

# Modelling and Assessment of the Capabilities of a Supermarket Refrigeration System for the Provision of Regulating Power

Niamh O'Connell<sup>a,\*</sup>, Henrik Madsen<sup>a</sup>, Pierre Pinson<sup>b</sup>, Mark O'Malley<sup>c</sup>

<sup>a</sup> Department of Applied Mathematics and Computer Science, Technical University of Denmark

<sup>b</sup> Department of Electrical Engineering, Technical University of Denmark

<sup>c</sup> Department of Electrical Engineering, University College Dublin

---

## Abstract

This report presents an analysis of the demand response capabilities of a supermarket refrigeration system, with a particular focus on the suitability of this resource for participation in the regulating power market. An ARMAX model of the system is identified from experimental data, and the model is found to have time constants at 10 and 0.12 hours, indicating the potential for the system to provide flexibility in both the long- and short-term. Direct- and indirect-control architectures are employed to simulate the demand response attainable from the refrigeration system. A number of complexities are revealed that would complicate the task of devising bids on a conventional power market. These complexities are incurred due to the physical characteristics and constraints of the system as well as the particular characteristics of the control frameworks employed. Simulations considering the provision of up- and down-regulation reveal that allowing the system to occupy any state within its feasible region results in a complex behaviour. This would require intensive monitoring and control and would be excessively complicated to communicate to a market operator. By restricting the operating region of the system this behaviour can be simplified. These restrictions result in a loss of optimality, but a result in a resource that can be communicated to the market operator in the form of a bid containing a quantity of power for up- or down-regulation and the duration for which the service can be provided.

**Keywords:** Demand Response, Electricity Markets, ARMAX Modelling, MPC, Regulating Power

---

## 1 Introduction

The advent of renewable power generation as a central participant in global power systems has brought about a paradigm shift in the power sector. Power system flexibility is now a key concern, as the operating complexities incurred due to the fluctuating nature of renewable generation must be overcome. The activation of the inherent flexibility of certain electrical loads, or demand response, is a potential source of power system flexibility. As a power system resource, demand response is in its infancy. There

---

*\*Corresponding author:*

Technical University of Denmark

Department of Applied Mathematics and Computer Science

Matematiktorvet, Building 322, DK-2800 Kgs. Lyngby, Denmark

Tel: +45 4525 3369, email: noco@dtu.dk

is a great deal of research effort investigating the value that it brings to the power system, how it should be modelled and controlled, how it should participate optimally in a competitive electricity market, and the additional complexities that it might introduce to the power system (see [1, 2] and references therein).

Scale, technical ability and financial incentive are the three key components for a profitable demand response program. Residential households are frequently proposed as a potential demand response resource due primarily to the scale of the potential resource of thousands of houses [3, 4, 5, 6], however the financial return for individual households from optimised electrical heating consumption is an insufficient incentive [7]. Furthermore, activating such a resource would frequently mean replacing a prevailing heat source (e.g. gas or district heating) with electric heating, which in itself could take many years.

Supermarket refrigeration systems on the other hand fulfil all three of the requirements for an effective demand response resource. The electrical consumption of supermarkets in Denmark is approximately 550 GWh per year, which corresponds to 2% of the annual Danish electricity consumption [8], while those in Sweden have an energy intensity of 471kWh/m<sup>2</sup>/year [9], which corresponds to a consumption of 1.8TWh/year for all Swedish supermarkets, and 3% of Swedish power consumption. The electrical consumption of the average supermarket is comprised of a number of electrical sinks, including lighting and indoor heating, however refrigeration accounts for the largest share, at up to 47% [10]. As a comparison, the balancing energy required to compensate for forecast errors in wind power generation in Denmark in 2011 was 1.3 TWh [11]. The electrical consumption of Danish supermarket refrigeration systems is approximately 20% of this imbalance. This establishes that supermarkets have sufficient scale of electrical demand to effect change in the operation of a power system through the provision of demand response from their refrigeration system. Regarding the financial incentive, the cost of consumption electricity only accounts for 1% of the operating costs of a supermarket, however as the typical profit margin is only 3%, any improvement in the cost efficiency of energy consumption corresponds to a sizeable increase in revenue to the supermarket operator [9]. Finally, considering the technical feasibility of achieving demand response, the supermarket refrigeration system is ideally suited to this. The refrigeration system and the foodstuff within it are a substantial thermal mass which can be harnessed to store energy through chilling. This allows the electrical demand of the system to be adjusted or optimised towards energy or cost efficiency, as well as for the provision of conventional power system services such as regulating power. This can be achieved while maintaining the temperature of the food within acceptable limits to prevent spoiling, and the consequent loss of revenue to the supermarket operator. Demand response can be considered as a secondary revenue stream or business model for a supermarket chain, where its existing structure and the established potential for demand response provide the basis of a virtual power plant, or aggregator, that can both interface directly with the electricity market and control the response from its population of refrigeration systems. This is yet another advantage over residential demand response where a third party aggregator would be required to facilitate a connection to the market.

A limited body of work has examined the demand response potential of supermarket refrigeration systems. Non-convex model predictive control (MPC), a form of indirect-control based demand response, was employed in [8] to optimally schedule the operation of a supermarket refrigeration system with respect to the price of electricity. Direct-control based demand response from supermarkets is demonstrated briefly in [12], where an ordinary differential equation (ODE) based model of a supermarket system is simulated and successfully follows a power consumption reference. The use of refrigeration systems for the provision of frequency control from small drink chillers in local markets is proposed in [13]. Demand response applications outside of the supermarket field have attracted intensive research attention, in particular residential and building based demand response [3, 14], electric vehicles [15] and heat pumps [16].

The novel contributions of this work are three-fold:

- An ARMAX model for a supermarket refrigeration system is identified, with appropriate simplifications to ensure the key dynamics and characteristics of the system for demand response are maintained while eliminating any unnecessary model complexities. This model can be employed for simulation in studies beyond those considered in this work, and as a starting point for population studies.
- A range of demand response control architectures are developed and simulations are conducted using

the identified ARMAX model. This highlights the numerous complexities involved in modelling demand response, and the consequent barriers to achieving optimal participation in an electricity market. These control architectures can be applied to models of other potential demand response resources to explore their characteristics in a similar manner to that presented in this work.


- A simplified representation of the demand response characteristics of the supermarket refrigeration system is developed, such that participation in the electricity market can be facilitated through an aggregator which communicates these characteristics in a manner both understandable and implementable by a power system or market operator. This is an original contribution to this field, and an important step towards achieving full participation of demand response in a competitive electricity market.

This paper is structured as follows. Section 2 presents a model of the supermarket refrigeration system studied, including a description of the dataset and the modelling approach employed. This is followed by Section 3, where both direct and indirect control architectures are employed to illustrate the key characteristics and complexities of demand response from supermarket refrigeration systems. Section 4 explores how these complexities can be reduced or eliminated to result in a demand response behaviour that, while not optimal, can be communicated effectively to a system or market operator. Section 5 presents a discussion on the necessity and benefits of establishing a representation of the practical capabilities of a demand response resource (or population thereof). The conclusions to this work are presented in Section 6.

## 2 Modelling Supermarket Refrigeration System

### 2.1 Description of System and Data

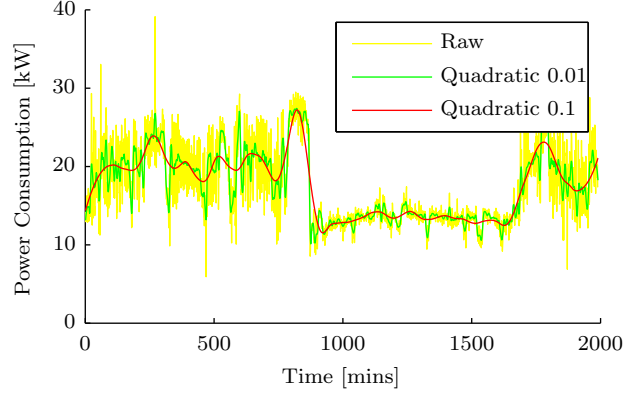
The system examined in this work is a Danfoss experimental refrigeration test centre. Data was sourced from on going experimental work at this test facility, wherein the dynamics of the system are examined by imposing step changes in the reference temperatures in the system. The test system comprises four medium temperature refrigeration display unit, and four low temperature units, two of which are included in this study. The dataset includes temperature data from a large number of sensors in each display unit, located at: the evaporator air outlet, food surface, inside the food, and at the evaporator air inlet (return). In addition, power consumption data is provided for each of the two compressors that drive the system. All data is provided at a time resolution of one minute. The data considered in this work spans eight days, from the 25<sup>th</sup> of March to the 1<sup>st</sup> of April 2013, during which period three step changes in temperature occur. A significant period of time is allowed between temperature step changes, allowing the identification of slow system dynamics. Figure 1 illustrates the placement of the sensors in each display unit. The food and food surface temperatures are denoted  $T_f$  and  $T_s$  respectively, while the evaporator inlet and outlet temperatures are given as  $T_i$  and  $T_o$ . Individual temperature control of each display unit is achieved by varying the open/closed position of the valve indicated in Figure 1. As the objective of this work is to gain insight into the overall power consumption characteristics of this system, consideration of the valve positions of individual units is deemed to be outside of the current scope.



**Figure 1:** *Simplified Display Case System Description*

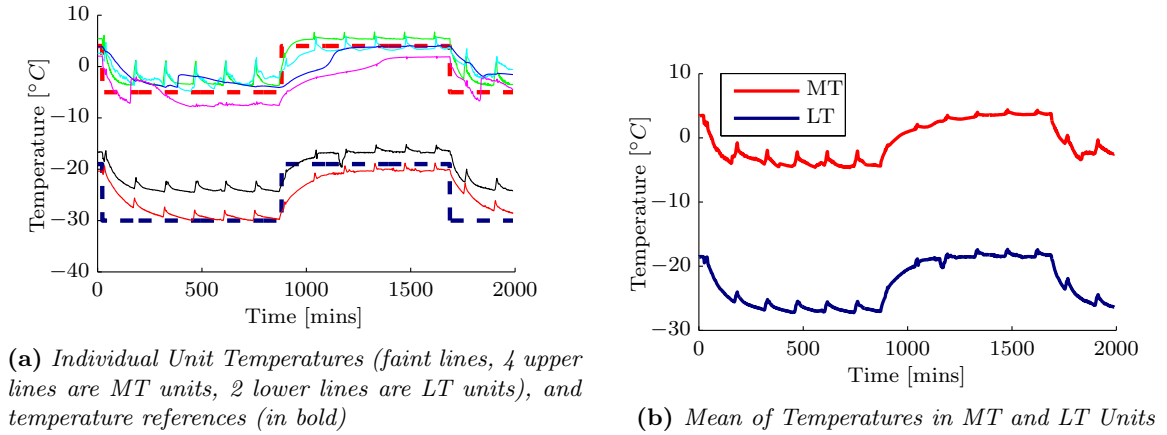
As the intention of this work is to identify the key dynamics and characteristics of this system in such a manner as to facilitate the aggregation of multiple supermarket systems, the resolution and detail of the provided dataset is considered excessive. The power consumption data is simplified by aggregating the consumption of the compressors, reducing the time resolution to five minutes, and smoothing the data using quadratic smoothing [17]. Figure 2 shows the impact of the smoothing operations. The raw power data contains high frequency elements resulting from the on/off switching operations of the compressors, which are not pertinent to the modelling of the slower dynamics relevant for demand response purposes, and which can therefore be removed through smoothing. Two smoothing spans are considered, 1% and 10% of the data; the smaller span is selected as it more accurately captures the dynamics of the power consumption changes around a temperature reference change (shown in Figure 3a).

The temperature data is similarly simplified. A single representative temperature sensor is selected for each display unit, and the six resulting temperature profiles (one for each of the 4 MT and 2 LT display units) are reduced to two representative temperature profiles, one each for the MT and LT display units. The sensor at the outlet of the evaporator is commonly considered by food safety authorities as that which



**Figure 2:** Total Power Consumption, 5 minute resolution

must remain strictly within set limits, however the relatively fast dynamics of this temperature indicate that limited demand response would be attainable with this control variable. Instead, the temperature at the food surface is selected, as this provides slower dynamics due to the thermal storage of the food, and a more realistic limit for flexibility in temperature (and consequently power consumption) while avoiding food spoilage. Figure 3a shows the temperatures in each of the display units, which follow the temperature references indicated by the bold dashed lines. Figure 3b shows the reduced temperature profiles. These reduced profiles are simply the mean of the MT and LT temperatures. Principal component analysis (PCA) was considered as a technique to determine the best representation of the aggregated MT and LT temperatures, however the presence of defrosting operations rendered these representations poor. Defrosting occurs regularly in the refrigeration system, where the temperature of the system is raised above the normal operating range to allow accumulated ice to melt and ensure the efficiency of the system. Defrosting can be observed in Figure 3a, where periodic spikes in temperature occur. Defrosting is achieved by opening the valves in the display cases fully and placing air heaters below the evaporators in each of the display units. As the power consumption of these heaters is separate from the reported power consumption of the compressors, these dynamics cannot be explained by a model that considers only the power consumption and temperature data. It is therefore advantageous to reduce the influence of defrosting operations on the overall time series of temperature. This is a further advantage of selecting the food surface temperature sensor over that located at the evaporator outlet, where the influence of defrosting exceeds that of the reference temperature change and would complicate modelling of the intended dynamics relating to the change in reference temperature.



**Figure 3:** Temperatures

## 2.2 ARMAX Model

The data provided is employed to identify an ARMAX (Auto-Regressive Moving Average with eXogenous Input), single-input, two-output model of the system. The two outputs of the system are the representative MT and LT temperatures, while the total power consumption is the considered input.

The ARMAX model has the form:

$$\phi(B)Y_t = \omega(B)X_t + \theta(B)\epsilon_t \quad (1)$$

where  $B$  represents a time lag and each of  $\phi(B)$ ,  $\omega(B)$  and  $\theta(B)$  is a polynomial, whose order is specified in the model fitting process. The MATLAB system identification toolbox [18] is used to determine the parameter values for each polynomial, using the least squares method with respect to the one-step prediction errors. Note that in this case  $Y_t$  is a vector, containing both output temperatures, and there are separate polynomials linking each output to the input power, as well as polynomials representing the dependence of the MT and LT temperatures on each other, as shown here:.

$$\begin{bmatrix} \phi_{11}(B) & \phi_{12}(B) \\ \phi_{21}(B) & \phi_{22}(B) \end{bmatrix} Y_t = \begin{bmatrix} \omega_1(B) \\ \omega_2(B) \end{bmatrix} X_t + \begin{bmatrix} \theta_1(B) \\ \theta_2(B) \end{bmatrix} \epsilon_t \quad (2)$$


A number of different model forms (number of lags in each of the  $\phi(B)$ ,  $\omega(B)$  and  $\theta(B)$  polynomials) were investigated, and evaluated based on the proximity of the resulting model residuals to Gaussian white noise; perfectly normally distributed residuals would indicate that the dynamics of the system are adequately represented in a model.

The model selected here is of the form ARMAX(2,1,2), that is, it contains a second order AR component and a first order MA component, and considers the current and directly previous input values. Note that the leading coefficient on the off-diagonal AR processes ( $\phi$  matrix) is zero, ensuring causality in the system. The off-diagonal elements in the  $\phi$  matrix represent a link in the system between the temperatures in the MT and LT units; while this link is not physically present in the system, the model indicates that the data exhibits such a link.

## 2.3 Modelling Results


The performance of this model for one-step prediction is demonstrated in Figure 4. As expected, it performs quite well, mostly due to the very limited extent of the prediction, and the fact that the model was optimised for this purpose. More interestingly, Figure 5 shows the performance of this model for simulation. In this case, the model is not provided with any updated information on the actual realisation of temperatures within the system, but estimates them based on the provided power consumption data. Assessment of the performance and accuracy of this model is important here, as an accurate model would indicate accurate simulations.

Table 1 provides some indicative metrics for this model. The mean squared error (MSE) is provided for both one-step prediction and simulation, as well as the time constants of the system. As the MT and LT temperatures are linked by the off-diagonal elements of the  $\phi$  matrix, their time constants are the same. The time constants of 10 hours and 0.12 hours are a promising indicator for the potential for the system to shift electrical demand over many hours, as well as providing instantaneous-type demand response products where the power consumption of the system is interrupted for brief periods with limited or no notice. The largest time constant is a numerical artefact, indicating possible over-fitting of the model. The calculated time constant values for this aggregated system model are supported by the findings in [8], where it is found that the system temperature with the slowest dynamics (in that case the temperature in a frost room, which is comparable to our LT temperature) increases from the minimum to maximum allowed temperatures in 11.5 hours if no cooling is applied. The system analysed in [8] cannot be directly



(a) *Model One-Step Prediction Performance, MT Unit*    (b) *Model One-Step Prediction Performance, LT Unit*

**Figure 4: Model Prediction**



(a) *Model Simulation Performance, MT Unit*

(b) *Model Simulation Performance, LT Unit*

**Figure 5: Model Simulation**

compared to that considered here as it is a separate system with distinct dynamics and constraints, however the similarity between the longer time constant in our system model and their dynamics is encouraging.

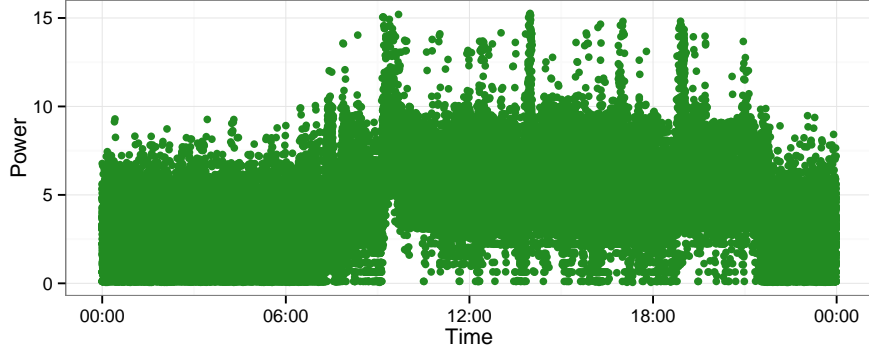
**Table 1: Model Key Metrics**

	MT	LT
MSE (Prediction) [ $^{\circ}\text{C}$ ]	0.0036	0.0021
MSE (Simulation) [ $^{\circ}\text{C}$ ]	1.1	1.44
Time Constants [hrs]	175.05	-
	10.02	-
	0.12	-
	0.05	-

## 2.4 Further Considerations


The model derived in this work is considered complete with respect to the provided data, however an accurate model of an operational supermarket would require further information, or additional inputs. There are a number of notable differences between an operational supermarket system and the Danfoss test centre considered in this work. An operational supermarket refrigeration system is subject to a number of stochastic stimuli, such as the presence of customers and staff in the shop (and their associated thermal interactions) and the removal and addition of foodstuffs. Furthermore, an operational supermarket has opening and closing hours, which will induce two separate regimes in a model. The impact of the

time of day on the power consumption in a supermarket refrigeration system is illustrated in Figure 6, where a clear change in the power consumption trend can be seen at approximately 9:00 am and 10:00 pm. In this case the data is sourced from a small supermarket located on the Danish island of Funen.



**Figure 6:** *Relationship between Power Consumption and Time of Day*

Furthermore, the limited extent of the data at the Danfoss test centre precludes the consideration of the impact of external temperature on the power consumption. Data from the operational supermarket is considered over one year from November 2011 to October 2012. Figure 7 shows the relationship between power consumption and external temperature. This is an intuitive relationship as it can be expected that a greater power consumption would be required when the external temperature is high, due to the larger temperature gradient between the external temperature and the temperature required in the refrigeration system.



**Figure 7:** *Relationship between Power Consumption and Outdoor Temperature*

The simulations presented in this work consider the somewhat limited case of the experimental refrigeration system. Thus, the simulations present an accurate indication of the general characteristics and abilities for the provision of demand response. Conducting similar experiments on an operational supermarket system over a longer time period would facilitate the identification of a more advanced model. This model could incorporate the diurnal and seasonal patterns in baseline power consumption, and the consequent changes to the volume of power and energy available for demand response.



### 3 System Controllers

The demand response capabilities of the Danfoss refrigeration test centre are explored here by placing the derived model in three different control environments: temperature reference tracking, power reference tracking and economic model predictive control. Simulations conducted with these controllers facilitate the identification of certain key characteristics of the system that should be taken into account when evaluating and analysing demand response from refrigeration system on the aggregate scale, where a population of supermarkets is considered. Furthermore they facilitate the identification of electricity market products for which this form of demand response may be suited, and the determination of methods for supermarkets to participate in such a market.

The models presented previously are stochastic models, and can be placed within stochastic control frameworks that would provide a measure of the uncertainty of the demand response resource, and consequently the reliability of the resource. For simplicity in this initial work, we restrict the consideration to deterministic simulations. Model predictive control (MPC) is considered in all cases, where the controller is provided with a forecast of the reference to follow, or the price in the case of economic MPC, and optimises the operation of the system correspondingly. Optimisations are conducted in a receding horizon manner, where the controller considers the time period from  $t$  to  $t + f$  and the control action for  $t$  is taken, the forecasts are then updated and the optimisation is repeated for  $t + 1$  to  $t + 1 + f$ , where again only the control action for  $t + 1$  is taken. This process continues for the duration of the simulation. The simulations are conducted in Matlab, with the optimisations solved in GAMS [19], using the CPLEX solver for linear problems, and the CONOPT solver for non-linear problems.

The model described in (2) is placed in a control environment with an objective function and constraints relating to the temperature limitations of the system and its power capacity. These are described in the following sections.

#### 3.1 Temperature Tracking

The objective of this controller is to have the refrigeration system track a temperature reference that is provided by an external body. This is an example of direct control demand response, where the controlling body (e.g. aggregator) has detailed knowledge of the system and can therefore issue state-based control directives directly.

The objective function in this case is the square of the difference between the achieved and the reference temperatures. Only the MT temperature is considered in the objective function here as the constraints on this temperature are the binding restriction on the flexibility of the system, as will be seen later in the simulation examples. If the LT temperature was required to follow a reference, the constraints on the MT temperature would prevent the LT temperature from ever reaching its reference in certain cases. This could be alleviated by relaxing the constraints on the MT temperature, but this is unrealistic. A proximal regularisation term (the second component of the objective function) is also included in the objective function to prevent rapid changes or oscillations in power consumption which would be damaging to the compressors, where  $\gamma$  is a constant weight to damp large steps in power consumption on each iteration. This problem is a quadratic (non-linear) problem. Equation (3c) is the system model, and defines the dynamics of the system within the control environment. The limits for the MT and LT temperatures are defined by (3d) and (3e) respectively, while (3f) defines the maximum power consumption of the system. These constraints ((3c) to (3f)) are common for all of the control environments.

The temperature limits for the MT temperature are set as  $-6^{\circ}\text{C}$  and  $6^{\circ}\text{C}$ . The limits for the LT temperature are  $-35^{\circ}\text{C}$  and  $-10^{\circ}\text{C}$ . The maximum power consumption of the system is set at 30kW. These constraint values were selected by examining the data provided from the experiments conducted at the Danfoss test centre (and can be compared with Figures 2 and 3a). It is acknowledged that these constraints, particularly the MT temperature constraints, may not be acceptable in a commercial system where food quality must be maintained. In the absence of operational data from a commercial supermar-

ket, the data from the Danfoss test system is the best approximation available.

$$\min_{\mathbf{P}} \sum_{t=1}^T \left( T_t^{MT} - T_t^{ref} \right)^2 + \gamma \sum_{t=1}^T (P_t - P_{t-1})^2 \quad (3a)$$

$$\text{s.t.} \quad \begin{bmatrix} \phi_{11}(B) & \phi_{12}(B) \\ \phi_{21}(B) & \phi_{22}(B) \end{bmatrix} \begin{bmatrix} T_t^{MT} \\ T_t^{LT} \end{bmatrix} = \begin{bmatrix} \omega_1(B) \\ \omega_2(B) \end{bmatrix} P_t + \begin{bmatrix} \theta_1(B) \\ \theta_2(B) \end{bmatrix} \epsilon_t, \quad (3b)$$

$$T_{min}^{MT} \leq T_t^{MT} \leq T_{max}^{MT}, \quad (3c)$$

$$T_{min}^{LT} \leq T_t^{LT} \leq T_{max}^{LT}, \quad (3d)$$

$$P_t \leq P_{max}, \quad (3e)$$

$$P_t \geq 0. \quad (3f)$$

Two sample simulations are presented in Figure 8. The left column shows the behaviour of the system when a step increase in temperature is imposed, while the right column shows the behaviour corresponding to a step decrease in temperature. The dashed green line shows the temperature reference to be tracked, while the dashed red lines show the temperature constraints on the system. In both cases, the LT temperature never reaches its bounds, illustrating the reason for its exclusion in the objective function of this controller.

This simulation illustrates an asymmetry in the response of the system. When a temperature increase is required from the minimum to the maximum allowed temperature, the compressors can simply shut down and allow the temperature to rise naturally. This corresponds to a reduction in power consumption of 11.57kW (this is the steady state power consumption required to maintain the minimum temperature). In comparison, when a temperature decrease is required, the compressors can increase output to their maximum level, this corresponds to an increase in power consumption of 19.87kW (a steady state power consumption of 10.13kW is required to maintain the maximum temperature). The duration for which this power change is maintained is naturally different in each case, with the power capacity of the system and the steady state power consumption required to maintain the temperature prior to the temperature change influencing the nature of the change in power consumption. The power change for the upwards temperature change is maintained for 80 minutes, while that for the downwards temperature change is only maintained for 45 minutes.


For this particular system model, the steady state power consumption required to maintain a given MT temperature is expressed as:

$$P_{ss} = 10.85 - 0.12T^{MT} \quad (4)$$


Note that this is not a general relationship, and in a realistic system this would be a more complex expression considering additional inputs such as time of day and ambient temperature.

The asymmetry illustrated here is an important consideration for any aggregator participating in demand response programs, as the imposed temperature change is simply a method to achieve a desired change in power consumption.


Furthermore, it is important that the aggregator considers the impact of the forecast extent on system behaviour. Figure 9 shows a simulation very similar to that presented in the left column of Figure 8 however in this case the forecast extent is 2 hours (previously it was 30 minutes). Comparing these two figures it can be seen that a controller with a longer forecast extent will anticipate the change in temperature reference much earlier and consequently a change in power consumption will be induced earlier than in the short forecast case. Thus, a longer forecast will extend the duration of the induced change in power consumption, however the magnitude of the power adjustment is reduced.




(a) Power Consumption,  $\Delta T > 0$




(b) Power Consumption,  $\Delta T < 0$




(c) MT Temperature,  $\Delta T > 0$



(d) MT Temperature,  $\Delta T < 0$




(e) LT Temperature,  $\Delta T > 0$




(f) LT Temperature,  $\Delta T < 0$


**Figure 8:** Temperature Reference Tracking Simulation



(a) *Power Consumption*



(b) *MT Temperature*



(c) *LT Temperature*

**Figure 9:** *Temperature Reference Tracking Simulation, with forecast of 2 hours*

### 3.2 Power Reference Tracking

The power reference tracking controller has a similar form to the temperature reference tracking controller. The key difference lies in the objective function, where the square of the difference between the achieved and reference power consumption is minimised. This is a quadratic programming problem. In this case there is no need for a penalty function to avoid oscillations as the power follows a reference, which is smooth.


The controller has the form:

$$\min_P \sum_{t=1}^T (P_t - P_t^{ref})^2 \quad (5)$$


subject to constraints given in equations (3c) to (3d).

This controller is a further example of direct control demand response, where the aggregator simply defines the power consumption to be achieved by the refrigeration system. In this case, the change in power consumption is a direct result of the control directive issued by the aggregator, whereas in the temperature reference controller it is induced as a result of the temperature reference. In a power reference controller it is important that the aggregator has sufficient understanding of the system to know if a power reference can be met, for example if the system is operating at its minimum allowed temperature no increase in power consumption can occur. In comparison, a temperature reference controller can always achieve the requested temperature reference (so long as it lies within the temperature bounds), though the resulting power consumption does not offer as fine control as a power reference tracker. For any step change in temperature reference, under a temperature tracking controller, the system will either fully curtail power consumption, or increase to its maximum power capacity, depending on the direction of the temperature change. The duration of this power change will depend on the temperature change requested, following which the power consumption will return to the level required to maintain the steady state temperature, as defined by (4).


Figure 10 shows two examples of changes in power reference; in the left column power consumption is curtailed fully and in the right it is curtailed to 5kW. The dashed green line shows the reference power to be tracked. In both cases the initial reference is 30kW, the capacity of the system, however this cannot be maintained as the minimum allowed temperature is reached. The system then reduces power consumption to 11.57kW, the required power consumption to maintain the minimum allowed temperature. This sample simulation illustrates the concept of response saturation, where a change in power reference can only be maintained for a finite period due to the physical constraints of the system. The time to saturation (duration for which a response can be maintained) is a function of the change in power reference in this particular case, where the system commences the power adjustment from steady state conditions. In other cases, the time to saturation also varies with the initial temperature and power consumption, and the forecast horizon available to the controller. Where power consumption is curtailed fully the time to saturation is 90 minutes, while that for curtailment to 5kW is 170 minutes. This is an important relationship to be considered if an aggregator wishes to participate in the regulating/balancing market as it must know how much power can be adjusted by, and for how long the adjustment can be maintained. This is further explored in Section 4.




(a) Power Consumption,  $P_{curt} = 0kW$




(b) Power Consumption,  $P_{curt} = 5kW$




(c) MT Temperature,  $P_{curt} = 0kW$



(d) MT Temperature,  $P_{curt} = 5kW$



(e) LT Temperature,  $P_{curt} = 0kW$



(f) LT Temperature,  $P_{curt} = 5kW$

**Figure 10:** Power Reference Tracking, Complete Curtailment

### 3.3 Economic MPC

Economic MPC is an example of a controller used for indirect control based demand response. In this case the aggregator simply issues prices and observes the response of the refrigeration system (or a group thereof). Typically this form of demand response control does not require in-depth knowledge of the system characteristics, instead a model can be formed and updated based on observations of power changes with respect to price. The objective function of this controller minimises the cost of consuming power, where  $\lambda_t$  is the price of electricity at time  $t$ . In addition, a function is included ( $\gamma T_T^{MT}$ ) which places a financial value on the final temperature in each of the sequential optimisations. This ensures that the optimisations do not tend towards achieving the maximum allowed temperature (minimum power consumption) to minimise the cost of power consumption. The value assigned to  $\gamma$  is essentially the value of storage of cooling energy in the system, as the resulting lower temperature in the system will provide more scope for flexibility by curtailing the power consumption of the system during high price periods. This concept is commonly employed when scheduling the operation of hydro storage power plants with associated reservoirs, where water is stored until it is economically attractive to use it to generate power [20].

The controller has the form:

$$\min_P \sum_{t=1}^T \lambda_t P_t + \gamma T_T^{MT} \quad (6)$$

subject to constraints given in equations (3c) to (3d).

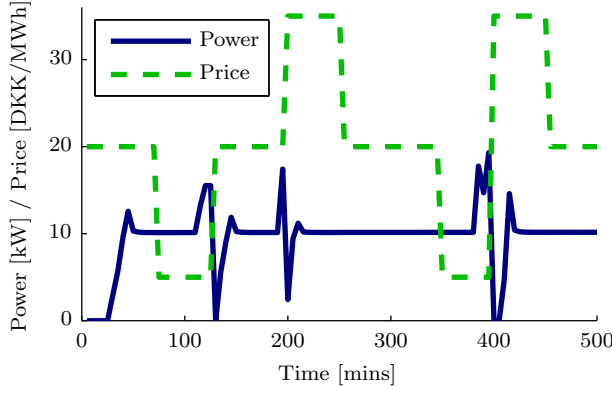
Figure 11 illustrates the performance of an economic MPC scheme, highlighting the impact of the value term in the objective function. The green dashed line shows the price profile in this case. The left column of Figure 11 shows the case where the value placed on the final temperature is zero, the right column has a value of 1000 DKK, which is arbitrarily selected for illustrative purposes. In a realistic controller this value could be tuned to reflect, for example, the day-ahead price so that the benefit of flexibility could be considered in terms of the price of consumption. It is clear that when there is a high value on the final temperature the system exhibits more flexible behaviour, as it tends towards the minimum allowed temperature in each of the sequential MPC optimisations, allowing the possibility store cooling energy during low price periods and stop power consumption during high price periods. The simulation with no value on the final temperature (left column) exhibits limited flexibility, with temperatures remaining almost constantly at the maximum allowed value. In both cases a price forecast of 30 minutes is available to the controllers, this is another factor for the very limited flexibility of the simulation with no value on the final temperature. As the controller only becomes aware of price increases 30 minutes prior to their occurrence, there is a limited time available to store cooling energy, resulting in the peaks in consumption visible just prior to each price increase. There is no corresponding decrease in power consumption price to a price decrease as the system is operating at the maximum allowable temperature and cannot reduce its consumption.

### 3.4 Key Characteristics

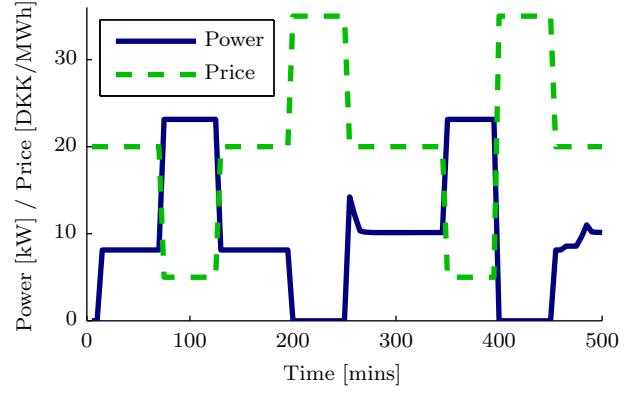
The simulations shown in the previous sections highlight a number of the complexities of demand response from refrigeration systems. These complexities can be attributed to either the physical constraints and characteristics of the system, or the control framework. Asymmetry in response, saturation of response, the impact of the forecast horizon extent and the value of the final temperature in the system (for EMPC) should all be considered by an aggregator when evaluating the potential of their portfolio of refrigeration systems to provide demand response. These are just some of the characteristics that should be considered when devising bidding strategies on a power market. Some are more difficult to consider than others, and may not be directly available to the aggregator, depending on its level of communication with the end-user (refrigeration system in this case).

In the following sections we examine the behaviour of the refrigeration system in more detail, focussing on the power reference controller and investigating the potential for participation in the regulating/balancing power market.







(a) Power Consumption




(b) Power Consumption




(c) MT Temperature



(d) MT Temperature



(e) LT Temperature



(f) LT Temperature

Figure 11: Economic MPC

## 4 Suitability of Supermarket Refrigeration Systems for Provision of Regulating Power

### 4.1 Nordic Regulating Power Market

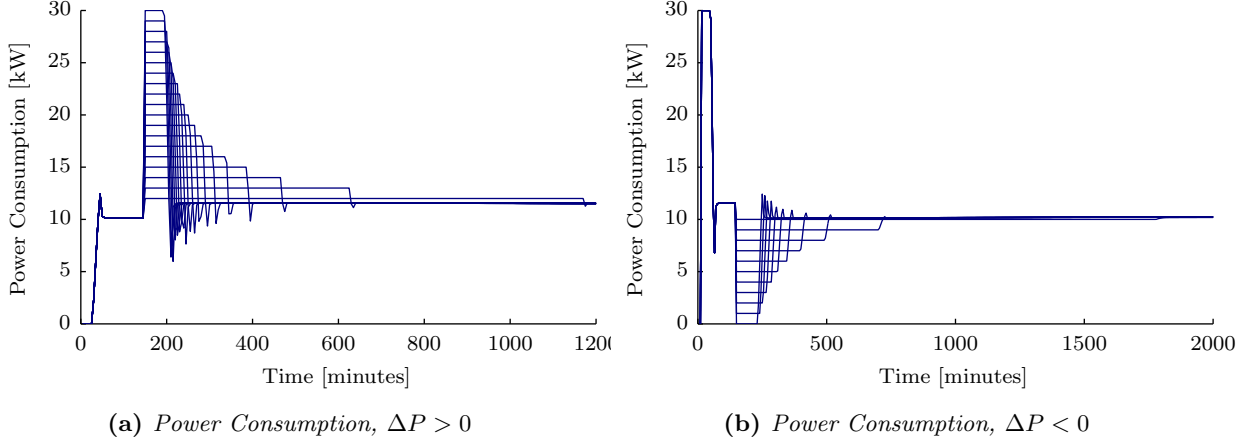
The Nordic power market consists of a number of complementary markets operating on different time scales. The regulating power market is a market for manual reserves which can be activated within 15 minutes to provide up- or down-regulation. Bids for regulating power can be submitted up to 45 minutes before the operating hour, but can be activated at any time during the operating hour. The body responsible for the provision of regulating power must be able to fully activate a given bid within a maximum of 15 minutes from receipt of the activation order. Regulation bids are activated based on a merit order bid list and the regulating power price is set by the highest activated bid for up-regulation during each operating hour (or lowest activated bid for down-regulation). The current minimum bid size on the regulating power market is 10MW, precluding individual refrigeration systems from participating [21]. Participation could be accommodated through an aggregator, which would act as a balance responsible party for a population of supermarkets. The power reference simulations shown in the previous section illustrate the suitability of the refrigeration system for provision of regulating power. The model is capable of maintaining a change in power consumption, with an associated time to saturation, or duration of response. This time to saturation will define the amount of energy that can be offered on the regulating power market as often a minimum duration of response is specified.

### 4.2 Steady-State Simulations


The basic requirement for declaration of a bid on the regulating power market is an understanding of the amount of power adjustment that can be achieved by the system under control, and for how long this can be maintained. The previous sample simulations have illustrated that this quantity (time to saturation) is dependent on power consumption prior to the adjustment, temperature prior to the adjustment, forecast extent available to the controller and the power adjustment required. Considering all of these factors results in a very complex expression. We simplify the situation by assuming that the system is in steady state, at either the maximum or minimum allowed temperature and the corresponding power consumption required to maintain this (10.13 and 11.57kW respectively). Furthermore, it is assumed that the controller has a forecast extent of 30 minutes. From an aggregator's perspective, the power reference tracking controller is the simplest solution for participation in the regulating power market as a direct command is issued to the refrigeration system and there is no need to estimate the relationship between either temperature or price and the resulting power consumption. This ensures that the requested regulating power can be provided with the greatest possible accuracy and reliability.

Here we demonstrate the relationship between a power adjustment and the time to saturation, for both up-regulation (reduction in power consumption) and down-regulation (increase in power consumption). A series of simulations are conducted, illustrating the power consumption and resulting temperature profiles for a range of steps in the power reference. Figure 12 shows the power consumption profiles in the case of down-regulation (left) and up-regulation (right), while the corresponding temperature profiles can be seen in Figure 13. In the down-regulation case, the system starts from a steady-state, at the maximum allowed temperature ( $6^{\circ}\text{C}$ ). Any downwards change in power reference would have no impact on the system as it is incapable of reducing its power consumption without breaching its temperature limits. The opposite holds for the up-regulation case, where the initial steady state condition is the minimum allowed temperature ( $-6^{\circ}\text{C}$ ). It is evident from Figure 12 that a change in power consumption occurs in all cases, though the duration for which this change can be maintained varies according to the level of adjustment in power consumption. The time to saturation is defined here as the time from the initial power change until the power adjustment can no longer be maintained and consumption deviates by more than 5% from the reference. From Figure 13 it can be seen that the systems enters a new steady-state at saturation, and in these cases the power consumption changes almost instantaneously from the reference power to the steady-state power consumption required to maintain the steady-state temperature, that is, there is no gradual ramping in power consumption. It can already be concluded from these figures that

there is a non-linear relationship between the time to saturation and the power adjustment.

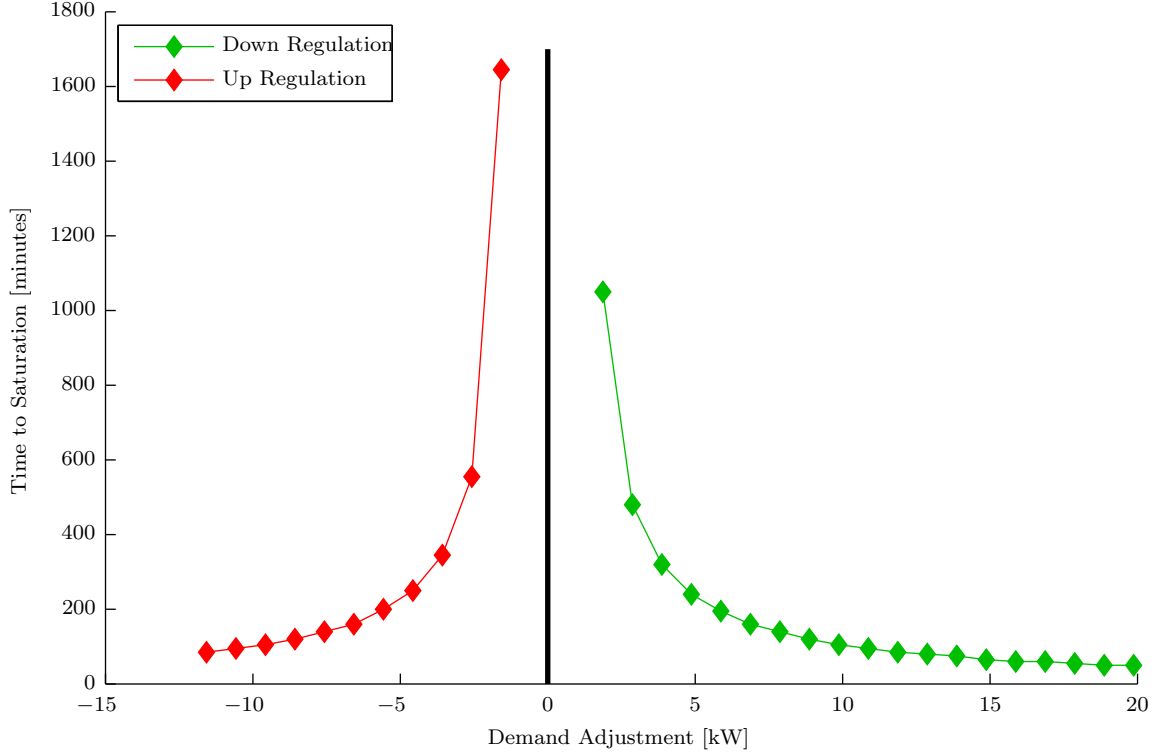


**Figure 12:** *Power Consumption Changes*



**Figure 13:** *Temperature*

Figure 14 shows the saturation time for a range of power reference adjustments, for both up- and down-regulation. The non-linear relationship can be seen very clearly here. In addition to the time to saturation, this figure also provides the recovery time of the system, another important factor to consider when participating in the regulation power market. This recovery time is the minimum period of time between the conclusion of one regulating power event and the commencement of the next, where both regulation events are in the same direction. Assume for example that a 6kW reduction in power consumption (up-regulation) is required, from a steady-state at the minimum allowed temperature. The system can maintain this for approximately 200 minutes (red curve). Following this event, and assuming that a new steady-state has been reached, the original steady state can be recovered by increasing the power consumption. If power consumption is increased by 20kW the recovery time is only 45 minutes (green curve). The time to recovery following an up-regulation event is the same as the time to saturation for down-regulation event. Alternatively, the new steady-state condition can be maintained if a down-regulation (consumption increase) event is expected. Note that these curves are symmetric between  $\pm 12\text{kW}$ , with the extended tail on down-regulation due to the large capacity of the refrigeration system. This is not unrealistic as refrigeration systems are commonly dimensioned to accommodate the highest thermal load day (i.e. warmest summer day) in 10 years, resulting in a large capacity overhead on regular operating days.




**Figure 14:** *Time to saturation of response, or duration for which a response event can be maintained, for an up-regulation event (red) and a down-regulation event (green)*


### 4.3 Impact of Forecast Horizon

In this section we extend the analysis of saturation time to consider multiple forecast extents. Power reference forecasts of 30 minutes, 2 hours, 4 hours and 6 hours are analysed. Figure 15 shows the power consumption with both up- and down-regulation for each of the forecast extents. Figure 16 shows the corresponding temperature profiles. In this case both up and down regulation are shown in a single plot for a given forecast extent. Of particular note here is that at the end of the response (when saturation is reached) power consumption does not change instantaneously from the reference adjustment to the steady-state power consumption as it did with only a 30 minutes forecast (Figure 12). With a longer forecast the change from the reference power consumption to the new steady-state is more gradual, as the controller has more advance warning of the approach to the temperature constraint and can respond appropriately. This is most clear when the power reference change is negative. Here, the system temperature will increase towards its maximum, and as this is approached power consumption will slowly ramp back upwards. Corresponding behaviour is not seen in the down-regulation case, however it is clear that in both cases the time to saturation reduces as the forecast extent increases.


Figure 17 shows the saturation time for the range of forecast extents analyses, while Figure 18 zooms into the lower section of the figure for clarity. The non-linear relationship is maintained in all cases, and in the down-regulation case there is a clear difference in saturation times for all levels of power reference adjustment. This difference increases as the power adjustment decreases. In comparison, the difference in saturation time for down-regulation converges as the power adjustment decreases. An aggregator with a portfolio of refrigeration systems would need to understand the impact of the forecast extent on the response it can expect.




(a) *Power Consumption, Forecast = 30 mins*



(b) *Power Consumption, Forecast = 2 hours*




(c) *Power Consumption, Forecast = 4 hours*




(d) *Power Consumption, Forecast = 6 hours*


**Figure 15:** *Power Consumption Changes with Varying Forecast Extent*




(a) *Temperature, Forecast = 30 mins*



(b) *Temperature, Forecast = 2 hours*




(c) *Temperature, Forecast = 4 hours*




(d) *Temperature, Forecast = 6 hours*

**Figure 16:** *Temperature Changes with Varying Forecast Extent*



**Figure 17:** *Saturation Time with Varying Forecast Extent*




**Figure 18:** *Saturation Time with Varying Forecast Extent, close-up*




#### 4.4 Complexity of Non-Steady State Operation


In this final section of simulations we extend the analysis to include non-steady state starting conditions, where the system temperature lies between the maximum and minimum values. Varying forecast extents are not considered, all simulations have a forecast of 30 minutes. Figures 19 and 20 show the power consumption and temperature profiles for the simulations, which are shown according to their starting conditions. In all cases power reference changes to induce both up- and down-regulation were issued, however in those cases where the system temperature had already reached a bound, one of the directions of power change was not possible and is therefore not seen in the figures.




(a) Power Consumption,  $P_0 = 30 \text{ kW}$




(b) Power Consumption,  $P_0 = 25 \text{ kW}$




(c) Power Consumption,  $P_0 = 20 \text{ kW}$



(d) Power Consumption,  $P_0 = 17 \text{ kW}$




(e) Power Consumption,  $P_0 = 15 \text{ kW}$




(f) Power Consumption,  $P_0 = 13 \text{ kW}$


**Figure 19:** Power Consumption Changes




(g) Power Consumption,  $P_0 = 11 \text{ kW}$



(h) Power Consumption,  $P_0 = 10 \text{ kW}$




(i) Power Consumption,  $P_0 = 9 \text{ kW}$




(j) Power Consumption,  $P_0 = 5 \text{ kW}$


**Figure 19:** Power Consumption Changes




(a) Temperature,  $P_0 = 30 \text{ kW}$




(b) Temperature,  $P_0 = 25 \text{ kW}$




(c) Temperature,  $P_0 = 20 \text{ kW}$



(d) Temperature,  $P_0 = 17 \text{ kW}$




(e) Temperature,  $P_0 = 15 \text{ kW}$




(f) Temperature,  $P_0 = 13 \text{ kW}$


**Figure 20:** Temperature Changes with Varying Forecast Extent




(g) Temperature,  $P_0 = 11 \text{ kW}$



(h) Temperature,  $P_0 = 10 \text{ kW}$



(i) Temperature,  $P_0 = 9 \text{ kW}$




(j) Temperature,  $P_0 = 5 \text{ kW}$

**Figure 20:** Temperature Changes with Varying Forecast Extent


These simulations are included to highlight the difficulties of forecasting and modelling the saturation time when the system is not starting from steady-state conditions. Figure 21 shows the saturation time for each of the simulations. Generally, for up-regulation, a lower starting temperature results in a longer saturation time, while the opposite holds for down-regulation. While the general non-linear trend observable in the previous simulations is present here, there are a number of notable exceptions. Consider in particular the simulation that commences from a starting temperature of  $-0.24^{\circ}\text{C}$  and a power consumption of 13kW; at large values of power reduction the trend in saturation time is as expected, however as the magnitude of the power reference reduction is reduced the saturation time suddenly decreases, which is not expected according to the previous trends. To understand this it is important to consider the power consumption levels required to maintain steady-state conditions at the minimum and maximum temperature limits, 11.57kW and 10.13kW respectively. Thus, a power consumption greater than 11.57kW will cause a decrease in temperature, while a consumption less than 10.13kW will cause a temperature increase. Power consumption between 10.13kW and 11.57kW will maintain a steady temperature, or a very slowly changing one. Therefore, for a starting power consumption of 13kW, a power reference decrease of greater than 2-3kW will cause the temperature to increase, whereas anything less will result in a new power consumption greater than 11kW, and thus temperature will decrease. In this case the system will saturate in two different directions, towards the maximum temperature for a power reference decrease of greater than 2-3kW and towards the minimum temperature otherwise. Figure 22 is an illustration of the temperature changes resulting from different power consumption levels.

This highlights the importance of the aggregator understanding the full nature of the refrigeration systems under its control, and the unexpected behaviour that can arise if the aggregator does not have this understanding.

These complexities can be avoided by operating the system at steady-state continuously, and reserving the spare upwards and downwards capacity for occasions where regulating power is required.



**Figure 21:** *Saturation Time, Non Steady-State Operation*



**Figure 22:** *Simple Explanation of Complexities of Saturation Time with Non Steady-State Operation*

## 5 Discussion

There is a wealth of literature examining methods to extract an optimal demand response behaviour from a large variety of applications [5, 22, 23], however the common trend is that the power system or market operator either has full knowledge of the characteristics and constraints of each responsive element [24, 25, 5, 26, 14], or they assume a simple linear price response curve [27, 28, 29]. Neither approach facilitates a realistic analysis of the participation of demand response in an electricity market or the consequent development of a business case for widespread demand response. The simulations provided in this work, in particular those employing Economic MPC, have clearly shown that the physical characteristics of a thermal energy (heating and cooling) demand response resource preclude a simple linear price response, a conclusion supported by numerous other studies [7, 30]. Furthermore, it is unrealistic to assume that the market operator has full knowledge of each of the demand response elements or applications on the system. Current electricity markets are operated at a very granular scale, where participating resources (typically generation) provide the market operator with a simple representation of its willingness to participate in the market, often as simple as a bid comprising of a price and the corresponding power level.

The role of the market operator is simply to clear the market, determining the market price. Participating resources will then respond according to whether their bid has been accepted or not. Its role is not to determine the capabilities of individual demand response resources and issue control signals to realise the required response. That is the role of an intermediary, or an aggregator, that will facilitate the participation of demand response resources in the market. An aggregator must have sufficient knowledge of its portfolio of demand response resources to understand their (dynamic) capabilities, constraints, uncertainties and the relevant control mechanisms necessary to realise their capabilities. Additionally, and crucially, the aggregator must be able to communicate this information, in understandable manner, to a market or power system operator, such that their resulting market clearing operation or optimisation (unit commitment) is numerically tractable and computationally feasible within the time allowed for such operations. This communication should also be suitable for the product the aggregator is providing, for example regulating power or flexible ramping capacity. Ideally, this should be the same, or similar to the manner in which traditional resources (generation, storage) participate in the market or power system. This simplification will naturally lead to a loss in optimality, as some of the capabilities of the demand response resource may not be representable in such a way. As an example, provision of regulating power from the examined supermarket refrigeration system at a non-steady state point results in greater flexibility, but quite complex behaviour. In comparison, simplifying the behaviour such that the system always commences a regulating power operation from steady state results in a simple non-linear curve describing the relationship between power adjustment and saturation time, a quantity that can be

communicated easily to the relevant body.

The work presented in this report focusses on how the behaviour of a single demand response resource (supermarket refrigeration system) can be simplified such that it can participate in the regulating power market by providing the market with information on the magnitude of power change it can achieve and for how long this can be maintained. How this response should be priced is a complex problem which is beyond the scope of this discussion. A similar analysis could consider the behaviour of a population of resources, and for other power markets or products. A group of supermarkets can be considered a homogeneous population, where similar model forms can be used to describe their behaviour, but each element has distinct characteristics such as power capacity, time constants and constraints. The aggregate behaviour exhibited by such a population will be different from that of a single unit, and will likely present a number of further opportunities for flexibility while still maintaining simplicity. Taking the regulating power market as an example, the saturation time in this case will be limited by the shortest saturation time within the population, and the power adjustment will be followed by a gradual ramp as other elements of the population reach their saturation time. This behaviour can be optimised by activating the power adjustment of individual elements in sequence. This could be used to either extend the duration of a regulating power service (at reduced magnitude) or provide a customised ramping rate, for example.

By performing analyses on a population of supermarkets in a similar manner to the analysis presented in this work, we can determine their capabilities, limitations, the complexities of operating a population of supermarkets for demand response, and the simplifications required to render the provision of demand response is practicable. Consequently, we can determine their suitability for participation in different electricity markets, or for provision of particular products. By determining a realistic view of the practical capabilities of a responsive population of supermarkets (or other application), we can evaluate the benefit that demand response will bring to the power system, the commercial feasibility of aggregators or other intermediaries, methods for market participation of aggregators, and consequently the economic feasibility of demand response, or business cases that would ensure its economic feasibility.

## 6 Conclusions

This report presents the initial work towards establishing a realistic model of the practical capabilities of demand response in a market environment, considering the particular case of supermarket refrigeration. An ARMAX model of a refrigeration test system has been identified, and the model time constants, 0.12 and 10 hours, indicate the potential for flexibility in power consumption over both short and long time horizons. Simulations have been conducted in a number of control environments to highlight the complexities of modelling the demand response behaviour that can be achieved from supermarket refrigeration systems. These complexities are induced because of the physical characteristics of the system and the control frameworks employed, and include asymmetric response in consumption to changes in temperature references, differing behaviour with different forecast horizons, the saturation of response, and the impact of control parameters on the flexibility attainable. These complexities impede the participation of demand response in a market environment, as market operators require participants to declare their willingness to participate in a very simple manner, often as simple as a quantity of energy and a corresponding price. Restrictions have been imposed on the operation of the supermarket refrigeration system considered in this work, so that a more simple representation of the demand response behaviour can be achieved. By imposing these simplifications, participation in a market can be facilitated, however with each simplification there is an associated loss of optimality in the demand response that can be achieved. The market considered in this work is the regulating, or balancing power market, in which market participants provide an adjustment of their power production/consumption for a given period of time. The refrigeration system has been shown to have demand response capabilities suitable for this market, where a response can be provided for up to 27.5 hours. Restricting the operation of the refrigeration system so that it commences each period of providing regulation from a steady state, we have shown that there is a non-linear relationship between the amount of power adjustment and the duration for which the regulation service can be provided. This relationship is sufficiently simple for a market intermediary, or aggregator representing a supermarket (or population thereof), to communicate its demand response capabilities to a market operator in the form of a bid. A representation of the practical



demand response capabilities is central to determining the value and establishing the economic feasibility of demand response.

## References

- [1] N. O’Connell, P. Pinson, H. Madsen, and M. O’Malley, “Benefits and challenges of electrical demand response: A critical analysis.” Under Review.
- [2] H. Madsen, J. Parvizi, R. Halvgaard, L. E. Sokoler, J. B. Jørgensen, L. H. Hansen, and K. B. Hilger, *Control of Electricity Loads in Future Electric Energy Systems*, vol. 4 of *Intelligent Energy Systems*. Wiley, 2014.
- [3] O. Corradi, H. Ochsenfeld, H. Madsen, and P. Pinson, “Controlling electricity consumption by forecasting its response to varying prices,” *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 421–429, 2013.
- [4] A. Molderink, V. Bakker, M. G. C. Bosman, J. Hurink, and G. J. M. Smit, “Management and control of domestic smart grid technology,” *Smart Grid, IEEE Transactions on*, vol. 1, pp. 109–119, Sept 2010.
- [5] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, and M. Morari, “Reducing peak electricity demand in building climate control using real-time pricing and model predictive control,” in *Decision and Control (CDC), 2010 49th IEEE Conference on*, pp. 1927–1932, Dec 2010.
- [6] J. M. Lujano-Rojas, C. Monteiro, R. Dufo-Lopez, and J. L. Bernal-Agustin, “Optimum residential load management strategy for real time pricing (rtp) demand response programs,” *Energy Policy*, vol. 45, no. 0, pp. 671 – 679, 2012.
- [7] H. Allcott, “Rethinking Real-Time Electricity Pricing,” *Resource and Energy Economics*, vol. 33, no. 4, pp. 820 – 842, 2011. Special Section: Sustainable Resource Use and Economic Dynamics.
- [8] T. Hovgaard, L. Larsen, J. Jørgensen, and S. Boyd, “Nonconvex model predictive control for commercial refrigeration,” *International Journal of Control*, vol. 86, no. 8, pp. 1349–1366, 2013.
- [9] J. Arias, *Energy Usage in Supermarkets - Modelling and Field Measurements*. PhD thesis, Division of Applied Thermodynamics and Refrigeration, Department of Energy Technology, Royal Institute of Technology, 2005.
- [10] R. Furberg and C. Norberg, “Energy efficiency in supermarkets - project work in business, technology and leadership.” Stockholm, Royal Institute of Technology, 2000.
- [11] P. Pinson, “Is the weather hustling electricity markets?.” Inaugural Professor Lecture, May 2013.
- [12] S. E. Shafiei, H. Rasmussen, and J. Stoustrup, “Modeling supermarket refrigeration systems for demand-side management,” *Energies*, vol. 6, pp. 900–920, 2013.
- [13] P. Nyeng, J. Ostergaard, M. Togeby, and J. Hethey, “Design and implementation of frequency-responsive thermostat control,” in *Universities Power Engineering Conference (UPEC), 2010 45th International*, pp. 1–6, 2010.
- [14] R. Halvgaard, N. Poulsen, H. Madsen, and J. Jorgensen, “Economic model predictive control for building climate control in a smart grid,” in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, pp. 1–6, Jan 2012.
- [15] J. Lopes, F. Soares, and P. Almeida, “Integration of electric vehicles in the electric power system,” *Proceedings of the IEEE*, vol. 99, pp. 168–183, Jan 2011.
- [16] T. Pedersen, P. Andersen, K. M. Nielsen, H. Starmose, and P. D. Pedersen, “Using heat pump energy storages in the power grid,” in *Control Applications (CCA), 2011 IEEE International Conference on*, pp. 1106–1111, Sept 2011.
- [17] T. Hastie and R. Tibshirani, *Generalized Additive Models*. Chapman & Hall, 1990.
- [18] MATLAB, *System Identification Toolbox, Matlab version 8.0 (R2012b)*. Natick, Massachusetts: The MathWorks Inc., 2012.
- [19] GAMS, *General Algebraic Modeling System (GAMS), version 23.7*. Washington D.C, USA: GAMS Development Corporation, 2011.
- [20] J. P. S. Catalao, S. J. P. S. Mariano, V. M. F. Mendes, and L. A. F. M. Ferreira, “Scheduling of head-sensitive cascaded hydro systems: A nonlinear approach,” *Power Systems, IEEE Transactions on*, vol. 24, pp. 337–346, Feb 2009.

- [21] C. Bang, F. Fock, and M. Togeby, “The existing nordic regulating power market,” tech. rep., EA Energy Analyses, 2012.
- [22] J. Kondoh, N. Lu, and D. Hammerstrom, “An evaluation of the water heater load potential for providing regulation service,” *Power Systems, IEEE Transactions on*, vol. 26, pp. 1309–1316, Aug 2011.
- [23] L. Thompson, K. Song, A. Lekov, and A. McKane, “Automated demand response opportunities in wastewater treatment facilities,” tech. rep., Ernest Orlando Lawrence Berkeley National Laboratory, 2008.
- [24] M. Zugno, J. Morales, P. Pinson, and H. Madsen, “A Bilevel Model for Electricity Retailers Participation in a Demand Response Market Environment,” *Energy Economics*, vol. 36, no. 0, pp. 182–197, 2012.
- [25] W. Zhang, K. Kalsi, J. Fuller, M. Elizondo, and D. Chassin, “Aggregate model for heterogeneous thermostatically controlled loads with demand response,” in *IEEE Power and Energy Society General Meeting*, pp. 1–8, July 2012.
- [26] Y. Zong, D. Kullmann, A. Thavlov, O. Gehrke, and H. Bindner, “Application of model predictive control for active load management in a distributed power system with high wind penetration,” *IEEE Transactions on Smart Grids*, vol. 3, pp. 1055–1062, June 2012.
- [27] C. De Jonghe, B. Hobbs, and R. Belmans, “Optimal generation mix with short-term demand response and wind penetration,” *IEEE Transactions on Power Systems*, vol. 27, pp. 830–839, May 2012.
- [28] S. H. Madaeni and R. Sioshansi, “The impacts of stochastic programming and demand response on wind integration,” *Energy Systems*, vol. -, pp. 1–16, 2012.
- [29] R. Sioshansi, “Evaluating the impacts of real-time pricing on the cost and value of wind generation,” *IEEE Transactions on Power Systems*, vol. 25, pp. 741–748, May 2010.
- [30] P. Thorsnes, J. Williams, and R. Lawson, “Consumer responses to time varying prices for electricity,” *Energy Policy*, vol. 49, no. 0, pp. 552–561, 2012. Special Section: Fuel Poverty Comes of Age: Commemorating 21 Years of Research and Policy.