

RESEARCH ARTICLE

Nonconvex Model Predictive Control for Commercial Refrigeration

Tobias Gybel Hovgaard^{a,b*}, Lars F.S. Larsen^a, John Bagterp Jørgensen^b, and Stephen Boyd^c

^aVestas Technology R&D, DK-8200 Aarhus N, Denmark; ^bDTU Informatics, Technical University of Denmark, DK-2800 Lyngby, Denmark; ^cInformation Systems Laboratory, Department of Electrical Engineering, Stanford University, 94305 Stanford, USA

(Submitted June 2012, Accepted October 2012)

We consider the control of a commercial multi-zone refrigeration system, consisting of several cooling units that share a common compressor, and is used to cool multiple areas or rooms. In each time period we choose cooling capacity to each unit and a common evaporation temperature. 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 due to the temperature dependence of thermodynamic efficiency. 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 and 15 minute time periods, using historical electricity prices and weather data, as well as random variations in thermal load. These simulations show substantial cost savings, on the order of 30%, compared to a standard thermostat-based control system. Perhaps more important, we see that the method exhibits sophisticated response to real-time variations in electricity prices. This demand response is critical to help balance real-time uncertainties in generation capacity associated with large penetration of intermittent renewable energy sources in a future smart grid.

Keywords: Energy management, Optimization methods, Predictive Control, Nonlinear control systems, Smart grids.

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. In, *e.g.*, Finn et al. (2011) facilitation of wind generated electricity by price optimized thermal storage was described. In contrast to the current centralized power generation system, the electricity grid will be a network of many independent power generators. The smart grid will be the future intelligent electricity grid that incorporates all these. The Danish transmission system operator (TSO) has the following definition of smart grids which we adopt in this work: “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). Different means of utilizing demand response in a smart grid setting have been investigated in an increasing number of publications, *e.g.*, Andersson et al. (2010), Han et al. (2010), Saele and Grande (2011), Molina-Garcia et al. (2011), for plug-in electric vehicles and heat pumps. Kirschen (2003) investigated demand response and price elasticity and Pina et al. (2012) analyzed different demand side management strategies.

*Corresponding author. Email: togho@vestas.com

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'. The thermostat (hysteresis) control policy most commonly used today does not exploit this and a large potential for energy and cost reductions exists. Preliminary investigations have been carried out in Larsen et al. (2007), Hovgaard et al. (2010).

We propose an economic optimizing model predictive controller, economic MPC, to address this for a commercial refrigeration system. Predictive control—also known as receding horizon control—for constrained systems has emerged during the last 30 years as one of the most successful methodologies to control industrial processes (Qin and Badgwell 2003) and is increasingly being considered to control both refrigeration and power systems (Sarabia et al. 2008, Edlund et al. 2011, Blarke and Dotzauer 2011). MPC based on optimizing economic objectives has only recently emerged as a general methodology with efficient numerical implementations and provable stability properties (Rawlings and Amrit 2009, Diehl et al. 2011, Angeli et al. 2011) and is now considered for smart grid related problems too, (Hindi et al. 2012, Halvgaard et al. 2012). We have previously demonstrated the capability of economic MPC in, *e.g.*, Hovgaard et al. (2010, 2011, 2012) to minimize the total cost of energy for a commercial refrigeration system while enabling it to participate in demand response schemes. Economic MPC has the ability to choose the optimal cooling strategy from predictions of the disturbances such as load, efficiency, and price of electricity. This is achieved by utilizing the thermal capacity to shift the consumption in time, while keeping the temperatures within certain bounds. We choose these bounds so that they have no consequences for food quality and safety. Van Harmelen (2001), Bush and Wolf (2009), Oldewurtel et al. (2010) also described the use of load shifting capabilities to reduce total energy consumption. For other 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).

An underlying challenge in applying MPC to vapor compression refrigeration systems is that the classical thermodynamics models are quite complex, and include many nonlinearities, such as temperature-dependent efficiencies. 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. (2012), 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 in the following ways. First, our formulation (choice of variables) results in an optimization problem with linear constraints, but an objective function that is nonconvex. Instead of a generic SQP (or other) method, we use a sequential convex programming (SCP) method, in which the objective is approximated by a convex function in each iteration; the equality and inequality constraints, which are convex, 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 needed in each iteration is quite straightforward. Unlike SQP, our method does not exhibit terminal quadratic convergence, but since our method converges in practice in just a handful of iterations, this does not seem to be an issue, at least in this application. 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 describe the method in detail, and report careful numerical simulations on a realistic supermarket refrigeration system. For prediction of outdoor temperatures and real-time electricity

prices we build models using three years of historical data. With 15-minute sample time and a prediction horizon of 48 steps CVXGEN transforms the original optimization problem into a standard form QP with 573 variables and 1248 constraints, which can be solved by the custom solver 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. Immediately we see cost savings on the order of 30%. We see that MPC does pre-cooling, *i.e.*, cools to a lower-than-normal temperature (without leaving the acceptable temperature range) to reduce cooling needed at times with higher electricity prices. By scatter plotting electricity price and energy consumption, 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, while maintaining temperatures within the required range.

1.1 Prior work

Leducq et al. (2006) used NMPC with an iterative routine to optimize the coefficient of performance (COP) for a refrigeration plant while maintaining a fixed cooling capacity. Since the focus was not on load-shifting, a quadratic cost function was used to track the cooling capacity. As the cooling capacity was not a decision variable the problem became convex in the cost function. Still due to the computational burden, the prediction horizon was limited to 3–4 sample periods. Elliott and Rasmussen (2008) controlled a multi-evaporator refrigeration system with MPC that tracked energy efficient set-points. By optimizing only over the cooling capacity from each evaporator and using a PI controller based on the most loaded unit for controlling the evaporation temperature, the optimization problem rendered convex and linear. But this strategy completely disregarded these two variables' interdependency on the system efficiency. As we will see in this study, the multivariate problem has to be taken into account. A sequential NMPC approach was also used in Sonntag et al. (2008) to minimize the compressor switching. Even though computational complexity is not reported directly, the authors state that “the approach does not yield satisfactory results for larger systems due to the combinatorial growth of the search space.”

Predictive control and optimization for energy cost reductions in vapor compression cycles have been investigated for building temperature regulation too. Ma et al. (2012) considered time of use pricing in that context. The problem was formulated as a linear program (LP) but no specific details were given on how the power consumption was approximated. Oldewurtel et al. (2010), Ma et al. (2012a,b) all used weather predictions to optimize the energy efficiency. In the first, the cost reduced to a linear function while stochastic disturbances were handled by affine disturbance feedback. In the latter two, power consumption was implemented as a 5-D lookup and a move-blocking strategy was used to reduce computational burden. An average computational time of 20 minutes was reported.

SQP is a well known method used for NMPC and, *e.g.*, Ma and Borrelli (2012) applied a tailored SQP algorithm to building temperature regulation. However, the energy consumption model was a static function of the load on the air-side and again the control decisions' influence on the COP was lost. 10–13 seconds' computation times on a 3 GHz dual-core processor were reported. In Oldewurtel et al. (2012) the studies from Oldewurtel et al. (2010) were extended and a sequential LP algorithm was used to deal with a bilinear cost. No computational times were reported in this study.

The need for computationally efficient optimization in MPC applied to systems with either fast sampling or limited computational resources are considered in an increasing number of publications. In Diehl et al. (2002) a direct multiple shooting method was presented, capable of solving an NMPC problem with 42 differential states and 122 algebraic states over 20 control intervals in 10 s and in Wang and Boyd (2010) a quadratic MPC problem with 12 states, 3 controls, and a horizon of 30 intervals was solved in 5 ms using warm-starting. Another approach to real-time MPC is the explicit methods as reported in, *e.g.*, Zeilinger et al. (2008) where the technique was used in combination with online optimization for solving QPs under restrictions

on the computational time. Grancharova et al. (2007) gives an extension to explicit NMPC. However, it was reported that it is troublesome to ensure stability if the problem is nonconvex, and in addition, the explicit methods are not suitable for larger problems due to extremely large state-spaces. Approaches to parallel implementation of MPC algorithms for real-time execution were shown in, *e.g.*, Jerez et al. (2011) where a problem with 32 states, 16 inputs and 10 control intervals was solved in 344 ms on an FPGA. For further reviews of numerical methods for solution of real-time optimal control problems in NMPC see, *e.g.*, Diehl et al. (2009).

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). Remarkable speed-ups achieved using tailored QP-solvers exported from CVXGEN have been reported in, *e.g.*, Kraning et al. (2011), Mattingley et al. (2011). In a recent report (O’Donoghue et al. 2012) a splitting technique to a generic linear-convex optimal control problem is introduced and computation times faster than what is obtained by CVXGEN are reported. This suggests that our method could speed up even further.

1.2 Outline

In §2 we describe the dynamic models used for the commercial multi-zone refrigeration system. We define variables and constraints and briefly describe the control policy most commonly used in commercial refrigeration today. In §3 we establish an MPC controller for the system and give details on the proposed iterative optimization scheme. We describe the method for obtaining a convex approximate objective function and how to solve this using CVXGEN. We demonstrate the method by simulation of a case study for which we describe the scenario, along with very simple predictors in §4. Following this, the results of the numerical examples appear. We simulate the system for a full year and report on computation time, convergence, cost savings, and demand response behavior. In §5 we give concluding remarks.

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. We describe the thermodynamics, the constraints of the system and the function reflecting the economic cost of operating the plant.

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: energy flows and power consumption are in Watts, temperatures are in degrees centigrade, pressures are in Pascal, enthalpies are in J/kg, and instantaneous electricity prices are in EUR/W. This fixes the units of all quantities used.

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 $T_e(t)$ (given by the pressure $P_e(t)$) has to be lower than the temperature in the cold reservoir $T_{\text{air}}(t)$ and the condensation temperature has to be higher than the temperature at the hot reservoir $T_a(t)$. Low pressure refrigerant, with the pressure $P_e(t)$, from the outlet of the evaporator is compressed in the compressors to a high pressure $P_c(t)$ 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, but 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 $T_{\text{food}}(t)$ and satisfies the differential equation,

$$m_{\text{food}}c_{p,\text{food}}\frac{dT_{\text{food}}}{dt} = \dot{Q}_{\text{food-air}}, \quad (1)$$

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

$$m_{\text{air}}c_{p,\text{air}}\frac{dT_{\text{air}}}{dt} = \dot{Q}_{\text{load}} - \dot{Q}_{\text{food-air}} - \dot{Q}_{\text{e}}, \quad (2)$$

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

$$\dot{Q}_{\text{food-air}} = k_{\text{food-air}}(T_{\text{air}} - T_{\text{food}}),$$

$$\dot{Q}_{\text{load}} = k_{\text{amb-cr}}(T_{\text{amb}} - T_{\text{air}}) + \dot{Q}_{\text{dist}},$$

$$\dot{Q}_{\text{e}} = k_{\text{evap}}(T_{\text{air}} - T_{\text{e}}),$$

where k is the constant overall heat transfer coefficient between two media, $T_{\text{amb}}(t)$ is the temperature of the ambient air which puts the heat load on the refrigeration system, and $\dot{Q}_{\text{dist}}(t)$ 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 $\dot{W}_{\text{c}}(t)$, dominates the power consumption in the system. It can be 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 REFQNS (Skovrup 2000), which models the thermodynamical properties of different refrigerants. We describe \dot{W}_{c} as

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

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.

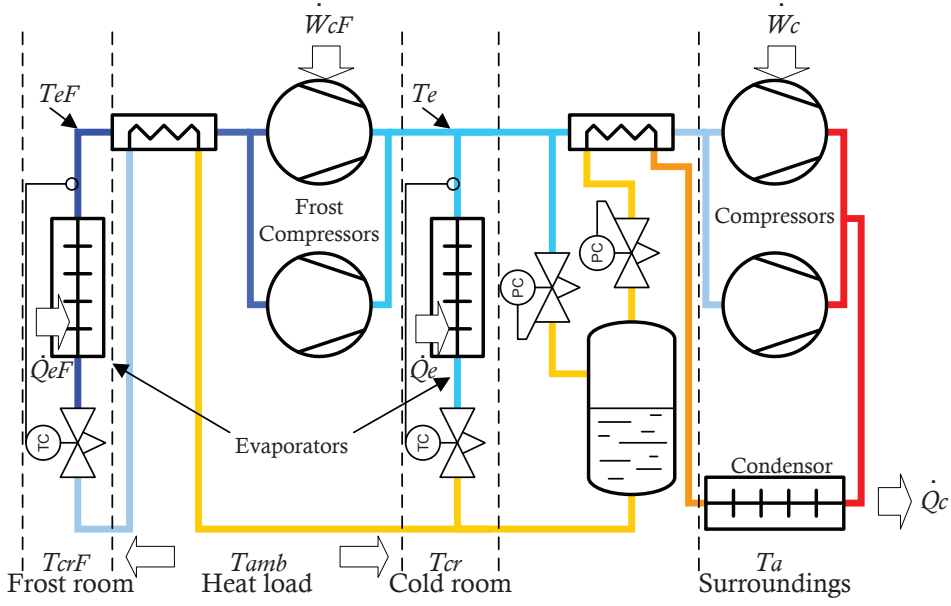


Figure 1. Schematic layout of basic refrigeration system.

The mass flow is determined as the ratio between cooling capacity and change of enthalpy over the evaporator ($h_{oe}(t) - h_{ie}(t)$):

$$m_{\text{ref}} = \frac{\dot{Q}_e}{h_{oe}(T_e) - h_{ie}(P_c)}.$$

The efficiency function η_{is} can be found in several ways. We used data from first principles thermodynamic calculations to fit a 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, \dots, c_5 are constant parameters. We found this approximation to be accurate within 1%. Fig. 2 shows η_{is} versus the pressure ratio $\alpha = P_c/P_e$.

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 can describe the work in the frost compressor by identical equations but the pressure at its outlet is determined by the evaporation temperature for the cooling part. The mass flow through the frost compressor adds to the flow through the cooling compressors. 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

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

For later reference we express (4) using the coefficients of performance, COP, ($\eta_{\text{COP}}(t)$ and $\eta_{\text{COP,F}}(t)$ respectively),

$$C = \int_{T_0}^{T_{\text{final}}} p_{\text{el}} \left(\frac{1}{\eta_{\text{COP}}} \dot{Q}_e + \frac{1}{\eta_{\text{COP,F}}} \dot{Q}_{eF} \right) dt.$$




Figure 2. Isentropic efficiency of the compressor as a function of the pressure ratio P_c/P_e .

$\eta_{\text{COP}}(t)$ and $\eta_{\text{COP,F}}(t)$ are complicated functions of the outdoor temperature and of the controllable variables \dot{Q}_e and T_e . For any given values of these variables we can, however, compute the coefficients of performance using the steps outlined in Algorithm 1.

Algorithm 1 Calculating the COP for a three-unit system.

Require:

1. Initial values: T_e and $\{\dot{Q}_{e,i}\}_{i=1}^3$.
2. Prediction of outdoor temperature T_a .

Compute:

1. Pressure in gas cooler P_c .
 2. Enthalpy into evaporator h_{ie} as a function of P_c .
 3. Enthalpy out from evaporators $h_{oe,i}$ as a function of T_e and $\dot{Q}_{e,i}$.
 4. Enthalpy into compressor h_{ic} using mass and energy balances to combine $h_{oe,i}$'s.
 5. Enthalpy out of compressor h_{oc} as a function of h_{ic} , T_e , and P_c .
 6. η_{is} as a function of T_e and P_c .
 7. COP as $\eta_{is}(1 - \eta_{\text{heat}})(h_{oei} - h_{ie})/(h_{oc} - h_{ic})$.
-

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

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

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,

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

where $(a)_+ = \max\{a, 0\}$. This objective term penalizes violations of the temperature range constraints. 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 by the nature of the system,

$$0 \leq \dot{Q}_e \leq k_{\text{evap,max}}(T_{\text{air}} - T_e), \quad (6)$$

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

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 $\dot{W}_{c,\text{max}}$ is the constant limit on maximum energy consumption in the compressors. We define the set Ω as all (\dot{Q}_e, T_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 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. 3 outlines the overall structure of the proposed method and in the following sections we describe the details of the controller.




Figure 3. Block diagram of the MPC 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

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

where the variables are \dot{Q}_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 $[T_0, T_{\text{final}}]$,

$$\dot{\mathbf{Q}}_e = \left\{ \dot{Q}_e^k \right\}_{k=0}^{N-1}, \quad \mathbf{T}_e = \left\{ T_e^k \right\}_{k=0}^{N-1}. \quad (9)$$

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. The predictions could come from any source, including national weather service, market or balance responsible parties on the power grid, etc. In this paper we use very simple implementations of predictors that we describe in §4.4.

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}_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,

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

where $\hat{\eta}_{\text{COP}}^i$ and $\hat{\eta}_{\text{COP,F}}^i$ are calculated for the i th iteration as in Algorithm 1 using \dot{Q}_e^{i-1} and T_e^{i-1} found in the previous iteration. Thus in each iteration we solve a convex optimization problem, which can be done very reliably and extremely quickly. Our approximation in each step is simple and natural: We use the coefficient of performance calculated for the last iteration trajectory.

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 2 outlines the method. In the algorithm, φ_{prox} and φ_{roc} are regularization terms which we describe in §3.3.

Algorithm 2 Iterative optimization with nonconvex objective.

Initialize

\dot{Q}_e^0 , T_e^0 , and $i = 1$.

Compute

$\hat{\eta}_{\text{COP}}^i$ and $\hat{\eta}_{\text{COP,F}}^i$, as functions of $\{\dot{Q}_e, T_e\}^{i-1}$ and T_a .

Solve

minimize $\hat{C}^i + V + \varphi_{\text{prox}} + \varphi_{\text{roc}}$,
 subject to $(\dot{Q}_e^i, T_e^i) \in \Omega$,
 $T_{\text{food}}^{\text{final},i} = (T_{\text{food,min}} + T_{\text{food,max}}) / 2$,

Update

\dot{Q}_e^i , T_e^i , and $i = i + 1$

Repeat until convergence.

In Hovgaard et al. (2012) 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

$$\varphi_{\text{prox}} = \rho_{\text{prox}} \sum_{k=0}^{N-1} \|\dot{Q}_e^k - \dot{Q}_e^{k,\text{prev}}\|_2^2, \quad (11)$$

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 solution 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 one iteration of the sequential optimization, greater amounts of cooling capacity are applied to time steps where the efficiency of the system is high. Doing this causes the mass flow of refrigerant, the pressure difference over the compressor, or both to increase, and thereby lowers the efficiency. If this effect is sufficiently powerful, the COP calculated in the following iteration might be completely different and the optimization will try to reduce cooling at those time steps and the outcome will differ greatly from the previous. Proximal regularization eliminates this oscillatory behavior.

Finally, we add a quadratic penalty on the rate-of-change of \dot{Q}_e ,

$$\varphi_{\text{roc}} = \rho_{\text{roc}} \sum_{k=1}^{N-1} \|\dot{Q}_e^k - \dot{Q}_e^{k-1}\|_2^2. \quad (12)$$

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 et al. (2004), Jørgensen (2005), Wang and Boyd (2010).

3.4 Non-homogeneous sampling

To benefit from the variations in outdoor temperature and electricity prices we want to have an effective prediction horizon of at least 12 hours. Since the tail of the control sequence calculated in open loop is typically not identical the optimal closed-loop sequence, we choose a sufficiently long prediction and control horizon of 24 hours.

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 the 24-hour prediction horizon. One way of reducing the problem size is non-homogeneous sampling over the prediction horizon, exploiting that accuracy becomes less important towards the end of the open-loop sequence. Hence, we are using a prediction horizon augmented of three sequences with increasing sample time.

4 Case study

By simulation of realistic case studies we have verified the functionality and performance of the proposed MPC controller. In this section we describe the scenarios used and present the outcomes of the simulations.

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. The cooling capacity is controlled individually for each unit and we index these variables as $\{\dot{Q}_{e,i}\}_{i=1}^3$. The refrigeration system that we monitored uses CO₂ as refrigerant. CO₂ is getting increasingly popular for supermarket refrigeration since it is non-poisonous and non-flammable and since several governments put restrictions on the usage of conventional HFC refrigerants. We use calculations of the power consumption capable of handling

Table 1. Key parameters for the refrigeration system used in the case study scenario.

UNIT 1: MILK COOLER		
$m_{\text{food}}c_{p,\text{food}}$	550.0	kJ/K
$m_{\text{air}}c_{p,\text{air}}$	80.0	kJ/K
$k_{\text{amb-cr}}$	8.0	W/K
$k_{\text{food-air}}$	45.0	W/K
$k_{\text{evap,max}}$	135.0	W/K
$T_{\text{food,min}}$	1.0	°C
$T_{\text{food,max}}$	4.0	°C
UNIT 2: VERTICAL DISPLAY		
$m_{\text{food}}c_{p,\text{food}}$	395	kJ/K
$m_{\text{air}}c_{p,\text{air}}$	100.0	kJ/K
$k_{\text{amb-cr}}$	11.0	W/K
$k_{\text{food-air}}$	80.0	W/K
$k_{\text{evap,max}}$	170.0	W/K
$T_{\text{food,min}}$	2.0	°C
$T_{\text{food,max}}$	3.0	°C
UNIT 3: FROST ROOM		
$m_{\text{food}}c_{p,\text{food}}$	775	kJ/K
$m_{\text{air}}c_{p,\text{air}}$	50.0	kJ/K
$k_{\text{amb-cr}}$	2.3	W/K
$k_{\text{food-air}}$	19.0	W/K
$k_{\text{evap,max}}$	88.0	W/K
$T_{\text{food,min}}$	-22.0	°C
$T_{\text{food,max}}$	-18.0	°C
COMMON		
T_{amb}	20.0	°C
$T_{e,\text{min}}$	-12.0	°C
$T_{eF,\text{min}}$	-35.0	°C
Compressor heat loss (η_{heat})	15	%

both sub- and super-critical operation of the CO₂ system. Table 4.1 gives the key parameters for the system. In Hovgaard et al. (2012) we demonstrated how to estimate the parameters and design an observer for the food temperatures in the refrigeration system. We convert the system in §2.1 to the discrete-time equivalent using these parameters. 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.

The temperature in the frost room (which has the slowest dynamics) increases from $T_{\text{food,min}}$ to $T_{\text{food,max}}$ in approximately 11.5 hours if no cooling is applied. This supports the need for a horizon of at least 12 hours as mentioned in §3.

4.2 Algorithm details

We use a prediction horizon of 24 hours, with nonhomogeneous sampling. The first 6-hour interval is sampled every 15 minutes, followed by the second 6-hour interval sampled every 30 minutes, and the last 12-hour interval is sampled every hour. This gives us a total of 48 values to describe the 24-hour period.

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. With these values of the regularization parameters, the sequential convex programming method typically converged in 5 or so steps. We found that early termination, after only 2 steps, still resulted in quite good closed-loop control performance.

Recent advances in convex optimization allow for convex QPs to be solved at millisecond and microsecond time-scales. We use CVXGEN (Mattingley and Boyd 2012) 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 Temperatures and prices

As outdoor temperatures and electricity prices affect the efficiency and the cost, respectively, of operating the system, they are important factors in the MPC formulation. 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 calendar year and use three years of data for training the predictors.

4.4 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. Only past values of such parameters can be available to the controller and in the present work we incorporate 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.

In the literature predictors are suggested for different purposes and with different levels of complexity. In Galanis and Anadranistakis (2002) a Kalman filter approach is taken to correct temperature forecasts and in Leephakpreeda (2012) a grey prediction model is used for outdoor temperatures as well. Mohsenian-Rad and Leon-Garcia (2010) used a correlation-based analysis to find coefficients for a polynomial estimator of real-time electricity prices. In this work we use predictors that are simple to find from historical data and require extremely little computational effort in the real-time closed-loop implementation. Predictions of both electricity prices and outdoor temperatures are computed in the same manner which we describe here.

We use the historical training data set to construct typical days that describe the mean daily variation for each month in the year. If, *e.g.*, price is sampled every hour we get 24 data points for each one of the 12 months. We compute a smooth baseline covering all 365 days in a year using linear interpolation of two adjacent months.

For the entire historical data set we calculate the residual (difference between baseline and historical data) and compute a residual predictor by solving the convex optimization problem

$$\text{minimize } \sum_{k=1}^K \|[R_{k-n}, \dots, R_k]X - [R_{k+1}, \dots, R_{k+N}]\|_2^2 + \lambda \|X\|_1, \quad (13)$$

for X , 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. The ℓ_1 regularization on the predictor, with positive parameter λ , yields a sparse predictor matrix (Boyd and Vandenberghe 2004). By cross-validation with the test data set we choose λ to minimize the validation error.

Now, we can compute the predictions online in each time increment by first predicting the N future residuals from the n past residuals (n past measurements subtracted the baseline) and adding these to the baseline of the corresponding time window.

Algorithm 3 summarizes this procedure. After experimenting with the data, we chose to use two days of past data for predicting the outdoor temperature (residual) 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. Fig. 4 shows the mean absolute prediction error for outdoor temperatures and for electricity prices over the prediction horizon. The temperature data cover a range from -11°C to 30°C with

Algorithm 3 Computing predictors from historical data.

Off-line:

1. D = historical data set.
2. Compute typical day for each month by averaging over D .
3. Compute yearly baseline (b) by linear interpolation.
4. Compute residual $R = D - b$.
4. Compute X by (13).

On-line:

1. R_{past} = past measurements $- b$.
 2. Predict residual as: $R_{\text{past}}^T X$.
 3. Compute prediction as: predicted residual + baseline.
-




Figure 4. Mean absolute value of prediction errors for outdoor temperature and electricity spot price with training set covering 2007–2009 and test set covering 2010. Prediction horizon is 24 hours.

an average of 6.3°C , and the price data cover a range from $-20\text{EUR}/\text{MWh}$ to $100\text{ EUR}/\text{MWh}$, with an average of $46\text{ EUR}/\text{MWh}$. For the baselines the mean absolute errors are 2.5°C and $13.2\text{ EUR}/\text{MWh}$ for temperature and price respectively.

We show an example with baseline, predicted values, and real measurements for a randomly chosen point of time in Fig. 5. Fig. 6 shows histograms for the prediction errors of the outdoor temperatures at 1, 4, 12 and 24 hours into the future and Fig. 7 gives the same for electricity prices.

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

4.5 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.

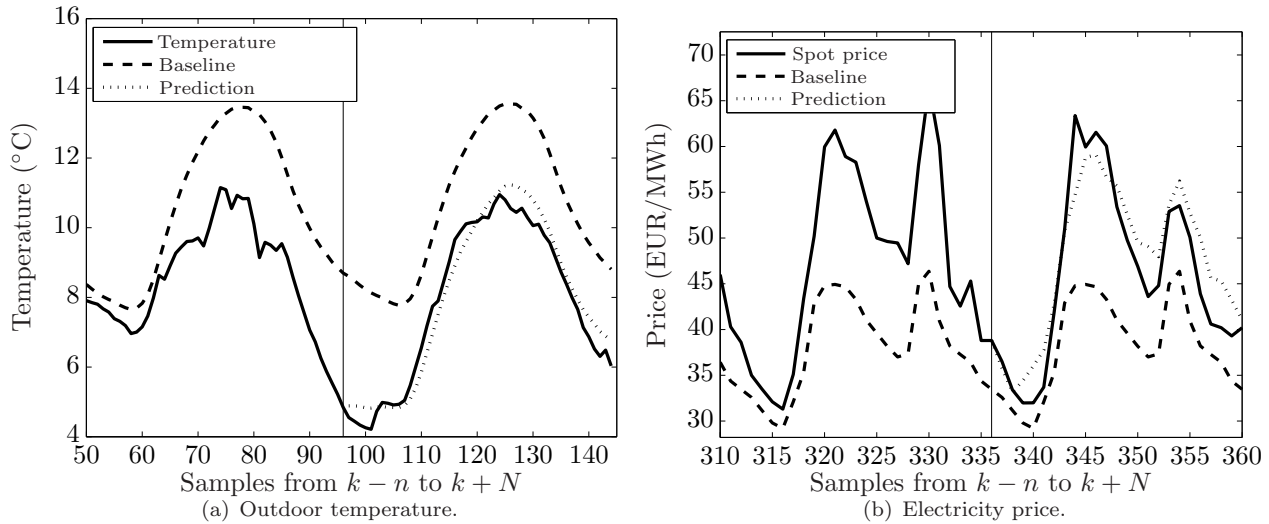


Figure 5. Measurements, baseline and prediction for a randomly chosen point of time, k . The vertical line indicates k and everything to the left of that point are past measurements used for prediction while the predictions are shown to the right of the line.




Figure 6. Prediction error histograms for outdoor temperature covering 2010.




Figure 7. Prediction error histograms for electricity price covering 2010.

4.6 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 by less than 1%.

4.7 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. 8 shows a segment of the simulated system with




Figure 8. Selected trajectory for food temperature and hourly cost of energy for the refrigeration system controlled by thermostat vs. the proposed MPC controller. In addition, the spot price is plotted for comparison (scaled to fit the range of the other variables).

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 compensated by the savings when the electricity prices go up.

In Fig. 9–10, resulting temperature distributions for selected units 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. A unit with larger thermal mass (Fig. 9) is utilized to a greater extent than a unit which has less storage capacity (Fig. 10).

We observe savings on the order of 40–50% for the simulations covering a full year (2010). However, a part of this comes from the ability to increase the evaporation temperature, and thereby the efficiency, significantly at times where there is almost no cooling demand. In an actual refrigeration system more units are expected and the chance of instances where all of them have an imperceptible cooling demand at the same time decreases. In addition, the most loaded unit might not even be able to participate with flexibility and will thus maintain its cooling demand at all times. A more realistic savings estimate is in the order of 30%.




Figure 9. Temperature distribution for selected unit. Simulation over the full year 2010.




Figure 10. Temperature distribution for selected unit. Simulation over the full year 2010.

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%.

Fig. 11 compares the cost-per-period distribution for the system controlled by thermostat and by MPC, respectively. In particular, we observe how a majority of the savings come from avoiding the most expensive instances, *e.g.*, above 0.006 EUR/period, when we use the MPC control policy.

4.8 Demand response

Fig. 12 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 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 vice versa, if we employ the proposed MPC scheme. A linear fit is made using a Huber function

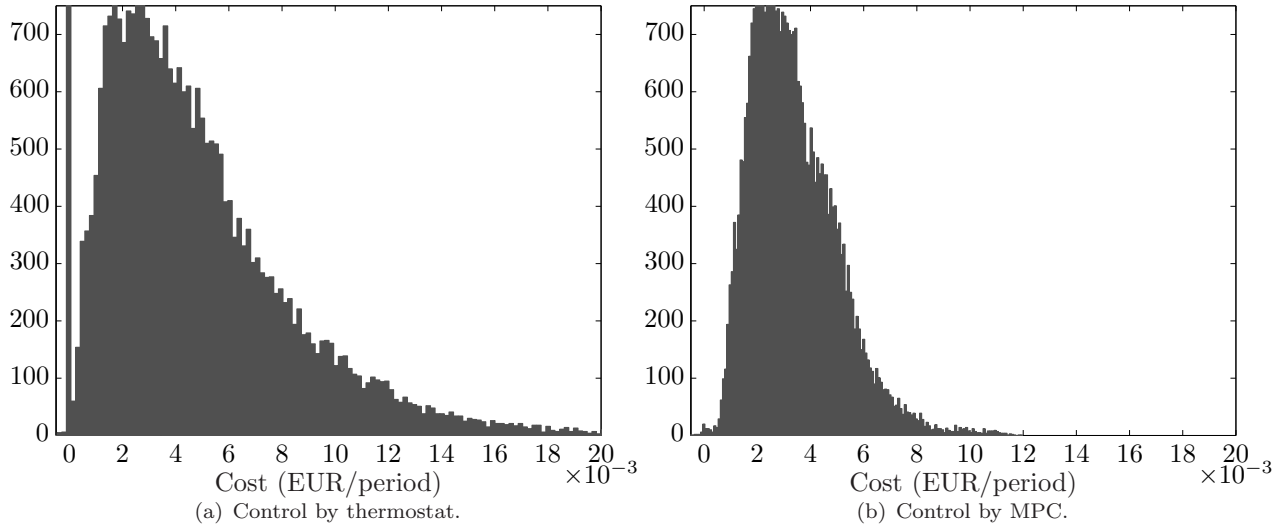



Figure 11. Distribution of cost-per-period. Simulation over the full year 2010.



(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 $-53 \text{ W}/(\text{EUR}/\text{MWh})$.

Figure 12. Illustration of demand response in systems controlled by MPC vs. thermostat control. The data-points are from simulation of the month of July 2010.

regression. The slope is around $-50 \text{ W}/(\text{EUR}/\text{MWh})$ 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.

4.9 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.

4.10 Plant perturbations

By re-running the simulations using the exact same controller but with reasonable perturbations in the plant parameters we observed that the proposed controller is quite robust. 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 and behaviors, like what we reported for the nominal system in Figures 9–12, appear.

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.

Acknowledgments

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

References

- Andersson, S.L., Elofsson, A.K., Galus, M.D., Göransson, L., Karlsson, S., Johnsson, F., and et al., (2010), “Plug-in hybrid electric vehicles as regulating power providers: Case studies of Sweden and Germany,” *Energy policy*, 38, 2751–2762.
- Angeli, D., Amrit, R., and Rawlings, J. (2011), “On average performance and stability of economic model predictive control,” *Automatic Control, IEEE Transactions on*, PP, 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, 371–389.
- Biegler, L. (2009), “Efficient Nonlinear Programming Algorithms for Chemical Process Control and Operations,” in *System Modeling and Optimization*, Vol. 312 of *IFIP Advances in Information and Communication Technology*, ed. Korytowski, A. and Malanowski, K. and Mitkowski, W. and Szymkat, M., Springer Boston, pp. 21–35.
- Blarke, M.B., and Dotzauer, E. (2011), “Intermittency-friendly and high-efficiency cogeneration: Operational optimisation of cogeneration with compression heat pump, flue gas heat recovery, and intermediate cold storage,” *Energy*, 36, 6867 – 6878.
- Boggs, P.T., and Tolle, J.W. (1995), “Sequential Quadratic Programming,” *Acta Numerica*, 4, 1–51.
- Boyd, S., and Vandenberghe, L., *Convex Optimization*, Cambridge University Press (2004).
- Bush, R., and Wolf, G. (2009), “Utilities Load Shift With Thermal Storage,” *Transmission & Distribution World*, p. 12.
- Camacho, E.F., Samad, T., Garcia-Sanz, M., and Hiskens, I. (2011), “Control for Renewable Energy and Smart Grids,” in *The Impact of Control Technology, Control Systems Society*, ed. T. Samad and A. M. Annaswamy, IEEE Control Systems Society, pp. 69–88.
- Diehl, M., Amrit, R., and Rawlings, J.B. (2011), “A Lyapunov Function for Economic Optimizing Model Predictive Control,” *Automatic Control, IEEE Transactions on*, 56, 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, 577–585.
- Diehl, M., Ferreau, H., and Haverbeke, N. (2009), “Efficient Numerical Methods for Nonlinear MPC and Moving Horizon Estimation,” in *Nonlinear Model Predictive Control*, Vol. 384 of *Lecture Notes in Control and Information Sciences*, ed. Magni, Lalo and Raimondo, Davide and Allgöwer, Frank, Springer Berlin / Heidelberg, pp. 391–417.
- Dinh, Q.T., Savorgnan, C., and Diehl, M. (2011), “Real-time sequential convex programming for nonlinear model predictive control and application to a hydro-power plant,” in *Proc. of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, pp. 5905–5910.
- Edlund, K., Bendtsen, J.D., and Jørgensen, J.B. (2011), “Hierarchical model-based predictive control of a power plant portfolio,” *Control Engineering Practice* 19, pp. 1126–1136.
- Elliott, M.S., and Rasmussen, B.P. (2008), “Model-based predictive control of a multi-evaporator vapor compression cooling cycle,” in *Proc. of the American Control Conference*, pp. 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.
- Finn, P., Fitzpatrick, C., Connolly, D., Leahy, M., and Relihan, L. (2011), “Facilitation of renewable electricity using price based appliance control in Irelands electricity market,” *Energy*, 36, 2952 – 2960.
- Galanis, G., and Anadranistakis, M. (2002), “A one-dimensional Kalman filter for the correction of near surface temperature forecasts,” *Meteorological Applications*, 9, 437–441.
- Grancharova, A., Johansen, T., and Tøndel, P. (2007), “Computational Aspects of Approximate Explicit Nonlinear Model Predictive Control,” in *Assessment and Future Directions of Nonlinear Model Predictive Control*, Vol. 358 of *Lecture Notes in Control and Information Sciences*, ed. Findeisen, R. and Allgöwer, F. and Biegler, L., Springer Berlin / Heidelberg, pp. 181–192.
- Halvgaard, R., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2012), “Economic Model Predictive Control for building climate control in a Smart Grid,” in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, pp. 1–6.
- Han, S., Han, S., and Sezaki, K. (2010), “Development of an Optimal Vehicle-to-Grid Aggregator for Frequency Regulation,” *IEEE Transactions on Smart Grid*, 1, 65–72.
- Hindi, H., Greene, D., and Laventall, C. (2012), “Coordinating Regulation and Demand Response in Electric Power Grids: Direct and Price-Based Tracking Using Multirate Economic Model Predictive Control,” in *Control and Optimization Methods for Electric Smart Grids*, Vol. 3 of *Power Electronics and Power Systems*, ed. Chakraborty, A. and Ilic, M. D., Springer US, pp. 111–131.
- 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., Edlund, K., and Jørgensen, J.B. (2010), “The Potential of Economic MPC for Power Management,” in *49th IEEE Conference on Decision and Control, 2010*, pp. 7533–7538.
- Hovgaard, T.G., Larsen, L.F.S., Edlund, K., and Jørgensen, J.B. (2012), “Model predictive control technologies for efficient and flexible power consumption in refrigeration systems,” *Energy*, 44, 105 – 116.
- Hovgaard, T.G., Larsen, L.F.S., and Jørgensen, J.B. (2011), “Flexible and Cost Efficient Power Consumption using Economic MPC - A Supermarket Refrigeration Benchmark,” in *Proc. of the 50th IEEE Conference on Decision and Control and European Control Conference*, pp. 848–854.
- 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*, pp. 1515–1520.
- Hovgaard, T.G., Larsen, L.F.S., Skovrup, M.J., and Jørgensen, J.B. (2012), “Analyzing Control Challenges for Thermal Energy Storage in Foodstuffs,” in *Proc. of the IEEE International Conference on Control Applications (CCA) part of the IEEE Multi-Conference on Systems and Control (MSC)*, pp. 956–961.
- Hovgaard, T.G., Larsen, L.F.S., Skovrup, M.J., and Jørgensen, J.B. (2012), “Optimal energy consumption in refrigeration systems - modelling and non-convex optimisation,” *The Canadian Journal of Chemical Engineering*, p. in press.
- Jerez, J.L., Constantinides, G.A., Kerrigan, E.C., and Ling, K.V. (2011), “Parallel MPC for real-time FPGA-based implementation,” in *Proc. of the 18th IFAC World Congress*, pp. 1338–1343.
- Jørgensen, J.B. (2005), “Moving Horizon Estimation and Control,” Department of Chemical Engineering, Technical University of Denmark.
- Jørgensen, J.B., Rawlings, J.B., and Jørgensen, S.B. (2004), “Numerical methods for large-scale moving horizon estimation and control,” in *Proc. of the Int. Symposium on Dynamics and Control Process Systems (DYCOPS)*.
- Kirschen, D.S. (2003), “Demand-side view of electricity markets,” *Power Systems, IEEE Transactions on*, 18, 520–527.
- Kraning, M., Wang, Y., Akuiyibo, E., and Boyd, S. (2011), “Operation and Configuration of a Storage Portfolio via Convex Optimization,” in *Proc. of the 18th IFAC World Congress*, pp. 10487–10492.
- Larsen, L.F.S., Thybo, C., and Rasmussen, H. (2007), “Potential Energy Savings Optimizing the Daily Operation of Refrigeration Systems,” *Proc. European Control Conference, Kos, Greece*, pp. 4759–4764.
- Leducq, D., Guilpart, J., and Trystram, G. (2006), “Non-linear predictive control of a vapour compression cycle,” *International Journal of Refrigeration*, 29, 761 – 772.
- Leephakpreeda, T. (2012), “Implementation of adaptive indoor comfort temperature control via embedded system for air-conditioning unit,” *Journal of Mechanical Science and Technology*, 26, 259–268.
- Ma, J., Qin, J., Salsbury, T., and Xu, P. (2012), “Demand reduction in building energy systems based on economic model predictive control,” *Chemical Engineering Science*, 67, 92–100.
- Ma, Y., and Borrelli, F. (2012), “Fast Stochastic Predictive Control for Building Temperature Regulation,” in *Proc. of the American Control Conference (ACC), 2012*, pp. 3075–3080.
- Ma, Y., Borrelli, F., Hencsey, B., Coffey, B., Bengesa, S., and Haves, P. (2012a), “Model Predictive Control for the Operation of Building Cooling Systems,” *IEEE Transactions on Control Systems Technology*, 20, 796–803.
- Ma, Y., Kelman, A., Daly, A., and Borrelli, F. (2012b), “Predictive Control for Energy Efficient Buildings with Thermal Storage: Modeling, Stimulation, and Experiments,” *IEEE Control Systems*, 32, 44–64.
- Mattingley, J., and Boyd, S. (2012), “CVXGEN: a code generator for embedded convex optimization,” *Optimization and Engineering*, 13, 1–27.
- Mattingley, J., W., Y., and Boyd, S. (2011), “Receding Horizon Control,” *IEEE Control Systems Magazine*, 31, 52–65.
- Mohsenian-Rad, A.H., and Leon-Garcia, A. (2010), “Optimal Residential Load Control With Price Prediction in Real-Time Electricity Pricing Environments,” *IEEE Transactions on Smart Grid*, 1, 120–133.
- Molina-Garcia, A., Kessler, M., Fuentes, J.A., and Gomez-Lazaro, E. (2011), “Probabilistic Characterization of Thermostatically Controlled Loads to Model the Impact of Demand Response Programs,” *IEEE Transactions on Power Systems*, 26, 241–251.
- O’Donoghue, B., Stathopoulos, G., and Boyd, S., “A Splitting Method for Optimal Control,” http://www.stanford.edu/~boyd/papers/oper_splt_ctrl.html (2012).

- Oldewurtel, F., Parisio, A., Jones, C.N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., and Morari, M. (2012), “Use of model predictive control and weather forecasts for energy efficient building climate control,” *Energy and Buildings*, 45, 15–27.
- Oldewurtel, F., Parisio, A., Jones, C.N., Morari, M., Gyalistras, D., Gwerder, M., and et al., (2010), “Energy Efficient Building Climate Control using Stochastic Model Predictive Control and Weather Predictions,” in *Proc. of American Control Conference 2010*, pp. 5100–5105.
- 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 *Proc of the American Control Conference (ACC)*, pp. 5100–5105.
- Pina, A., Silva, C., and Ferrão, P. (2012), “The impact of demand side management strategies in the penetration of renewable electricity,” *Energy*, 41, 128 – 137.
- Qin, S.J., and Badgwell, T.A. (2003), “A survey of industrial model predictive control technology,” *Control engineering practice*, 11, 733–764.
- Rawlings, J.B., and Amrit, R. (2009), “Optimizing Process Economic Performance Using Model Predictive Control,” *Nonlinear Model Predictive Control: Towards New Challenging Applications*, pp. 119–138.
- Saele, H., and Grande, O.S. (2011), “Demand Response From Household Customers: Experiences From a Pilot Study in Norway,” *IEEE Transactions on Smart Grid*, 2, 90–97.
- Sarabia, D., Capraro, F., Larsen, L.F.S., and de Prada, C. (2008), “Hybrid NMPC of supermarket display cases,” *Control Engineering Practice*, 17, 428–441.
- 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*, pp. 13901–13906.
- Van Harmelen, G.L. (2001), “The virtual power station targeting residential, industrial and commercial controllable loads,” *IFAC Conference on Technology Transfer in Developing Countries - Automation in Infrastructure Creation (DECOM-TT 2000) Proceedings volume from IFAC Conference*, pp. 45–48.
- Wang, Y., and Boyd, S. (2010), “Fast Model Predictive Control Using Online Optimization,” *IEEE Transactions on Control Systems Technology*, 18, 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 the 47th IEEE Conference on Decision and Control (CDC)*, pp. 4718–4723.