Preprints of 4th IFAC Nonlinear Model Predictive Control Conference
International Federation of Automatic Control
Noordwijkerhout, NL. August 23-27, 2012

Fast Nonconvex Model Predictive Control for Commercial Refrigeration

Tobias Gybel Hovgard *,** **Lars F. S. Larsen** *
John Bagterp Jørgensen ** **Stephen Boyd** ***

* *Vestas Technology R&D, DK-8200 Aarhus N, Denmark (e-mail: {togho,lfsla}@vestas.com).*
** *DTU Informatics, Technical University of Denmark, DK-2800 Lyngby, Denmark (e-mail: jbj@imm.dtu.dk)*
*** *Information Systems Laboratory, Department of Electrical Engineering, Stanford University, USA, (e-mail: boyd@stanford.edu)*

Abstract: We consider the control of a commercial multi-zone refrigeration system, consisting of several cooling units that share a common compressor. The goal is to minimize the total energy cost, using real-time electricity prices, while obeying temperature constraints on the zones. We propose a variation on model predictive control to achieve this goal. When the right variables are used, the dynamics of the system are linear, and the constraints are convex. The cost function, however, is nonconvex. To handle this nonconvexity we propose a sequential convex optimization method, which typically converges in fewer than 5 or so iterations. We employ a fast convex quadratic programming solver to carry out the iterations, which is more than fast enough to run in real-time. We demonstrate our method on a realistic model, with a full year simulation, using real historical data. These simulations show substantial cost savings, and reveal how the method exhibits sophisticated response to real-time variations in electricity prices. This demand response is critical to help balance real-time uncertainties associated with large penetration of intermittent renewable energy sources in a future smart grid.

Keywords: Predictive Control, Optimization, Nonlinear Control, Smart Power Applications.

1. INTRODUCTION

To obtain an increasing amount of electricity from intermittent energy sources such as solar and wind, we must not only control the production of electricity, but also the consumption, in an efficient, flexible and proactive manner. The smart grid will be the future intelligent electricity grid that incorporates this. The Danish transmission system operator (TSO) defines it as: “Intelligent electrical systems that can integrate the behavior and actions of all connected users—those who produce, those who consume and those who do both—to provide a sustainable, economical and reliable electricity supply efficiently” (Energinet.dk, 2011).

In Denmark around 4500 supermarkets consume more than 550,000 MWh annually. This corresponds roughly to 2% of the entire electricity consumption in the country. Refrigerated goods constitute a large capacity in which energy can be stored in the form of ’coldness’. As this is not exploited by the thermostat (hysteresis) control policy most commonly used today, we propose an economic optimizing model predictive controller, economic MPC, to address this. MPC based on optimizing economic objectives has only recently emerged as a general methodology with efficient numerical implementations and provable stability properties (Diehl et al., 2011; Angeli et al., 2011) and in, *e.g.*, Hovgaard et al. (2012a) we demonstrated its capability to minimize the total cost of energy for a commercial refrigeration system while enabling it to participate in demand response schemes. The economic MPC has the ability to choose the optimal cooling strategy by utilizing the thermal capacity to shift the consumption in time, while keeping the temperatures within certain bounds.

An underlying challenge in applying MPC to vapor compression refrigeration systems is that the classical thermodynamics models are quite complex, and include many nonlinearities. One approach, called nonlinear MPC (NMPC), is to accept the optimization problem to be solved as nonlinear and nonconvex, and use generic nonlinear optimization methods, such as sequential quadratic programming (SQP) (Boggs and Tolle, 1995). This is the approach taken in Hovgaard et al. (2012a), which used ACADO (Houska et al., 2010), a generic nonlinear optimal control code, to solve the optimization problems. NMPC is widely used in the chemical process industry (see, *e.g.*, Biegler (2009)) but in general it requires special attention to ensure (local) convergence, and the computational complexity can be prohibitively high. Our method differs from NMPC: Instead of a generic SQP (or other) method, we use a a sequential convex programming (SCP) method, in which the objective is approximated by a convex function in each iteration; the convex parts are preserved, giving us the speed and reliability of solvers for convex optimization (Boyd and Vandenberghe, 2004). Our method, like SQP, involves the solution of a sequence of (convex) quadratic programs (QPs), but differs very much in how the QPs are formed. In SQP, an approximation to the Lagrangian of the problem is used; the linearization required in each step can end up dominating the computation (Dinh et al., 2011). In our SCP method, the convexification step is quite straightforward. We use the tool CVXGEN (Mattingley and Boyd, 2012) to generate fast custom solvers for the QPs that arise in our method, achieving solution times measured in milliseconds.

We show careful numerical simulations on a realistic supermarket refrigeration system using prediction models for outdoor temperatures and real-time electricity prices based on actual data. CVXGEN transforms the original optimization problem into a standard form quadratic program that solves in a couple of milliseconds. This extreme speed allows us to carry out a simulation for a full year with 15-minute increments in around 4 minutes on a single-core processor. The results are quite interesting too. Immediately we see cost savings in the order of 30%. We show that our MPC controller exhibits a sophisticated form of demand response to prices, reducing consumption when the prices are high and pre-cooling when prices are low. Further details and results are provided in Hovgaard et al. (2012b).

Several publications have reported the use of NMPC to control refrigeration systems. See, *e.g.*, Leducq et al. (2006); Elliott and Rasmussen (2008); Sonntag et al. (2008). Predictive control for energy cost reductions in vapor compression cycles have to some extend been investigated for building temperature regulation too. Oldewurtel et al. (2010); Ma et al. (2012b,a) all use weather predictions or time of use pricing to optimize the energy efficiency. However, most of these confine themselves to simple descriptions of the energy consumption, disregarding the interdependency of the control variables and the efficiency. Ma and Borrelli (2012) uses sequential quadratic programming (SQP) to solve this problem. These methods yield long computational times, *e.g.*, starting from 10–13 seconds per step on a 3.00GHz dual-core processor. For general reviews of the use of thermal storage and for the importance of MPC in demand response schemes see, *e.g.*, Camacho et al. (2011); Arteconi et al. (2012). The need for computationally efficient optimization in MPC applied to systems with either fast sampling or limited computational resources is considered in an increasing number of publications such as Diehl et al. (2002); Zeilinger et al. (2008); Diehl et al. (2009); Wang and Boyd (2010). Embedded convex optimization applications have recently become more available to non-experts by the introduction of the automatic code generator CVXGEN (Mattingley and Boyd, 2012).

2. COMMERCIAL REFRIGERATION

In this section we describe the dynamic model of a commercial multi-zone refrigeration system. Such systems can include supermarkets, warehouses, or air-conditioning.

2.1 Model

The model describes a system with multiple cold rooms in which a certain temperature for the stored foodstuff has to be maintained. We describe the temperature dynamics and the energy cost of the system using SI units throughout.

The refrigeration system considered utilizes a vapor compression cycle in which a refrigerant circulates in a closed loop consisting of a compressor, an expansion valve and two heat exchangers, an evaporator in the cold storage room, as well as a condenser/gas cooler located in the surroundings. When the refrigerant evaporates, it absorbs heat from the cold reservoir which is rejected to the hot reservoir. To sustain these heat transfers, the evaporation temperature $T_{\mathrm{e}}(t)$ has to be lower than the temperature in the cold reservoir $T_{\mathrm{air}}(t)$ and the condensation temperature has to be higher than the temperature at the hot reservoir $T_{\mathrm{a}}(t)$. Low pressure refrigerant, with the pressure $P_{\mathrm{e}}(t)$, from the outlet of the evaporator is compressed in the compressors to a high pressure $P_{\mathrm{c}}(t)$ at the inlet to the condenser to increase the saturation temperature. In these expressions t denotes time. To lighten notation, we will drop the time argument (t) in time-dependent functions in the sequel. The setup is sketched in Fig. 1, with one cold storage room and one frost room connected to the system. Usually, several cold storage rooms, *e.g.*, display cases, connect to a common compressor rack and condensing unit. Because of this, the individual display cases see the same evaporation temperature; whereas each unit has its own inlet valve for individual temperature control.

2.2 Temperature dynamics

We use a first principles model and describe the dynamics in the cold room by simple energy balances. The temperature of the foodstuff is denoted by $T_{\mathrm{food}}(t)$ and satisfies the differential equation,

$$m_{\mathrm{food}} c_{\mathrm{p,food}} \frac{dT_{\mathrm{food}}}{dt} = \dot{Q}_{\mathrm{food-air}}, \tag{1}$$

where $\dot{Q}_{\mathrm{food-air}}(t)$ is the energy flow from the air in the cold room to the foodstuff, m_{food} is the (assumed constant) mass of food, and $c_{\mathrm{p,food}}$ is the constant specific heat capacity of the food. The temperature of the air in the cold room $T_{\mathrm{air}}(t)$ satisfies the differential equation,

$$m_{\mathrm{air}} c_{\mathrm{p,air}} \frac{dT_{\mathrm{air}}}{dt} = \dot{Q}_{\mathrm{load}} - \dot{Q}_{\mathrm{food-air}} - \dot{Q}_{\mathrm{e}}, \tag{2}$$

where $\dot{Q}_{\mathrm{food-air}}(t)$ is the energy flow from the air to the foodstuff, $\dot{Q}_{\mathrm{e}}(t)$ is the applied cooling capacity (energy absorbed in the evaporator), $\dot{Q}_{\mathrm{load}}(t)$ is heat load from the surroundings to the air, m_{air} is the constant mass of air, and $c_{\mathrm{p,air}}$ is the constant specific heat capacity of the air. We describe the heat flows using Newton's law of cooling,

$$\begin{aligned} \dot{Q}_{\mathrm{food-air}} &= k_{\mathrm{food-air}}(T_{\mathrm{air}} - T_{\mathrm{food}}), \\ \dot{Q}_{\mathrm{load}} &= k_{\mathrm{amb-cr}}(T_{\mathrm{amb}} - T_{\mathrm{air}}) + \dot{Q}_{\mathrm{dist}}, \\ \dot{Q}_{\mathrm{e}} &= k_{\mathrm{evap}}(T_{\mathrm{air}} - T_{\mathrm{e}}), \end{aligned}$$

where k is the constant overall heat transfer coefficient between two media, $T_{\mathrm{amb}}(t)$ is the temperature of the ambient air which puts the heat load on the refrigeration system, and $\dot{Q}_{\mathrm{dist}}(t)$ is a disturbance to the load (*e.g.*, an injection of heat into the cold room).

2.3 Energy cost

The energy used by the compressor, denoted $\dot{W}_{\mathrm{c}}(t)$, dominates the power consumption in the system. It can be

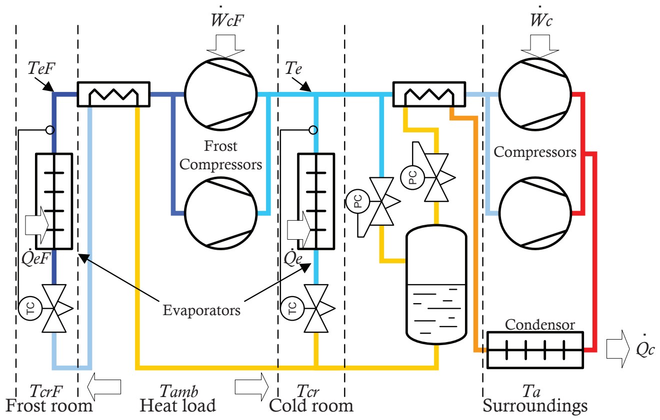

Fig. 1. Schematic layout of basic refrigeration system.

expressed by the mass flow of refrigerant $m_{\text{ref}}(t)$ and the change in energy content. We describe energy content by the enthalpy of the refrigerant at the inlet and at the outlet of the compressor ($h_{\text{ic}}(t)$ and $h_{\text{oc}}(t)$, respectively). These enthalpies are refrigerant-dependent functions of T_{e} and P_{c} (or equivalently, outdoor temperature T_{a}) as denoted in (3). They are computed using, *e.g.*, the software package REFEQNS (Skovrup, 2000), which models the thermodynamical properties of different refrigerants. We describe $\dot{W}_{\text{c}}$ as

$$\dot{W}_{\text{c}} = \frac{m_{\text{ref}}\left(h_{\text{oc}}(T_{\text{e}}, P_{\text{c}}) - h_{\text{ic}}(T_{\text{e}})\right)}{\eta_{\text{is}}(P_{\text{c}}/P_{\text{e}})(1 - \eta_{\text{heat}})}, \qquad (3)$$

where the isentropic efficiency $\eta_{\text{is}}(t)$ is a function mapping the pressure ratio over the compressor into compression efficiency and η_{heat} is a constant heat loss (in per cent) from the compressor. The mass flow is determined as the ratio between cooling capacity and change of enthalpy over the evaporator $(h_{\text{oe}}(t) - h_{\text{ie}}(t))$:

$$m_{\text{ref}} = \frac{\dot{Q}_{\text{e}}}{h_{\text{oe}}(T_{\text{e}}) - h_{\text{ie}}(P_{\text{c}})}.$$

For the efficiency function η_{is} we fitted a polynomial model of the form,

$$\eta_{\text{is}}(\alpha) = c_1 + c_2\alpha + c_3\alpha^{1.5} + c_4\alpha^3 + c_5\alpha^{-1.5},$$

where $c_1, \ldots, c_5$ are constant parameters. We found this description to be accurate within 1%.

Another compressor sits between the frost evaporator and the suction side of the other compressors, as seen in Fig. 1. This compressor decreases the evaporation temperature for the frost part of the system to a lower level. We use the subscript F to denote variables related to the frost part.

We describe the instantaneous energy cost of operating the system by multiplying power consumption by the real-time electricity price $p_{\text{el}}(t)$. The energy cost C over the period $[T_0, T_{\text{final}}]$ is

$$C = \int_{T_0}^{T_{\text{final}}} p_{\text{el}}\left(\dot{W}_{\text{c}} + \dot{W}_{\text{cF}}\right) dt. \qquad (4)$$

2.4 Control

Manipulated variables: Our controller manipulates the cooling capacity in each zone and the evaporation temperatures T_{e} and T_{eF}. The latter two are common for the entire refrigeration part and the entire frost part, respectively. In practice this is achieved by setting the set-points for inner control loops which operate with a high sample rate (compared to our control). This fast local control system allows us to ignore the complex and highly nonlinear behavior in the gas-liquid mixture in the evaporator.

Measured variables: The controller bases its decisions on measurements of air and food temperatures in each unit, on the known current outdoor temperature and electricity price, and on the predicted future values of the latter two. The heat disturbances are unknown.

2.5 Constraints

We would like the food temperatures to satisfy the inequalities

$$T_{\text{food,min}} \leq T_{\text{food}} \leq T_{\text{food,max}}, \qquad (5)$$

where $T_{\text{food,min}}$ and $T_{\text{food,max}}$ are a given allowable range given for each of the individual units. With randomly occurring load disturbances, it is not possible to guarantee that the temperatures are always in this range. So in lieu of imposing the constraints, we encode (5) as a set of soft constraints, *i.e.*, as a term added to the cost function,

$$V = \int_{T_0}^{T_{\text{final}}} \rho_{\text{soft,max}}(T_{\text{food}} - T_{\text{food,max}})_+ + \rho_{\text{soft,min}}(T_{\text{food,min}} - T_{\text{food}})_+ dt.$$

We choose the positive constants $\rho_{\text{soft,max}}$ and $\rho_{\text{soft,min}}$ so that violations are very infrequent in closed-loop operation. This formulation ensures a feasible problem even in the presence of uncertain loads. In a stochastic formulation, such as the one presented in Hovgaard et al. (2011), probabilistic constraints guarantee a feasible problem.

In addition, two constraints that cannot be violated are given,

$$0 \leq \dot{Q}_{\text{e}} \leq k_{\text{evap,max}}(T_{\text{air}} - T_{\text{e}}), \qquad (6)$$

$$0 \leq \dot{W}_{\text{c}} \leq \dot{W}_{\text{c,max}}, \qquad (7)$$

where $k_{\text{evap,max}}$ is the constant overall heat transfer coefficient from the refrigerant to the air when the evaporator is completely full and $\dot{W}_{\text{c,max}}$ is the constant limit on maximum energy consumption in the compressors. We define the set Ω as all $(\dot{Q}_{\text{e}}, T_{\text{e}})$ that satisfy the system dynamics (1)–(2) and the constraints (6)–(7).

2.6 Thermostat control

Today, most display cases and cold rooms are controlled by a thermostat. This means that maximum cooling is applied when the cold room temperature reaches an upper limit and shut off when the lower limit is reached. The advantage of this control policy is that it is simple

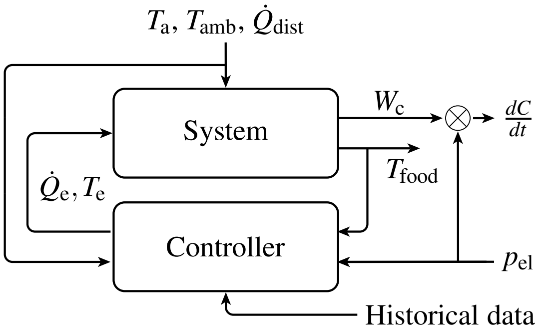

Fig. 2. Block diagram of the MPC controller.

and robust. The disadvantages, however, include: a high operating cost since the controller is completely unaware of system efficiency and electricity prices, no capability of demand response, and no specific handling of disturbances. All of these are addressed in our proposed method by intelligently exploiting the thermal capacity in the refrigerated mass.

3. METHOD

Fig. 2 outlines the overall structure of the proposed method and in the following sections we describe the details of the controller.

3.1 Economic MPC controller

The refrigeration system is influenced by a number of disturbances which we can predict (with some uncertainty) over a time horizon into the future. The controller must obey certain constraints, while minimizing the cost of operation. Economic MPC addresses all these concerns. Whereas the cost function in MPC traditionally penalizes a deviation from a set-point, the proposed economic MPC directly reflects the actual costs of operating the plant. This formulation is tractable for refrigeration systems, where we are interested in keeping the outputs (cold room temperatures) within certain ranges, while minimizing the cost of doing so.

Like in traditional MPC, we implement the controller in a receding horizon manner, where an optimization problem over N time steps (the control and prediction horizon) is solved at each step. The result is an optimal input sequence for the entire horizon, out of which only the first step is implemented. The controller aims at minimizing the electricity cost of operation. This cost relates to the energy consumption but we do not aim specifically at minimizing this, nor do we focus on tracking certain temperatures in the cold rooms. The optimization problem is thus formulated as

$$\begin{aligned} &\text{minimize} \quad C + V, \\ &\text{subject to} \ (\dot{\mathbf{Q}}_\text{e}, \mathbf{T}_\text{e}) \in \Omega, \\ &\qquad\qquad T_\text{food}^{T_\text{final}} = (T_\text{food,min} + T_\text{food,max})/2, \end{aligned} \tag{8}$$

where the variables are $\dot{Q}_\text{e}$ and T_e (both functions of time). The feasible set Ω imposes the system dynamics and constraints, and is defined by (1)–(2) and (6)–(7). We add a terminal constraint that the final food temperature $T_\text{food}^{T_\text{final}}$ must be at the midpoint of the allowable range of temperatures.

Instead of (8) we solve a discretized version with N steps over the time interval $[T_0, T_\text{final}]$,

$$\dot{\mathbf{Q}}_\text{e} = \left\{\dot{Q}_\text{e}^k\right\}_{k=0}^{N-1}, \quad \mathbf{T}_\text{e} = \left\{T_\text{e}^k\right\}_{k=0}^{N-1}. \tag{9}$$

The MPC feedback law is the first move in (9). The controller uses the initial state as well as predictions of the real-time electricity cost, the outdoor temperature and the injected heat loads for the time interval.

3.2 Sequential convex programming method

The feasible set Ω, the terminal constraint, and the cost function term V are all convex. Unfortunately, as C is nonconvex in the controllable variables $\dot{Q}_\text{e}$ and T_e, the problem in (8) is not convex. Instead of using a generic nonlinear optimization tool, we choose to solve the optimization problem iteratively using convex programming, replacing the nonconvex cost function C with a convex approximation. We express (4) using the coefficients of performance, COP,

$$\hat{C}^i = \int_{T_0}^{T_\text{final}} p_\text{el} \left(\frac{1}{\hat{\eta}_\text{COP}^i} \dot{Q}_\text{e} + \frac{1}{\hat{\eta}_\text{COP,F}^i} \dot{Q}_\text{eF} \right) dt, \tag{10}$$

where the COPs, $\hat{\eta}_\text{COP}^i$ and $\hat{\eta}_\text{COP,F}^i$, are complicated functions of the outdoor temperature and of the controllable variables $\dot{Q}_\text{e}$ and T_e. For any given values of these variables we can, however, compute the COP. Our approximation in each step is simple and natural: We use the COP calculated for the last iteration trajectory. Thus in each iteration we solve a convex optimization problem, which can be done very reliably and extremely quickly.

While our proposed method gives no theoretical guarantee on the performance, we must remember that the optimization problem is nothing but a heuristic for computing a good control and that the quality of closed-loop control with MPC is generally good without solving each problem accurately. Indeed, we have found that very early termination of this sequential convex programming method, well before convergence, still yields very good quality closed-loop control.

Algorithm 1 outlines the method. In the algorithm, φ_prox and φ_roc are regularization terms which we describe in §3.3.

Algorithm 1 Iterative optimization with nonconvex objective.

Initialize
$\dot{Q}_\text{e}^0$, T_e^0, and $i = 1$.
Compute
$\hat{\eta}_\text{COP}^i$ and $\hat{\eta}_\text{COP,F}^i$, as functions of $\{\dot{Q}_\text{e}, T_\text{e}\}^{i-1}$ and T_a.
Solve
$$\begin{aligned} &\text{minimize} \quad \hat{C}^i + V + \varphi_\text{prox} + \varphi_\text{roc}, \\ &\text{subject to} \ (\dot{Q}_\text{e}^i, T_\text{e}^i) \in \Omega, \\ &\qquad\qquad T_\text{food}^{T_\text{final},i} = \big(T_\text{food,min} + T_\text{food,max}\big)/2, \end{aligned}$$
Update
$\dot{Q}_\text{e}^i$, T_e^i, and $i = i + 1$
Repeat until convergence.

In Hovgaard et al. (2012c) we concluded that a unique minimum of the power consumption function exists within the feasible region. This assures that an iterative approach will converge to the intended extremum point.

3.3 Regularization

We use two different types of regularization in the optimization problem. To avoid oscillations from iteration to iteration we add proximal regularization of the form

$$\varphi_{\text{prox}} = \rho_{\text{prox}} \sum_{k=0}^{N-1} \|\dot{Q}_{\text{e}}^{k} - \dot{Q}_{\text{e}}^{k,\text{prev}}\|_2^2, \tag{11}$$

where the superscript 'prev' indicates that it is the solution from the previous iteration and ρ_{prox} is a constant weight chosen to damp large steps in each iteration. Smaller steps will of course increase the number of iterations required for the sequential convex programming method to converge, but, since we warm-start the algorithm from the soliution in the previous time step, the difference is negligible. Without proximal regularization oscillatory behavior can occur due to the nature of the thermodynamics in the refrigeration system. In addition, we add a quadratic penalty on the rate-of-change (roc) of $\dot{Q}_{\text{e}}$,

$$\varphi_{\text{roc}} = \rho_{\text{roc}} \sum_{k=1}^{N-1} \|\dot{Q}_{\text{e}}^{k} - \dot{Q}_{\text{e}}^{k-1}\|_2^2. \tag{12}$$

This regularization term serves two purposes: it improves the convergence of the sequential programming method, and also discourages rapid changes or switches in compressor levels, which helps reduce wear and tear of the compressor. Adding (11) and (12) to the linear objective formed by $\hat{C} + V$ results in a QP which we must solve once in each iteration. Due to the special structure of the MPC problem this QP is sparse; see, *e.g.*, Jørgensen (2005); Wang and Boyd (2010).

3.4 Non-homogeneous sampling

Speed of computation is a major concern in this work and we want to limit the size of the QPs that we solve in each iteration. A sampling time of 15 minutes directly gives 96 steps to be computed for a 24-hour prediction horizon. By using non-homogeneous sampling over the prediction horizon, exploiting that great accuracy becomes less important towards the end of the open-loop sequence, the number of steps can be reduced.

4. CASE STUDY

By simulation of realistic case studies we have verified the functionality and performance of the proposed MPC controller.

4.1 Scenario

Data from supermarkets actually in operation in Denmark have been collected. From these data, typical parameters such as time constants, heat loads, temperature ranges, capacities, and normal control policies have been estimated for three very different units; a milk cold room, a vertical shelving display case and a frost storage room. These units differ widely in load, mass of goods, and temperature demands.

We convert the system in §2.1 to the discrete-time equivalent. Since inner control loops are in place we have found that a sampling time of 15 minutes for the MPC controller is appropriate.

We model a contribution from the uncertain load by a 40% increase in the normal heat load. The increase occurs at random instances in 25% of the 15-minute periods. To account for this, back-offs from the temperature limits are introduced. We adjust these such that violations of the limits occur only 0.5–1% of the time. Less than 0.1° is often sufficient.

In our scenario we use temperature measurements from a meteorological station in the Danish city Sorø sampled every 30 minutes, along with hourly electricity spot prices downloaded from the Nordic electricity market, Nordpool. We simulate the scenario with data covering an entire calender year and use three years of data for training the predictors.

4.2 Algorithm details

We use a prediction horizon augmented of three sequences with increasing sample time; a 6-hour interval sampled every 15 minutes, a 6-hour interval sampled every 30 minutes and a 12-hour interval sampled every hour—resulting in 48 steps to be computed.

For regularization of the optimization problems the best behavior was observed with parameters in the order of $\rho_{\text{prox}} = 0.08$ and $\rho_{\text{roc}} = 0.06$; however, the method seems to be quite robust to changes in these values.

Recent advances in convex optimization allow for convex QPs to be solved at millisecond and microsecond time-scales. We use CVXGEN to generate a custom embedded solver for ultra fast computation of each convex QP in the sequential approach. CVXGEN transformed the original optimization problem into a standard form QP with 573 variables and 1248 constraints. In CVXGEN we specify and exploit the sparsity of the special problem structure.

4.3 Predictors

A prerequisite to solve the problem in (8) is to have available predictions of the outdoor temperatures and the electricity prices for the chosen prediction horizon, N. Only past values of such parameters can be available to the controller and in the present work we incorporate extremely simple predictors that can provide a sufficiently good estimate of the disturbances using a series of past measurements. We use historical data to train these predictors.

First, we use the historical training data set to create a baseline trajectory. For each month in a year we construct a typical day that describes the mean daily variation. If, *e.g.*, price is sampled every hour we get 24 prices for each of the 12 months. Next we calculate a residual (difference between baseline and historical data) for each one of the 12 baselines. For each of these, residual predictors are computed by

$$\text{minimize} \sum_{k=1}^{K} \|[R_{k-n}, \ldots, R_k]X - [R_{k+1}, \ldots, R_{k+N}]\|_2^2,$$

where K is the number of data points in the training data set, n is the number of past data points used for prediction, N is the number of future data points that we want to predict, X is the $n+1 \times N$ predictor matrix and R are the residuals. We employ an ℓ_1 regularization to avoid numerical instability that could lead to high variance models. Following this, a smoothed baseline is computed using interpolation of two adjacent months. Now, we can compute the predictions by first predicting the residuals of two adjacent months, interpolating these and adding them to the interpolated baseline of the same time window. We have chosen to use two days of past data for predicting the outdoor temperature and seven days for the price prediction. We use an entire week for the latter since the price pattern is different from weekdays to weekends. For both outdoor temperatures and electricity prices the training sets are defined from 1 January 2007 until 31 December 2009 and the simulation/test set covers the entire year of 2010.

For the unknown disturbance in the heat load we use a very simple predictor, namely the expected mean value of the random heat injection.

4.4 Computation times

We have simulated the proposed method with the case study described in the previous sections. The optimization problems solve in the order of a handful of milliseconds per MPC step which is more than fast enough for real-time implementation. A full year simulates in less than 4 minutes on a 2.8GHz Intel Core i7, excluding the time needed outside the optimization routine for predictors etc. The same problem with a generic solver such as ACADO takes around 4 minutes per MPC step on the same processor. For implementation in embedded industrial hardware a rough estimate of the computation time is around 1000 times of what we have observed here. This is still way below 10 seconds per time step which certainly allows for real-time implementation.

4.5 Convergence

When cold-started the proposed method generally converges in 10–20 iterations. In MPC, however, the open-loop trajectory from the previous run of the optimizer, shifted one time-step, is an excellent guess on the next outcome and is well-suited for warm-starting the algorithm. Using this warm start initialization, the method generally converges in fewer than 5 iterations. In addition, we find that early termination after, *e.g.*, 2–3 iterations generally gives good results, degrading the overall performance with less than 1%.

4.6 Savings

To benchmark the savings gained by introducing the proposed MPC controller, we have performed a simulation for the same system and conditions but using the conventional thermostat control policy. As in real systems the air temperature surrounding the foodstuff in each unit is the variable used in the thermostat. We have defined upper and lower bounds for switching on and off, such that the interval corresponds to what is normally observed in real operation. Besides, we determine the upper bound such that cooling quality is maintained at a minimal cost, *i.e.*, such that the food temperatures only violate the upper allowable limit in 0.5–1% of the time (to be comparable with the MPC control).

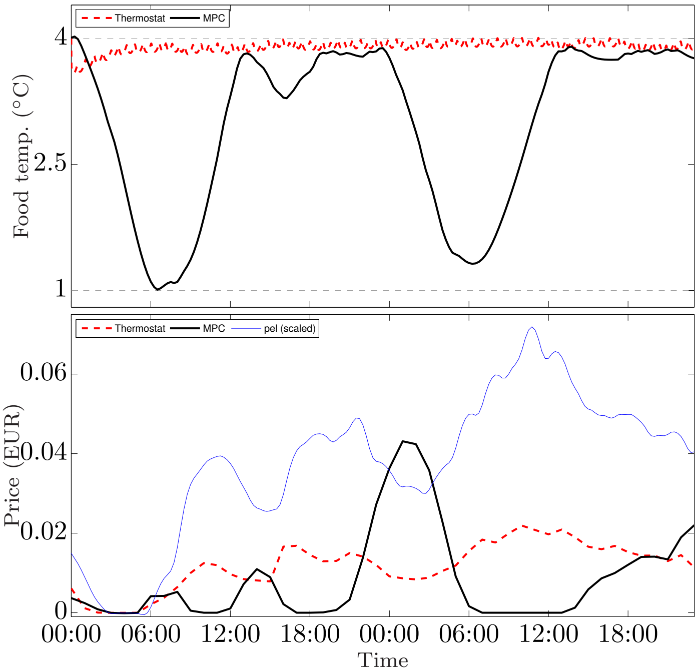

Fig. 3. Selected trajectory for food temperature and hourly cost of energy for control by thermostat vs. the proposed MPC.

Fig. 3 shows a segment of the simulated system with thermostat control versus the proposed MPC controller. We show the trajectory for one unit only and we observe how the food temperature is pulled down by the MPC controller at times with low electricity prices, meaning that pre-cooling is applied. At such times the instantaneous cost of operating the system might be higher than if the conventional thermostat is used, as can be seen on the figure. But this is, however, more than counteracted by the savings when the electricity prices go up.

In Fig. 4, resulting temperature distributions for a selected unit are shown for both control by thermostat and by MPC. While both control policies tend to keep the temperatures close to the upper limit most of the time, we observe how the MPC controller makes use of the entire range for storing coldness.

We observe savings in the order of 30%. Adding the uncertain heat load injections and the appropriate back-offs from the temperature limits, as described in §4, increases the overall cost by approximately 10%.

4.7 Demand response

Fig. 5 shows the total cooling energy applied to all three units plotted as a function of the electricity price at the time of use. We have selected one month to limit the

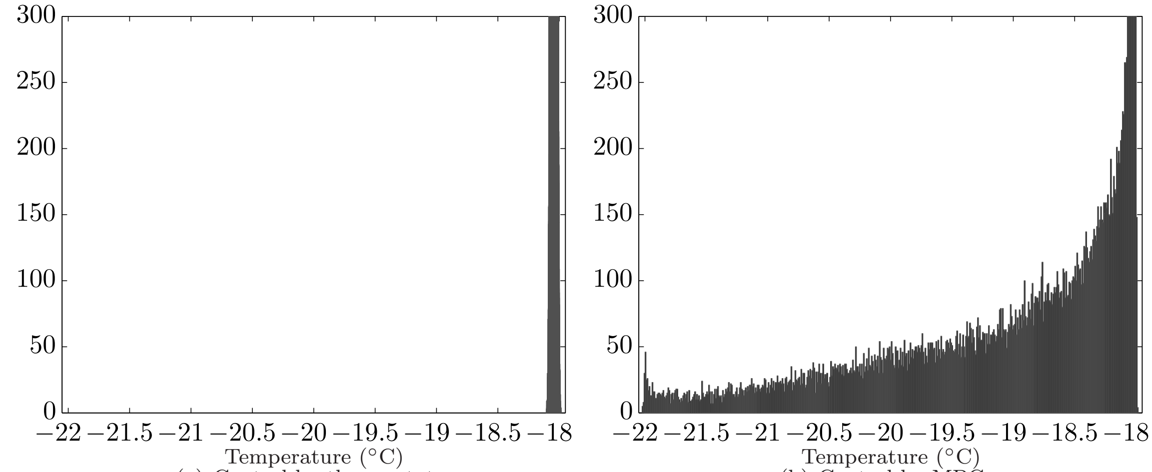

(a) Control by thermostat. (b) Control by MPC.

Fig. 4. Temperature distribution for selected unit. Simulation over the full year 2010.

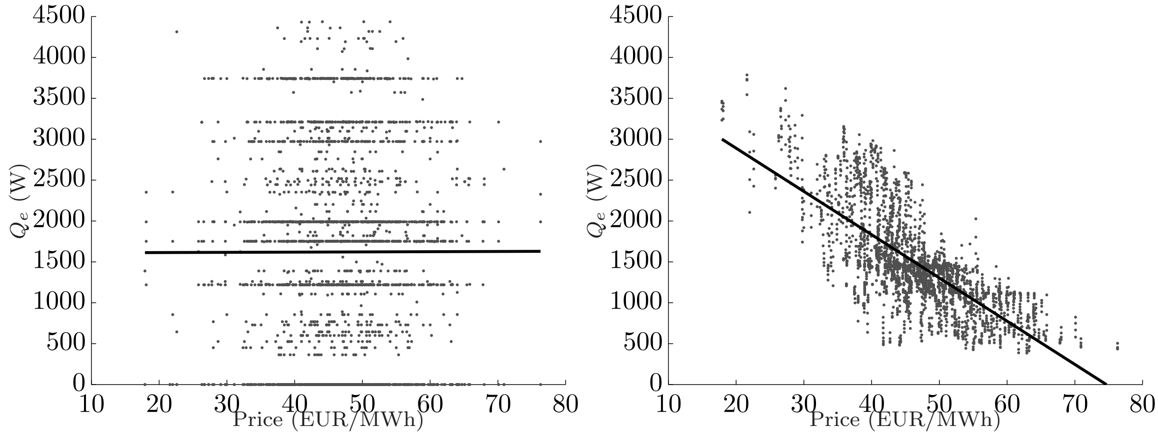

(a) Control by thermostat. The solid line is a linear fit with almost zero slope.

(b) Control by MPC. The solid line is a linear fit with a slope of −53 W/(EUR/MWh).

Fig. 5. Illustration of demand response in systems controlled by MPC vs. thermostat control.

number of data-points but the picture is almost identical for the entire year of simulation: We observe no correlation between energy consumption and electricity prices when the thermostat controls the refrigeration system while we see a clear tendency to apply more cooling at times with low prices, and vise versa, if we employ the proposed MPC scheme. A linear fit is made using a Huber function regression. The slope is around −50 W/(EUR/MWh) for the MPC controlled system as opposed to 0 for the thermostat which clearly illustrates the demand response behavior of the system. We should remember that the spot price used here is just an example and not a prerequisite of our method. In a smart grid the price signal could be artificially made by the balance responsible party to promote demand response.

4.8 Plant perturbations

With perturbations of up to at least 20–30% in parameters such as mass of the refrigerated foodstuff and the heat transfer coefficients we see essentially no changes in the closed-loop dynamics.

4.9 Perfect predictions

By again simulating over the full year of 2010, but this time with a prescient setting assuming knowledge of the exact future conditions instead of using their predictions, we are able to compare the performance of the simple predictors and give a rough judgment on how much the method relies on the availability of accurate predictions. We have observed that the extra savings gained by having the full information available are in the order of 1-2%. This justifies the use of simple predictors.

5. CONCLUSION

In this paper we have presented an MPC controller for a commercial multi-zone refrigeration system. We have based our method on convex optimization, solved iteratively to treat a nonconvex cost function. By employing a fast convex quadratic programming solver to carry out the iterations, the method is more than fast enough to run in real-time. Simulation on a realistic scenario reveal significant savings as well as convincing demand response capabilities suitable for implementation with smart grid schemes.

ACKNOWLEDGEMENTS

This work was carried out in collaboration with Danfoss Electronic Controls R&D, Refrigeration and Airconditioning, Nordborgvej 81, DK-6430 Nordborg, Denmark. We thank Ed Cazalet for helpful suggestions.

REFERENCES

Angeli, D., Amrit, R., and Rawlings, J. (2011). On average performance and stability of economic model predictive control. *Automatic Control, IEEE Transactions on*, PP(99), in press.

Arteconi, A., Hewitt, N.J., and Polonara, F. (2012). State of the art of thermal storage for demand-side management. *Applied Energy*, 93(0), 371 – 389.

Biegler, L. (2009). Efficient nonlinear programming algorithms for chemical process control and operations. In A. Korytowski, K. Malanowski, W. Mitkowski, and M. Szymkat (eds.), *System Modeling and Optimization*, volume 312 of *IFIP Advances in Information and Communication Technology*, 21–35. Springer Boston.

Boggs, P.T. and Tolle, J.W. (1995). Sequential quadratic programming. *Acta Numerica*, 4, 1–51.

Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.

Camacho, E.F., Samad, T., Garcia-Sanz, M., and Hiskens, I. (2011). Control for renewable energy and smart grids. In T. Samad and A.M. Annaswamy (eds.), *The Impact of Control Technology, Control Systems Society*, 69–88. IEEE Control Systems Society.

Diehl, M., Amrit, R., and Rawlings, J.B. (2011). A lyapunov function for economic optimizing model predictive control. *Automatic Control, IEEE Transactions on*, 56(3), 703 –707.

Diehl, M., Bock, H.G., Schlo, J.P., Findeisen, R., Nagy, Z., and Allgöwer, F. (2002). Real-time optimization and nonlinear model predictive control of processes governed by differential-algebraic equations. *Journal of Process Control*, 12(4), 577–585.

Diehl, M., Ferreau, H., and Haverbeke, N. (2009). Efficient numerical methods for nonlinear MPC and moving horizon estimation. In L. Magni, D. Raimondo, and F. Allgwer (eds.), *Nonlinear Model Predictive Control*, volume 384 of *Lecture Notes in Control and Information Sciences*, 391–417. Springer Berlin / Heidelberg.

Dinh, Q.T., Savorgnan, C., and Diehl, M. (2011). Real-time sequential convex programming for nonlinear model predictive control and application to a hydropower plant. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, 5905–5910.

Elliott, M.S. and Rasmussen, B.P. (2008). Model-based predictive control of a multi-evaporator vapor compression cooling cycle. In *American Control Conference, 2008*, 1463 –1468.

Energinet.dk (2011). Potential and opportunities for flexible electricity consumption with special focus on individual heat pumps (in Danish). Technical report, Energinet.dk, The Danish TSO owned by the Danish Climate and Energy Ministry. Denmark.

Houska, B., Ferreau, H., and Diehl, M. (2010). ACADO Toolkit–An Open Source Framework for Automatic Control and Dynamic Optimization. *Optimal Control Applications and Methods*.

Hovgaard, T.G., Larsen, L.F.S., Edlund, K., and Jørgensen, J.B. (2012a). Model predictive control technologies for efficient and flexible power consumption in refrigeration systems. *Energy*, in press.

Hovgaard, T.G., Larsen, L.F.S., and Jørgensen, J.B. (2011). Robust Economic MPC for a Power Management Scenario with Uncertainties. In *50th IEEE Conference on Decision and Control and European Control Conference*, 1515–1520.

Hovgaard, T.G., Larsen, L.F.S., Jørgensen, J.B., and Boyd, S. (2012b). Nonconvex model predictive control for commercial refrigeration. `http://www.stanford.edu/~boyd/papers/noncvx_mpc_refr.html`.

Hovgaard, T.G., Larsen, L.F.S., Skovrup, M.J., and Jørgensen, J.B. (2012c). Optimal energy consumption in refrigeration systems - modelling and non-convex optimisation. *The Canadian Journal of Chemical Engineering*, in press.

Jørgensen, J.B. (2005). *Moving Horizon Estimation and Control*. Ph.D. thesis, Department of Chemical Engineering, Technical University of Denmark.

Leducq, D., Guilpart, J., and Trystram, G. (2006). Nonlinear predictive control of a vapour compression cycle. *International Journal of Refrigeration*, 29(5), 761 – 772.

Ma, J., Qin, J., Salsbury, T., and Xu, P. (2012a). Demand reduction in building energy systems based on economic model predictive control. *Chemical Engineering Science*, 67(1), 92 – 100.

Ma, Y. and Borrelli, F. (2012). Fast stochastic predictive control for building temperature regulation. In *American Control Conference (ACC), 2012*, accepted.

Ma, Y., Borrelli, F., Hencey, B., Coffey, B., Bengea, S., and Haves, P. (2012b). Model predictive control for the operation of building cooling systems. *Control Systems Technology, IEEE Transactions on*, 20(3), 796 –803.

Mattingley, J. and Boyd, S. (2012). CVXGEN: a code generator for embedded convex optimization. *Optimization and Engineering*, 13, 1–27.

Oldewurtel, F., Parisio, A., Jones, C., Morari, M., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., and Wirth, K. (2010). Energy Efficient Building Climate Control using Stochastic Model Predictive Control and Weather Predictions. In *American Control Conference (ACC)*, 5100 – 5105.

Skovrup, M.J. (2000). Thermodynamic and thermophysical properties of refrigerants - software package in borland delphi. Technical report, Department of Energy Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark.

Sonntag, C., Devanathan, A., and Engell, S. (2008). Hybrid NMPC of a supermarket refrigeration system using sequential optimization. In *Proc of 17th IFAC World Congress*, 13901–13906.

Wang, Y. and Boyd, S. (2010). Fast model predictive control using online optimization. *Control Systems Technology, IEEE Transactions on*, 18(2), 267–278.

Zeilinger, M.N., Jones, C.N., and Morari, M. (2008). Real-time suboptimal model predictive control using a combination of explicit MPC and online optimization. In *Proc of. 47th IEEE Conference on Decision and Control (CDC)*, 4718–4723.