

4.3. Complexity

The dynamic program loops over each antecedent v_t of each vertex $v_{t+1} \in \mathcal{V}$ for $t \in \mathcal{T}^{-T}$ and computes the cost $H_{v_t, v_{t+1}}$ of each corresponding arc $(v_t, v_{t+1}) \in \mathcal{A}$.

For each $t > 0$, the number of vertices for period t is $|\hat{\mathcal{L}}| \cdot |\mathcal{S}'| \cdot \prod_{k=1}^K |\mathcal{E}'_k|$. For $t = 0$, there is only one vertex, v_0 . Therefore, the total number of vertices is $|\mathcal{V}| = T \cdot |\hat{\mathcal{L}}| \cdot |\mathcal{S}'| \cdot \prod_{k=1}^K |\mathcal{E}'_k| + 1$.

Moreover, each vertex in \mathcal{V} has at most $|\hat{\mathcal{L}}| \cdot \prod_{k=1}^K |\mathcal{L}_k|$ successors. Therefore, the total number of arcs in the graph is $|\mathcal{A}| \leq |\hat{\mathcal{L}}| \cdot \prod_{k=1}^K |\mathcal{L}_k| \cdot |\mathcal{V}|$.

Since computing the cost $H_{v_t, v_{t+1}}$ requires summing the electrical consumptions of the evaporators of each cold room, the corresponding time complexity is $O(K)$.

Therefore, the time complexity of the dynamic program is

$$\begin{aligned} O(K \cdot |\mathcal{A}|) &= O(|K| \cdot |\hat{\mathcal{L}}| \cdot \prod_{k=1}^K |\mathcal{L}_k| \cdot |\mathcal{V}|) \\ &= O(K \cdot |\hat{\mathcal{L}}| \cdot \prod_{k=1}^K |\mathcal{L}_k| \cdot (T \cdot |\hat{\mathcal{L}}| \cdot |\mathcal{S}'| \cdot \prod_{k=1}^K |\mathcal{E}'_k| + 1)) \\ &= O(K \cdot T \cdot |\mathcal{S}'| \cdot |\hat{\mathcal{L}}|^2 \cdot \prod_{k=1}^K |\mathcal{L}_k| |\mathcal{E}'_k|) \end{aligned}$$

As a result, it is not realistic to solve a problem with more than one cold room (i.e. not realistic for $|\mathcal{K}| > 1$).

5. Solution method

As explained in the previous section, it is not possible to solve the general problem with a linear programming solver because the model includes an ANN in some constraints. It is not moreover possible to use the discretised dynamic programming formulation with more than one cold room because the graph would be too large.

We use a matheuristic that combines a set partitioning formulation of our problem with the dynamic programming applied cold room per cold room, and a local search.

5.1. The matheuristic framework

Algorithm 1 presents our matheuristic framework. We first initialise the pool of production profiles Ω_k for each cold room $k \in \mathcal{K}$ (line 2). The initialisation is done using

Algorithm 1: The matheuristic framework

```
1 for all cold room  $k \in \mathcal{K}$  do
2   | initialise  $\Omega_k$ , the pool of production profiles for room  $k$  ;
3 end
4  $s := \text{solveSPP}(\{\Omega_l | l \in \mathcal{K}\}, \emptyset)$ , the current solution ;
5  $\text{tabu} := \{s\}$ , the list of previous solutions ;
6  $s^* := s$ , the best known solution;
7 while Stopping criteria not met do
8   | for all cold room  $k \in \mathcal{K}$  do
9     |  $\Omega_k := \{\omega(s^*, k), \text{generateNeighborhood}(s, k)\}$  ;
10  | end
11   $s := \text{solveSPP}(\{\Omega_k | k \in \mathcal{K}\}, \text{tabu})$  ;
12   $\text{tabu} := \text{tabu} \cup \{s\}$  ;
13  if  $f(s) < f(s^*)$  then
14    |  $s^* := s$  ;
15  | end
16 end
17 Return  $s^*$ ;
```

the dynamic programming independently on each room. This procedure is presented in Section 5.2. With the initial pools $\Omega_k \forall k \in \mathcal{K}$, the Set Partitioning Problem SPP is solved (line 4). Section 5.3 presents how our problem can be reformulated to a SPP with fixed costs and how it is solved in our case. The the initialisation of the solution then finishes in line 5 where a tabu list is initialised with the first solution s .

For each iteration of the algorithm, the following procedure is followed. Every pool of production profiles Ω_k is recomposed with two production profiles. The first production profile is $\omega(s^*, k)$ the production profile of room k in the best known solution s^* . The second is a generated neighbor of the current solution s (see Section 5.4). Then, the SPP presented in Section 5.3 is solved (line 11) with a time limit \bar{t}^{SPP} . All the solutions contained in the tabu list are forbidden in the SPP to make sure of generating a new current solution at every iteration of the algorithm. The new current solution found by the SPP is added to the tabu list (line 12) and the best known solution is updated if the cost of the current solution is lower than the cost of the previous best-known solution (line 14). The iterations continue until the stopping criterion is met. In our case the stopping criterion is a time limit.

5.2. Initial pools creation

The pool of production profiles Ω_k for a given cold room $k \in \mathcal{K}$ is initialised using dynamic programming. Because the graph becomes too large when more than one cold room is considered, we simplify the problem: the other cold rooms $k' \in \mathcal{K}$ with $k' \neq k$ are neglected and the frost and electrical consumption of these cold rooms are set to 0 for each period $t \in \mathcal{T}$. As a result, the set of vertices $\mathcal{V} = \{\langle t, l_t, s_t, \theta_{k,t} \rangle | t \in \mathcal{T}\}$ is reduced.

5.3. Set partitioning problem

With $\lambda_{k,\omega} \in \{0,1\}$ for all $k \in \mathcal{K}$, where $\omega \in \Omega_k$ is a decision variable that indicates if the energy production profile ω is selected, and $E_{\omega,t}$ the evaporator electrical power consumption of ω at time step t , the problem can be reformulated as a set partitioning problem with fixed costs as follows.

$$\min f = \sum_{t \in \mathcal{T}} \left[\sum_{l \in \hat{\mathcal{L}}} \sum_{l' \in \hat{\mathcal{L}}} z_{l,l',t} M_{l,l'} + C_t \left(\sum_{l \in \hat{\mathcal{L}}} x_{l,t} \hat{E}_l + \sum_{k \in \mathcal{K}} \sum_{\omega \in \Omega_k} \lambda_{k,\omega} E_{\omega,t} \right) + s_t^{relax} C^{relax} \right] \quad (16)$$

$$s_t + \sum_{l \in \hat{\mathcal{L}}} x_{l,t} \hat{P}_l = s_{t+1} + \sum_{k \in \mathcal{K}} \sum_{\omega \in \Omega_k} \lambda_{k,\omega} P_{k,\omega,t} \quad \forall t \in \mathcal{T}^{-T} \quad (17)$$

$$\sum_{l \in \hat{\mathcal{L}}} x_{l,t} = 1 \quad \forall t \in \mathcal{T} \quad (18)$$

$$\sum_{\omega \in \Omega_k} \lambda_{k,\omega} = 1 \quad \forall k \in \mathcal{K} \quad (19)$$

$$x_{l,t} + x_{l',t+1} - 1 \leq z_{l,l',t+1} \quad \forall l, l' \in \hat{\mathcal{L}}, t \in \mathcal{T}^{-T} \quad (20)$$

$$x_{\hat{L}_{0,0}} = 1 \quad (21)$$

$$s_0 = S_0 \quad (22)$$

$$-s_t^{relax} \leq s_t \leq \bar{S} \quad \forall t \in \mathcal{T} \quad (23)$$

$$\sum_{k \in \mathcal{K} | \omega(sol,k) \in \Omega_k} \lambda_{k,\omega(sol,k)} \leq |\mathcal{K}| \quad \forall sol \in tabu \quad (24)$$

$$x_{l,t}, z_{l,l',t}, \lambda_{k,\omega} \in \{0,1\} \quad \forall k \in \mathcal{K}, l, l' \in \hat{\mathcal{L}}, t \in \mathcal{T}, \omega \in \Omega_k \quad (25)$$

$$s_t^{relax} \geq 0 \quad t \in \mathcal{T} \quad (26)$$

Note that this formulation is now linear because it does not contain the nonlinear forecast function F . To ensure the feasibility of the problem, the mathematical program in Section 3.2 is slightly modified by introducing a relaxation of the constraints $0 \leq s_t$. Hence a decision variable $s_t^{relax} \geq 0$ is defined for the relaxation, and constraints (23) replace the constraints $0 \leq s_t$. Then a large penalty C^{relax} cost per unit of relaxation is added to the objective function.

With $\omega(sol,k)$ denoting the production profile of room k in the solution sol , constraints (24) ensure that all solutions in the *tabu* set of previously found solutions cannot be selected again as an optimal solution.

In Algorithm 1, the function `solveSPP`($\{\Omega_k | k \in \mathcal{K}\}$) is called. This function solves the set partitioning problem (with the relaxation on minimal stock, and tabu constraints) with a MILP solver and yields the best solution found within a maximal execution time.

5.4. Generate neighborhood

At each iteration of Algorithm 1, the function `generateNeighborhood`(s, k) is called to generate a new production profile for the cold room $k \in \mathcal{K}$ given the current solution s .

The generation of the neighborhood uses a slightly modified version of the dynamic programming presented in Section 4. As for the initial pool creation, the dynamic pro-

gram is solved with only one cold room, $\mathcal{V} = \{ \langle t, l_t, s_t, \theta_{k,t} \rangle | t \in \mathcal{T} \}$. However, the consumption of electricity and frost power from the other cold rooms is not considered null. With

$$\bar{P}_{k,t}^s = \sum_{k' \in \mathcal{K} | k' \neq k} P_{k',l_{k'},t}(s)$$

the sum of the frost power consumption of the other rooms for the current solution s , and

$$\bar{E}_{k,t}^s = \sum_{k' \in \mathcal{K} | k' \neq k} E_{k',l_{k'},t}(s)$$

the sum of the electric power consumption of the other rooms for the current solution s , the equations (14) and (15) are modified as follows

$$S(s, l, l_1, \dots, l_K) = s + \hat{P}_l - P_{k,l_k} - \bar{P}_{k,t}^s \quad (27)$$

and

$$H_{v_t, v_{t+1}} = C_t \left(\hat{E}_{l_t} + E_{k,l_{k,t}} + \bar{E}_{k,t}^s \right) + M_{l_t, l_{t+1}}. \quad (28)$$

Note that the values $\bar{P}_{k,t}^s$ and $\bar{E}_{k,t}^s$ are considered constant when the dynamic program for room k is solved. The production profile generated when solving this dynamic program is returned by the function `generateNeighborhood(s, k)`.

6. Temperature Forecast

Several types of heat exchange exist in the system. Without being exhaustive, the air inside the cold room exchanges heat with the outside air when the doors are opened, and the walls of the cold room themselves exchange heat with the air of the cold rooms located on either side. Heat exchanges also take place between the air and the staff in the cold room, the foodstuffs in the cold room, the floor of the room that is heated to avoid slipping on it, and so on. Given all these heat exchanges and given that some of them are uncertain, the temperature forecast is complex.

Another problem is access to the data. Our data history is limited to the temperature of the air in the cold room, the temperature outside the building and the power injected into the cold room at each period.

As presented in Section 2.2, three main types of models can be found in the literature: simple thermodynamic models, complex thermodynamic models, and Artificial Neural Networks. As a simple thermodynamic model does not perform an efficient temperature forecast, and a complex thermodynamic model requires data that are harder to collect and are specific to each cold store, an Artificial Neural Network model has been chosen. An ANN model can work with generic inputs that are monitored in every cold store, such as the internal temperature of the cold rooms and the power consumed by the evaporators.

6.1. Artificial Neural Networks

As Artificial Neural Networks have been proved effective to model complex thermodynamic systems in the context of temperature control [Ruano et al., 2006]. One ANN per cold room has been trained with one year of historical data and a 10-minute time step. The ANN training and hyper-parameter tuning follow a standard cross-validation procedure: the days of data are randomly divided into three sets of equal size: a training set, a validation set, and a testing set. This ensures that the variety of data is similar in each data set: hot days, week days, holidays, and so on. The ANNs with different sets of hyperparameter values are trained on the training set, and their performances are compared on the validation set. Once the best model has been selected, the test set can be used to provide an unbiased estimation of its performance.

The resulting ANNs have two layers : one hidden layer containing between 1 and 10 neurons with a sigmoid activation function and one output layer with a single neuron and a linear activation function. To output a temperature prediction at time $t + 1$, the ANN uses as inputs the temperature value at time t , the frost power injected into the cold room at time t and the outdoor temperature at time t .

In the next subsection, a linear approximation of the ANN model is presented.

6.2. Linear Forecast approximation

Figure 4 presents an instance of linearisation for a real cold room. The figure shows the relative evolution of the temperature at the next time step given the current temperature and power level of the evaporators. The linearisation of the temperature forecast for room $k \in \mathcal{K}$ is done with a simple linear regression with the same slope for each power level a_k and a constant offset $b_{k,l}$ per power level $l \in \mathcal{L}_k$. Since this linearisation depends on the value of the outdoor temperature Θ_t^{ext} which varies at each time period in our model, it is necessary to compute T linearisations, one for each time period. These computations can be ran offline. In the mathematical program, the linearisation coefficients depend on the time period: a_k and $b_{k,l}$ become $a_{k,t}$ and $b_{k,l,t}$.




Figure 4. Linear approximation (solid lines) of an ANN forecast with $\Theta_t^{ext} = 15^\circ\text{C}$.

As shown in Figure 5, this new approximated temperature forecast model is not as efficient as the ANN forecast. However it is the simplest feasible linearisation and it




Figure 5. An ANN linear approximation prediction compared to real temperature and ANN prediction on a given day. Temperature limits are denoted by solid lines.

allows us to replace the nonlinear constraints (3) in the model presented in Section 3.2 by the constraints

$$\theta_{k,t+1} = a_{k,t}\theta_{k,t} + \sum_{l \in \mathcal{K}_l} b_{k,l,t}y_{k,l,t} \quad \forall k \in \mathcal{K}, t \in \mathcal{T}^{-T}$$

and then use a MILP solver.

This linearisation will be useful in the next section to validate the solving method.

7. Numerical validation of the matheuristic

In this section we present some numerical results which allow us to validate the matheuristic proposed in Section 5. Sixteen instances are designed to model a wide range of industrial-sized cold stores. First, the interest of the two main mechanisms of the algorithm, besides the initialisation, is demonstrated. Then, the matheuristic is compared on these instances to a MILP solver within a time limit, and is shown to outperform it most of the time.

The program used for the comparison was coded in Java 1.8.0_45 and runs on a CPU Intel(R) Core(TM) i5-4690 CPU @ 3.50GHz. The Set partitioning Problem was solved with CPLEX 12.8. The same version of CPLEX was used as the MILP solver.

7.1. Description of instances

Modelling a cold store requires two kinds of information: characteristics of the physical elements of the frost production and distribution system (evaporators, compressors, storage, etc.) and historical data from the cold rooms in order to train the temperature model.

During this work, we had access to the characteristics and historical data of one real, industrial cold storage warehouse located in the north of France, with three cold rooms. In order to test the matheuristic algorithm on a set of diverse instances, 16 instances

were generated by extrapolation from the available data.

These instances use technical characteristics of real frost production systems. However, since the data of only three cold rooms was available, the three models trained on historical data were tweaked to obtain a higher variety of cold-room models. This was done by applying a multiplicative coefficient to the thermal losses of one of the three original models, which produced a new model.

The resulting models were used to form the 16 instances which model industrial-size cold stores with one to ten cold rooms.

The instances are named following a nomenclature based on the number of cold rooms it contains. The first letter of the instance name "S", "M", "L" and "H" indicates whether the cold storage represented by the instance is small sized, medium sized, large, or huge.

Along with those instances, real electric price data from the European Power Exchange is used. The price used is the SPOT Price, which varies hourly. The SPOT price is different every day, and its value is highly dependent on the hour of the day, as well as the season. Sixteen days of SPOT price were selected semi-randomly from historical data between 2012 and 2017, so as to span a variety of months of the year.

For all instances, the fixed-cost matrix $M_{l,l'}$ is defined as follows: $M_{l,l'} = mM$, where m is the number of compressor start-ups induced by the transition from state l to state l' . Significantly, when the fact of transitioning from state l to state l' induces no compressor start-up, then $M_{l,l'} = 0$. The value used for M is 7 € as a rough estimation of the cost of one start-up, based on the cost of a compressor and the average number of start-ups in a compressor's lifespan.

Table 1 presents some characteristics of the instances used as well as the SPOT day selected for each instance.

Table 1. Instances and SPOT price scenarios.

Instance	SPOT day	$ \mathcal{K} $	$ \mathcal{L} $	$\sum_k \mathcal{L}_k $	\bar{S} (kWh)
S1	01/07/2014	1	2	3	17
S2	28/10/2015	1	3	3	25
S3	15/07/2012	1	6	4	25
S4	30/04/2013	1	4	5	17
S5	27/04/2016	1	7	5	17
M1	12/03/2015	3	4	11	50
M2	25/01/2016	3	7	12	50
M3	15/07/2013	3	7	15	50
M4	30/06/2015	3	7	13	33
M5	22/04/2017	4	7	12	50
L1	19/03/2014	5	7	19	33
L2	29/10/2014	6	10	24	42
L3	14/08/2012	7	10	28	50
H1	01/01/2017	9	10	38	67
H2	07/09/2017	10	10	42	83
H3	27/02/2012	10	10	39	100

7.2. Parameters of the matheuristics

For the numerical experiments presented thereafter, the following parameter values are used in the matheuristic.

The optimisation horizon is 24 hours with a 10-minute time step, resulting in 144 time steps. The time step duration is based on industrial practices. Moreover, the historical data collected from the cold rooms use a 10-minute time step.

The number of temperature steps N_k^θ and the number of storage steps N^s for the discretisations presented in Section 4.2 are $N_k^\theta = 600$ and $N^s = 15$.

The time limit used depends on the size of the instance. Time limits are 60 seconds, 10 minutes, 30 minutes, and 2 hours for instance sizes S, M, L, and H respectively.

7.3. *Evaluation of mechanisms used in the matheuristic*

Table 2. Evaluation of mechanisms

Instance	Standard (euros)	Without local search	Without tabu
S1	60.04	0.00%	0.00%
S2	65.08	0.00%	0.00%
S3	54.83	0.00%	0.00%
S4	24.2	0.00%	0.00%
S5	19.86	0.00%	0.00%
M1	188.45	1.64%	0.83%
M2	179.47	3.17%	0.71%
M3	0.62	0.52%	0.52%
M4	106.75	3.78%	2.12%
M5	186.82	0.00%	0.00%
L1	189.31	4.48%	3.95%
L2	291.32	0.00%	0.00%
L3	300.52	0.78%	0.63%
H1	377.6	0.00%	0.00%
H2	284.7	0.37%	0.37%
H3	449.59	0.49%	0.49%
Average		0.95%	0.60%

Table 2 presents the results obtained with different configurations of the algorithm. Column "Standard" contains the value obtained with the algorithm presented in Section 5.

The following column contains the gap between the value of the first column and the value obtained after the initialisation step, without local search. The gap is always positive: since the local search is performed after initialisation, the value obtained with the local search can only be equal to or better than the initialisation value.

The last column contains the gap between the value of the first column and the value obtained after a local search but without using a tabu list. The value obtained using a tabu list is always equal to or better than the value without a tabu list.

This table demonstrates that the best configuration uses both a local search and a tabu list.

7.4. *Performance of the matheuristic: comparison with a MILP solver*

In this section the performance of the matheuristic is tested. Given that the nonlinear temperature forecast does not allow us to use a mixed integer linear programming solver directly, the benchmark must be made with a linear temperature model, to be able to compare the matheuristic with a MILP model solved with CPLEX. Therefore, in this section, both the matheuristic and the MILP solver use the simplified and linearised version of the temperature forecast presented in Section 6.2.

Table 3 presents a comparison of the results obtained by the matheuristic and a linear programming algorithm in equal time over 144 time steps. The time limits used depend on the instance and are those presented in Table ??.

Column "Math. Cost" contains the results of the optimisation with the matheuristic, while column "MILP Cost" contains the results of the optimisation with the MILP solver. The best lower bound obtained by the solver is presented in column "MILP Gap". Finally, Column "Math. vs MILP" presents the performance gap between the matheuristic and the linear solver.

Table 3. Comparing MILP Solver and Matheuristic over a 24-hour horizon

Instance	Math. Cost	MILP Cost	MILP Gap	Math. vs MILP
S1	38.06	41.04	38.04%	-7.26%
S2	38.53	42.52	40.87%	-9.38%
S3	106.04	106.91	38.39%	-0.81%
S4	40.19	50.50	49.71%	-20.42%
S5	31.34	38.76	44.74%	-19.14%
M1	260.74	293.18	26.53%	-11.06%
M2	338.32	357.53	20.55%	-5.37%
M3	137.67	143.16	27.50%	-3.83%
M4	181.93	199.36	32.63%	-8.74%
M5	230.87	284.27	37.85%	-18.78%
L1	251.58	262.08	23.36%	-4.01%
L2	430.10	467.67	24.98%	-8.03%
L3	166.61	169.92	17.97%	-1.95%
H1	642.34	698.40	22.02%	-8.03%
H2	455.66	520.70	26.23%	-12.49%
H3	676.52	770.38	22.67%	-12.18%
Average				-9.47%

The results show that the matheuristic performs significantly better than the MILP solver within the allowed time, with an average difference of 9.47% in favor of the matheuristic.

However, this table does not allow us to estimate the absolute performance of the matheuristic. The MILP gap between the value of the best integer solution and the best lower bound is still very large, with values ranging from 17.97% and 49.71%, depending on the instance. Therefore, even if the matheuristic performs significantly better than the MILP solver, there is no guarantee that the solution obtained by the matheuristic is close to optimal.

In order to compare the matheuristic and the MILP solver with a smaller MILP gap, the same comparison is performed over 72 time steps, that is, an optimisation horizon of 12 hours instead of 24 hours. While the same days of SPOT price as those presented in Table 1 are used, the starting time alternates between 12 am and 12 pm so as to cover both halves of the day equally. Table 4 presents the results of this comparison.

Table 4. Comparing MILP Solver and Matheuristic over a 12-hour horizon

Instance	Start time	Math. Cost	MILP Cost	MILP Gap	Math. vs MILP
S1	12 am	16.26	16.26	0.00%	0.00%
S2	12 pm	21.57	21.57	0.00%	0.00%
S3	12 am	37.46	36.92	0.00%	1.45%
S4	12 pm	23.51	23.51	0.00%	0.00%
S5	12 am	14.11	14.08	0.00%	0.23%
M1	12 pm	111.30	113.45	9.64%	-1.89%
M2	12 am	143.71	144.30	7.01%	-0.41%
M3	12 pm	75.07	76.36	17.08%	-1.69%
M4	12 am	68.85	68.94	16.02%	-0.12%
M5	12 pm	92.39	92.74	21.07%	-0.38%
L1	12 am	90.53	91.13	13.10%	-0.66%
L2	12 pm	215.87	226.39	19.47%	-4.65%
L3	12 am	166.61	169.92	17.97%	-1.95%
H1	12 pm	296.31	299.05	10.71%	-0.92%
H2	12 am	163.65	163.10	12.64%	0.34%
H3	12 pm	330.51	333.02	11.57%	-0.76%
Average					-0.71 %

When optimising over a time horizon of 12 hours instead of 24 hours, it can be seen that the matheuristic is still a bit better than the MILP solver on average, but the performance gap between both resolution methods is largely reduced. This time, the MILP solver is able to find certified optimal solutions for the smallest instances. On most of these instances, the matheuristic finds optimal (S1, S2, S4) or almost optimal (S5) solutions. However, a performance gap exists on instance S3.

This gap is due to the discretisation described in Section 4.2, which is performed to apply the dynamic programming algorithm. In this paper, parameter values $N_k^\theta = 600$ and $N^s = 15$ are used as a compromise between execution time and solution quality. When using parameter values $N_k^\theta = 1200$ and $N^s = 45$, the matheuristic is able to find an optimal solution in 10 minutes. However, we could not use those values for the comparison with the MILP solver, since the algorithm would then take more than 12 hours to complete the initialisation step on the largest instances.

Note that although the linear approximation temperature model presented in Section 6.2 is a very simple one, the matheuristic still performs better than the MILP solver overall. In addition, the matheuristic allows us to use the more efficient nonlinear ANN temperature forecast. Therefore, it is a better optimisation method for an industrial application.

8. Managerial insights

In order to present results for a realistic medium-sized industrial cold store within a tractable computation time, Instance M3 is used for the simulations presented in this section. An example of an optimised production schedule obtained with the matheuristic algorithm is presented in Figure 6.

The following data are plotted, in this order: the electrical consumption of the compressor, the electrical consumption of the evaporators, the quantity of cold energy contained in the storage, the cumulative cost, and the SPOT price over the day.

In the rest of this section, the matheuristic is used to perform analyses on 6 years of SPOT data. In order to reduce the high computation time required by the simulations, we perform a clustering of the SPOT price daily scenarios, and the medoids are used as representatives of each cluster.

Two analyses are then conducted: the first analysis focuses on the design of the tank, by considering the evolution of the operating costs depending on the size of the tank; and the second analysis focuses on the modeling of the wear of the compressors.

8.1. Price clustering

The clustering of the 2189 days of SPOT price data is computed using a Partition Around Medoids (PAM) algorithm. The number of 10 clusters is used as a good trade-off between simulation time and quality of the approximation, as a smaller number of clusters seems to simply discriminate only on the basis of the mean value of the prices, ignoring the variance. The resulting clustering is shown in Figure 7.

8.2. Storage design

The cold store model considered in the present work includes an active energy storage that acts as a buffer between the production and distribution of cold. This storage can be used in addition to the thermal inertia of the cold rooms to take advantage of the variations of the electricity price, by storing energy when the price is low and using the stored energy when the price is high.

The optimised production profiles for the medoid of each cluster are computed for different storage capacities. The cost values are presented as a function of the storage capacity for each medoid in Figure 8. The cost values of each medoid obtained are




Figure 6. Example of a production schedule.

then multiplied by the cardinal of the corresponding cluster. Summing the resulting quantities over all clusters gives an approximation of the operating cost from 2012 to 2017, for each storage capacity.

Figure 9 shows the estimated annual expenses as a function of storage capacity. Using the energy storage as a buffer for avoiding electricity consumption during moments of peak prices can allow savings of up to 50 % for the considered medium-sized cold store. In the considered example, the savings stop growing when the storage capacity goes over 125 kWh, which is significantly higher than standard storage capacities for a cold store of this size, but not unrealistic. Therefore, if such a cold store has a small storage capacity, investing in a larger storage could result in considerable savings in operating costs.

8.3. Compressor start-ups

As expressed in Section 1.2, our model includes a fixed cost payment at each change in the production level of the compressors. In these numerical experiments, we set the fixed cost matrix so that the fixed cost is paid only at compressor start-ups, which is

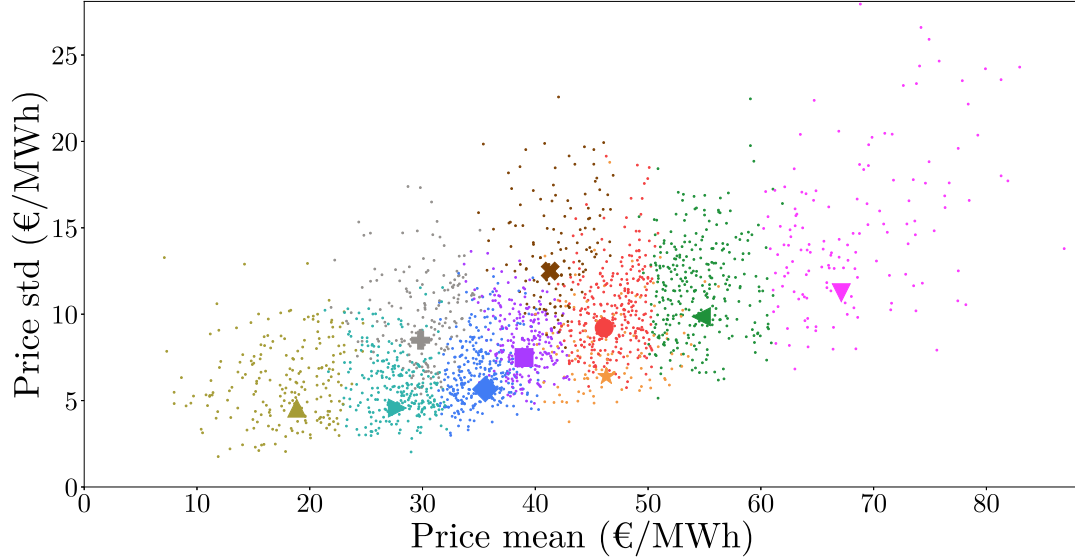


Figure 7. Clustering of SPOT prices from 2012 to 2017. The symbols denote the medoid of each cluster.

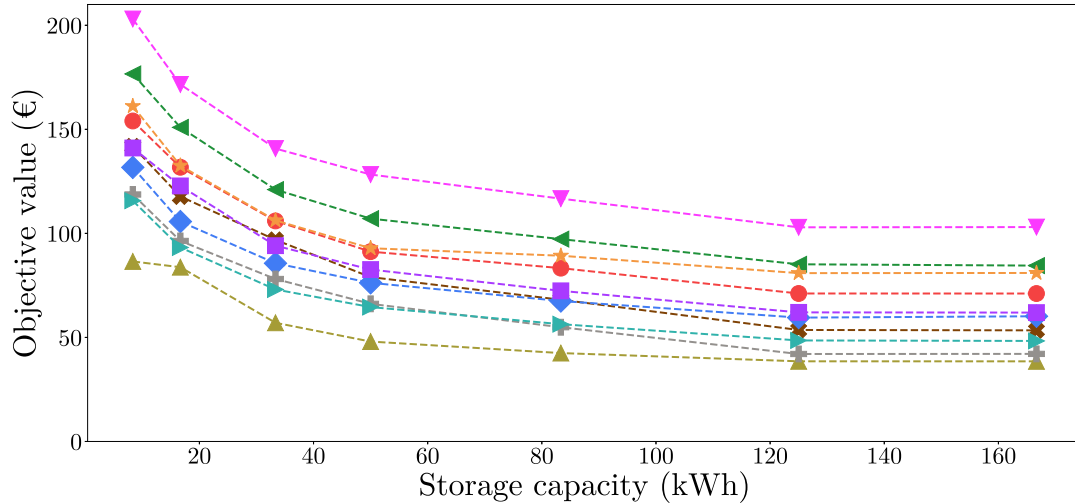


Figure 8. Evolution of optimised cost depending on storage capacity for each medoid price. The same symbols as in the previous figure are used for each medoid.

where the risk of failure is most important. In the previous numerical results, this value was fixed at €7, as the result of an estimation of the cost of a compressor start-up based on the replacement cost of a compressor, and on the average number of start-ups in a compressor lifespan. In this section, we evaluate the impact of the start-up cost value on the time between consecutive start-ups and on the electricity cost part of the cost function. As in the previous section, the optimisation is performed on each medoid price scenario, with different start-up cost values. The resulting profiles are then counted as many times as there are elements in the corresponding cluster. This produces estimations of statistics on the time between consecutive start-ups and estimations of annual electricity costs.

As seen in Figure 10, the average time between start-ups rises sharply as soon as the start-up cost is set to a non-null value, as it jumps from 3 hours with no start-up cost




Figure 9. Evolution of annual expenditures depending on storage capacity.

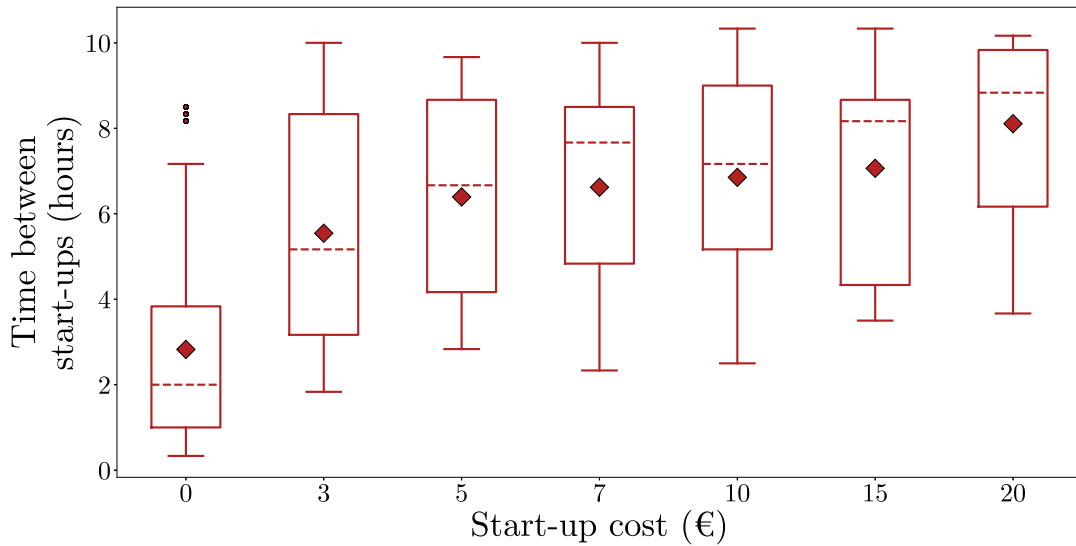


Figure 10. Time between start-ups as a function of start-up cost. Dotted lines denote the median values and symbols denote the average values.

to 5.5 hours with a €3 fixed cost. Little change can be seen however when the start-up cost ranges from €3 to €15. Then when using a €20 value, the average time between start-ups grows over 8 hours. The maximum time between start-ups jumps from 7 to 10 hours when the start-up costs become strictly positive, and barely changes when the fixed cost increases further, suggesting a limit imposed by the thermal inertia of the cold room. The minimal time also changes from 20 minutes to 2 to 4 hours when a start-up cost is used. Therefore, we can conclude that the inclusion of a start-up cost decreases the total number of start-ups and increases the time between start-ups.

Figure 11 displays the annual electricity cost as a function of the start-up cost value. As expected, the function is increasing in the start-up cost value. As the start-up cost increases, the relative weight of the electricity cost in the objective function decreases. However, the function seems to display different growth rates. It grows at a regular

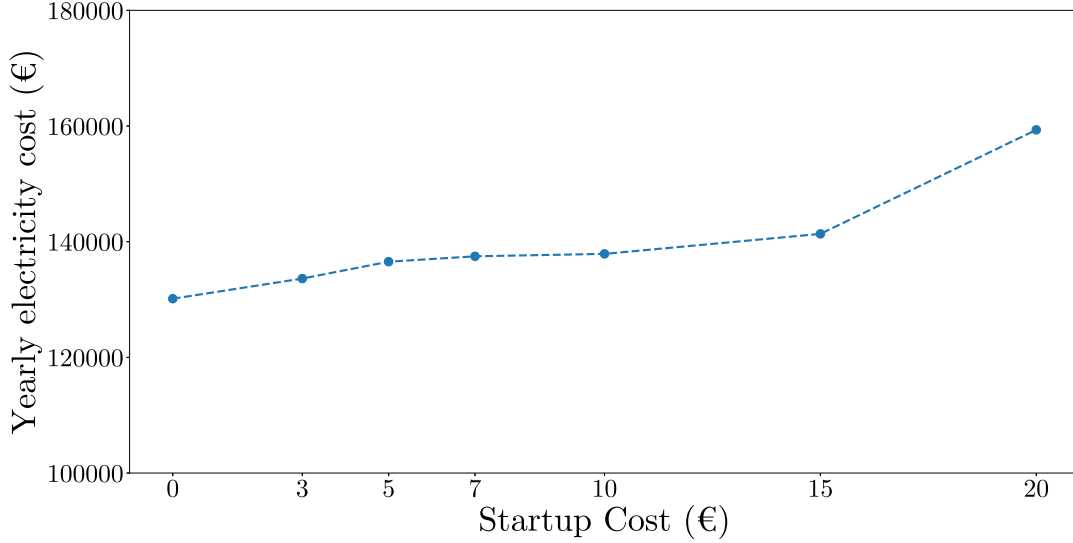


Figure 11. Annual electricity cost as a function of start-up cost value.

rate from €0 to €5, then slows down to a lower rate between €5 and €10. Finally, the growth rate increases slightly between €10 and €15 and reaches its maximal value between €15 and €20. Therefore, even though the chosen value of €7 was the result of a rough approximation, it can be deduced from the two previous graphs that any value between €5 and €10 gives similar production profiles, in terms of time between consecutive start-ups as well as in electricity cost. In relation to the annual cost with no start-up cost, using a €5 start-up cost results in a 5% increase in electricity costs, while using a €10 start-up cost results in a 6% increase.

In conclusion, we have shown that the start-up cost is an efficient way to avoid having many compressor start-ups in a short amount of time, and that the fixed cost value does not significantly influence the electricity cost.

9. Conclusion

In this paper a column generation-based matheuristic algorithm is presented for solving the temperature control problem in a cold store with several cold rooms. The objective is to minimise energy and maintenance costs by acting on both production (compressors) and distribution (evaporators) of cold energy, and using thermal energy storage as buffer. The proposed matheuristic algorithm uses dynamic programming to find production schedules for each cold room separately. It then uses mathematical programming to solve a set-partitioning problem that employs the results of dynamic programming as a pool of exploitable columns. After the initialisation phase, a local search phase is performed, again using dynamic programming and solving a set-partitioning problem to produce neighbors suitable to the current solution. A tabu list is maintained during execution to ensure solution diversity.

The matheuristic algorithm is compared to a MILP solver over instances of various industrial sizes built with data from real cold stores, using real SPOT electricity prices. The comparison shows that the matheuristic performs better than a MILP solver used with a time limit on both a 24-hour time horizon and, to a lesser extent, a 12-hour horizon. On instances where the MILP finds an optimal solution, the solution found

by the matheuristic is either optimal or within a few percentage points of optimality. Moreover, as opposed to the MILP solver, the matheuristic algorithm can also be used with a nonlinear temperature forecast. In our case, nonlinear artificial neural networks are trained on historical data and produce a better forecast than a first order linear forecast based on thermodynamic considerations.

Finally, the algorithm is used to estimate annual costs of one medium-sized instance of a cold store. It allows us to conclude that using dedicated energy storage in order to decrease electricity consumption during peak prices allows for savings of up to 50% of the running costs, including electricity costs as well as a fixed start-up cost that models the maintenance over the long term due to starting up the compressors many times. We show that the fixed start-up costs for the compressors significantly increase the average time between compressor start-up, from 3 to 6 hours, without producing a significant increase in electricity consumption costs.

The main implications of our research work concerns the evaluation of the schedules produced by our algorithm in real cold stores. While the algorithm produces cheap production schedules given our cold store model, it is yet to prove that those production schedules are indeed efficient in practice. Another implication is the use of a probabilistic temperature forecast. The temperature evolution is subject to variations due to random events, such as door openings. Moreover, the dynamic program used to create production schedules for individual cold rooms is compatible with the use of a probabilistic model. Finally, an application of this research to the optimisation of temperature control outside of cold storage, such as the residential sector, office buildings, or data centers, could be considered.

References

- N. Absi, S. Dauzère-Pérès, S. Kedad-Sidhoum, B. Penz, and C. Rapine. The single-item green lot-sizing problem with fixed carbon emissions. *European Journal of Operational Research*, 248(3):849–855, Feb. 2016.
- M. O. Ball. Heuristics based on mathematical programming. *Surveys in Operations Research and Management Science*, 16(1):21–38, Jan. 2011.
- T. Bayley, H. Süral, and J. H. Bookbinder. A hybrid Benders approach for coordinated capacitated lot-sizing of multiple product families with set-up times. *International Journal of Production Research*, 56(3):1326–1344, Feb. 2018.
- N. Brahimi, N. Absi, S. Dauzère-Pérès, and A. Nordli. Single-item dynamic lot-sizing problems: An updated survey. *European Journal of Operational Research*, 263(3): 838–863, Dec. 2017.
- D. F. Cook and R. E. Shannon. A predictive neural network modelling system for manufacturing process parameters. *International Journal of Production Research*, 30(7):1537–1550, July 1992.
- P. Ferreira, A. Ruano, S. Silva, and E. Conceição. Neural networks based predictive control for thermal comfort and energy savings in public buildings. *Energy and Buildings*, 55:238–251, Dec. 2012.
- J.-P. Gayon, G. Massonnet, C. Rapine, and G. Stauffer. Fast Approximation Algorithms for the One-Warehouse Multi-Retailer Problem Under General Cost Structures and Capacity Constraints. *Mathematics of Operations Research*, 42(3):854–875, Aug. 2017.
- L. Guimarães, D. Klabjan, and B. Almada-Lobo. Pricing, relaxing and fixing under lot sizing and scheduling. *European Journal of Operational Research*, 230(2):399–411,

- Oct. 2013.
- L. D. Ha, S. Ploix, E. Zamai, and M. Jacomino. Tabu search for the optimization of household energy consumption. In *Information Reuse and Integration, 2006 IEEE International Conference on*, pages 86–92. IEEE, 2006.
- T. G. Hovgaard, S. Boyd, L. F. Larsen, and J. B. Jørgensen. Nonconvex model predictive control for commercial refrigeration. *International Journal of Control*, 86(8):1349–1366, Aug. 2013.
- R. Jans and Z. Degraeve. Modeling industrial lot sizing problems: a review. *International Journal of Production Research*, 46(6):1619–1643, Mar. 2008.
- F. Jolai, R. Tavakkoli-Moghaddam, M. Rabbani, and M. R. Sadoughian. An economic production lot size model with deteriorating items, stock-dependent demand, inflation, and partial backlogging. *Applied Mathematics and Computation*, 181(1):380–389, Oct. 2006.
- L. Lu, W. Cai, L. Xie, S. Li, and Y. C. Soh. HVAC system optimization—in-building section. *Energy and Buildings*, 37(1):11–22, Jan. 2005.
- G. May, B. Stahl, and M. Taisch. Energy management in manufacturing: Toward eco-factories of the future – A focus group study. *Applied Energy*, 164:628–638, Feb. 2016.
- A. K. Pal, A. K. Bhunia, and R. N. Mukherjee. Optimal lot size model for deteriorating items with demand rate dependent on displayed stock level (DSL) and partial backordering. *European Journal of Operational Research*, 175(2):977–991, Dec. 2006.
- D. Panda, M. K. Maiti, and M. Maiti. Two warehouse inventory models for single vendor multiple retailers with price and stock dependent demand. *Applied Mathematical Modelling*, 34(11):3571–3585, Nov. 2010.
- A. Ruano, E. Crispim, E. Conceição, and M. Lúcio. Prediction of building’s temperature using neural networks models. *Energy and Buildings*, 38(6):682–694, June 2006.
- P. H. Shaikh, N. B. M. Nor, P. Nallagownden, I. Elamvazuthi, and T. Ibrahim. A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renewable and Sustainable Energy Reviews*, 34:409–429, June 2014.
- R. C. Sonderegger. Diagnostic tests determining the thermal response of a house. *ResearchGate*, 84(1), Nov. 1977.
- J.-T. Teng and C.-T. Chang. Economic production quantity models for deteriorating items with price- and stock-dependent demand. *Computers & Operations Research*, 32(2):297–308, Feb. 2005.