Data-driven Online Adaptation of the Scenario-tree in Multistage Model Predictive Control ⋆

Mandar Thombre, Dinesh Krishnamoorthy, Johannes Jäschke

Department of Chemical Engineering, Norwegian University of Science & Technology , NO-7491 Trondheim , (e-mail: mandar.thombre@ntnu.no, dinesh.krishnamoorthy@ntnu.no, johannes.jaschke@ntnu.no).

Abstract: Multistage model predictive control (MPC) is based on the enumeration of scenarios that represent the uncertainty in the system. Scenario selection is important in multistage MPC since the choice of scenarios determines the degree of conservativeness of the optimal solution. We propose a data-driven approach based on principal component analysis (PCA) to dynamically select the scenarios, leading to reduced conservativeness. When time-varying uncertainty is considered, PCA can be performed online to select new scenarios whenever the uncertainty data is updated. The results of the approach are demonstrated for a two-plant system with a thermal storage tank. The solution obtained is less conservative than with standard multistage MPC. This is because the online PCA-based approach accounts for the most recent, and thus more representative, uncertainty realizations.

Keywords: Multistage MPC, online PCA, scenario selection, parametric uncertainty

1. INTRODUCTION

A lot of energy in chemical process plants is lost in the form of industrial waste heat. In industrial clusters of multiple process plants, energy resources such as steam, cooling water, raw materials, etc. are often shared. Optimal energy efficiency in such a system demands flexible operation of the plants, so that surplus energy from one plant is transferred to another plant in need of it. The plants in such clusters can act both as sources of heat (those having surplus heat) and sinks of heat (those having energy demands). Examples of industrial clusters in Scandinavia include Mo Industripark in Norway and Kalundborg Symbiosis in Denmark. The surplus heat streams from different sources are available at varying temperatures. Moreover, supply of surplus heat and energy demand is often asynchronous, while also operating at differing time scales. To mitigate some of these issues, employing a thermal energy storage system is an attractive option. Such a system creates a buffer between the energy suppliers and users. For example, domestic thermal energy storage in buildings was considered by de Oliveira et al. (2016).

Process plants often exhibit highly nonlinear dynamics. In addition, they also have to contend with disturbances in process parameters like temperatures and flow rates during operation. In the context of energy exchange, this consequently results in an uncertain yield/consumption of surplus heat on the supply/demand side. The existence of such uncertainties presents a significant challenge for operating these plants at an optimal point.

Model predictive control (MPC) is a powerful tool for operational optimization that is widely used in the process industry. However, nominal MPC does not explicitly model the uncertainties in the system. To rigorously account for uncertainty, methods classified under "Robust MPC" have been receiving some attention in control literature. Based on the theory of robust optimization, the so called min-max MPC approach (first proposed by Campo and Morari (1987)) minimizes the cost of the worst-case realization of the uncertainty. The notion of feedback in min-max MPC was first introduced by Scokaert and Mayne (1998). In this approach, the closed-loop optimization is sought over different sequences of control inputs for different realizations of the uncertainty. This idea was further extended by Lucia et al. (2013) for multistage nonlinear MPC, based on the concepts of stochastic programming. Here, the evolution of uncertainty is assumed to be modeled by a scenario tree, which is generated from the discrete realizations of the uncertainty. By optimizing over each branch of this scenario tree, the idea that new information will be available in future stages is explicitly accounted for.

The convention in multistage MPC is to assume that the uncertainty information is known *a-priori*, that it is known perfectly and that it is discrete. However, the selection of scenarios that build the scenario tree is very important in the practical implementation, and has not been sufficiently addressed in the context of multistage MPC. Traditionally, scenario-based stochastic programming methods involve a two step process: estimating a probability distribution function (PDF) from a finite data-set, and subsequently discretizing the PDF to generate scenarios (Birge and Louveaux, 2011). Another approach is to skip the PDF

⋆ M.T and J.J gratefully acknowledge the financial support from FME HighEFF. D.K and J.J gratefully acknowledge the financial support from SFI SUBPRO. Corresponding author: J.J.

estimation step and go directly from data to scenarios, i.e. the discrete scenarios can directly be chosen from the available data samples. After all, the different data samples represent the discrete measurements of uncertainty in the system. Ideally, then, any set of selected scenarios should be a subset of this data set for the best representation of uncertainty.

Having decided on selecting scenarios directly from available data, the next question to consider is *which* data samples to select as scenarios. The size of the multistage MPC problem increases exponentially with increasing number of scenarios. Hence it is important to capture maximum uncertainty information with minimum number of scenarios, in order to be computationally efficient.

Uncertain parameters in a system often exhibit correlations. Sampling methods like the Monte Carlo or the Latin hypercube sampling emphasize randomness of sampling to maximize information, but ignore correlations. Therefore, these scenario selection methods may not be the best if we want to exploit the correlation to reduce the number of scenarios. To overcome this, multivariate data-analysis methods like the principal component analysis (PCA) can be used to detect any hidden correlations within the available data samples. The scenarios chosen using these multivariate methods explicitly take into account the interdependence between the parameters. Dimensionality reduction methods such as PCA explain the parametric variation in a data set in fewer dimensions - referred to as principal components. This means that lesser number of scenarios are able to effectively describe the parametric variation in the system, leading to a compact scenario tree formulation.

In this paper, we propose an online PCA-based approach for systems with time-varying uncertainty, that can be performed dynamically in the multistage MPC implementation. The idea is to select new scenarios online whenever new uncertainty data becomes available, and to systematically reformulate the optimization problem in anticipation of a predicted change in uncertainty. This can provide an additional hedge against uncertainty, since the "latest" data is constantly being used to select the scenarios. Moreover, we propose that if the parametric variation can be explained by a small number of "dominant" principal components (shown by PCA), it suffices to select the scenarios only along these components to sufficiently explain the uncertainty. The proposed approach is applied to a simple thermal energy storage model.

The paper is structured as follows. Section 2 describes the formulation of the multistage MPC problem and the dynamic scenario selection strategy using PCA. Section 5 illustrates the modeling for energy storage used as a case study for the demonstration of the multistage MPC. Simulation results are presented in Section 6 and the conclusions and recommendations are stated in Section 7.

2. PRELIMINARIES - MULTISTAGE MPC

Consider sets $\mathcal{S} = \{1, \ldots, S\}$ and $\mathcal{J} = \{0, \ldots, N-1\}$, where S is the total number of scenarios and N is the length of the prediction horizon. The scenario-based multistage MPC problem can be formulated as follows:

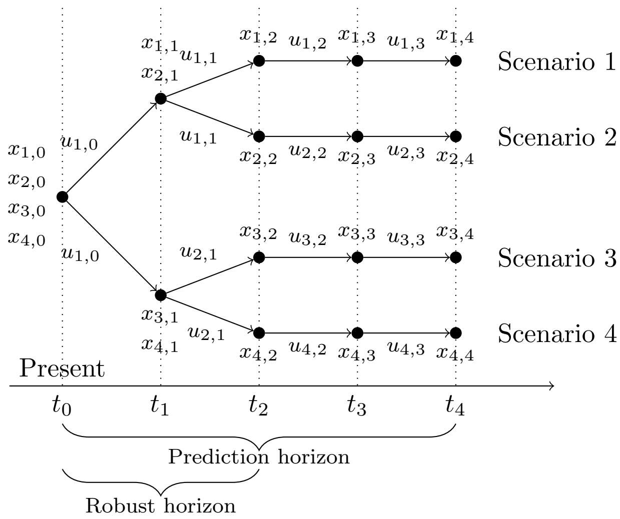

Fig. 1. Illustration of the scenario tree at time t_0 with $M = 2$ realizations of uncertainty. The prediction horizon $N = 4$ and the robust horizon $N_r = 2$.

$$\min_{\mathbf{x}_{i,j},\mathbf{u}_{i,j}} \quad \sum_{i\in\mathcal{S}} \omega_i \sum_{j\in\mathcal{J}} \phi(\mathbf{x}_{i,j}, \mathbf{u}_{i,j}) \tag{1a}$$

s.t.

$$\mathbf{x}_{i,0} = \mathbf{x}^{init}, \qquad \forall i \in \mathcal{S} \tag{1b}$$

$$\mathbf{x}_{i,j+1} = \mathbf{f}(\mathbf{x}_{i,j}, \mathbf{u}_{i,j}, \boldsymbol{\pi}_i), \qquad \forall j \in \mathcal{J}, \forall i \in \mathcal{S} \tag{1c}$$

$$\mathbf{g}(\mathbf{x}_{i,j}, \mathbf{u}_{i,j}) \leq 0, \qquad \forall j \in \mathcal{J}, \forall i \in \mathcal{S} \tag{1d}$$

$$\sum_{i\in\mathcal{S}} \mathbf{E}_i \mathbf{u}_i = 0, \qquad \forall j \in \mathcal{J}, \forall i \in \mathcal{S} \tag{1e}$$

The states and inputs for the ith scenario and jth time step are denoted by $\mathbf{x}_{i,j} \in \mathbb{R}^{n_x}$ and $\mathbf{u}_{i,j} \in \mathbb{R}^{n_u}$ respectively. The uncertain parameters for the ith scenario are denoted by $\pi_i \in \mathbb{R}^{n_\pi}$. Equation (1a) represents the cost to minimize, where ω_i is the weight assigned to the ith scenario and $\phi : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}$ is the objective function. Equation (1b) represents the initial condition constraints, with $\mathbf{x}^{init}$ being the vector of starting points for the states. Equation (1c) represents the model of the nonlinear dynamic system described by the vector of state equations $\mathbf{f} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_\pi} \to \mathbb{R}^{n_x}$, and Equation (1d) represents the constraints in the system, denoted by $\mathbf{g} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}$.

Equation (1e) represents the non-anticipativity constraints, which enforce that all control inputs applicable to branches of the same node, are equal. This is because, in real applications, the uncertainty is realized *after* the control input is applied. In other words, one cannot anticipate how the scenario tree is going to branch out at a node before a decision is taken at that node. The reader is referred to Krishnamoorthy et al. (2018a) for details on the construction of the non-anticipativity matrix $\mathbf{E_i}$.

To curb the rapid growth of the scenario tree, the uncertainty is assumed to be constant after a certain point in the prediction horizon so as to reduce the computational cost, as justified by Lucia and Engell (2013). This point represents the so-called robust horizon of the problem, with a length N_r. The scenario tree evolution showing the prediction horizon and the robust horizon is illustrated in Fig. 1. Therefore if we have M discrete realizations of the uncertainty, then this results in $S = M^{N_r}$ scenarios.

3. DYNAMICALLY ADJUSTING THE SCENARIO-TREE

Many systems have to contend with uncertainty that is time-varying. In this paper, we propose to update the scenario-tree dynamically to acknowledge this time-varying uncertainty. Our approach takes into account that the considered scenarios may change during the operation of the system. That is, the uncertainty may have very different characteristics during different points in time.

For example, consider the intra-day variation of energy demand in a district heating network, with high peaks in the morning and the afternoon. Clearly, the scenarios depicting the demand during the peak hours will be different from the scenarios during low and medium demand periods.

To reflect these changes, the scenario-tree can be updated during operation as new uncertainty data becomes available. We propose two update strategies for the scenario-tree:

(1) Adjust the scenarios, i.e the parameter values, at every time step as new information becomes available.
(2) Extend the length of the robust horizon at the current time step if a change in uncertainty information at a future time step can be anticipated.

The first strategy recognizes that multistage MPC is performed on a moving horizon, where the optimization problem is repeatedly solved at every time step with an updated initial condition. We propose, in addition, to also update the scenarios themselves at every time step in recognition of newly available uncertainty information.

For instance, assume that at time t_0, the scenario-tree is as given in Fig. 1. This tree is used in the multistage-MPC controller until updated information is available. If at time t_2 new information about the uncertain parameters becomes available, the scenario-tree for the optimization problem will be updated using the new data. This updated scenario-tree will be used in the multistage MPC implementation from time t_2 until newer information about the scenarios becomes available.

With respect to the second strategy, if it is known a-priori that uncertainty data will be updated at a future time step in the horizon, the robust horizon can be accordingly modified to take this into account. Farther way from the point of update, a shorter robust horizon can be used to reduce computational burden. As the predicted point of update comes closer, the robust horizon can be extended to include the new scenarios reflecting the update.

Consider again the scenario tree shown in Fig. 1. Here $N_r = 2$ at time t_0. However, if it known at time t_0 that new uncertainty information is available at at time t_2, then the robust horizon can be increased from $N_r = 2$ to $N_r = 3$ at time t_0, in order to accommodate this extra branching at time t_2. This would thus lead to consideration of additional scenarios (in this case, $2^3 = 8$ scenarios).

These updates of scenarios and the robust horizon can be done dynamically within multistage MPC. The procedures can be performed via online computations at any time step which presents new information about the uncertainty.

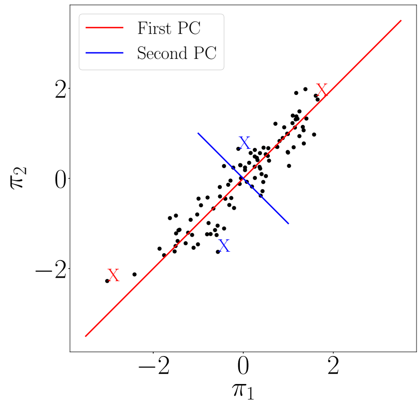

Fig. 2. Selected scenarios, marked by 'X', correspond to the maximum and minimum scores along the two principal components.

4. DATA-DRIVEN SCENARIO SELECTION

Given a data set representing the uncertain parameters, the naive approach of selecting scenarios would be to take the combinations of the minimum and maximum values of each parameter to maximize uncertainty information. Similar to random sampling methods, this approach ignores parameter correlation. PCA can be leveraged to judiciously choose those samples from the set that incorporate information about correlations between the parameters. This is especially true for large data sets with many parameters where detecting correlations between parameters, if they exist, is impractical via simple tools such as univariate analysis.

PCA employs a mathematical procedure that transforms a data set with multiple, possibly correlated, variables into a lesser number of uncorrelated variables, known as principal components. Essentially, it is an orthogonal linear transformation of the data set into a new coordinate system with each new axis representing a principal component. The first principal component points in the direction of maximum variance within the data set. Subsequent principal components account for as much of the remaining variance as possible, in decreasing order. This dimensionality reduction helps explain the parametric variation in the data using smaller number of components.

Consider a data matrix $\mathbf{X} \in \mathbb{R}^{n_o} \times \mathbb{R}^{n_\pi}$, where rows of the matrix represent observations and columns represent the (possibly correlated) parameters. To remove arbitrary biases from the measurements, the data is mean-centered and scaled, resulting in the data matrix $\mathbf{X}_{sc} \in \mathbb{R}^{n_o} \times \mathbb{R}^{n_\pi}$. Performing PCA on $\mathbf{X}_{sc}$ results in the output $\mathbf{Y} \in \mathbb{R}^{n_o} \times \mathbb{R}^{n_{\pi'}}$, with $\pi' \leq \pi$, according to:

$$\mathbf{Y} = \mathbf{X}_{sc}\mathbf{P} \tag{2}$$

where, $\mathbf{P} \in \mathbb{R}^{n_\pi} \times \mathbb{R}^{n_{\pi'}}$ is the projection matrix with each column representing a principal component. In other words, each column of $\mathbf{P}$ contains the coefficients that project the original data point to the new coordinate system $(\mathbf{Y})$ of π' principal components. These are also referred to as *loadings*. The matrix $\mathbf{Y}$ is called the *scores* matrix. The score of a data point along a principal component represents the distance of that data point from the mean along the direction of that principal component. For an overview of PCA and its algorithms, the reader is referred to Jolliffe and Cadima (2016).

Scenarios can be chosen by leveraging information from this transformed data set. Since the principal components are orthogonal to each other, scenarios can be chosen along the direction of these principal components to obtain maximum uncertainty information, as demonstrated in our previous paper (Krishnamoorthy et al., 2018b). For example, the chosen scenarios could correspond to the maximum and minimum scores along the dominant principal components that explain the variations in the data sufficiently, as shown in Fig. 2.

The PCA may result in principal components such that some components dominate over the others, in terms of how much data variability they explain. We propose to select scenarios only along these dominant principal components, since the chosen scenarios can then account for maximum variation in the uncertainty.

For instance, Fig. 2 shows scenarios (marked by red and blue "X"s) selected along both the principal components. Instead, since it can be seen clearly that the first principal component is dominant [1], scenarios can be chosen only in that direction (marked by red "X"s). Thus, instead of choosing 4 scenarios, only 2 scenarios can encompass most of the parametric variation in the data shown in Fig. 2, without any significant loss in explained variability. Reducing the number of scenarios in this manner can thus make the size of resulting multistage MPC problem significantly smaller, reducing the computational effort.

5. CASE STUDY: SIMPLE ENERGY STORAGE SYSTEM

We consider a simple two-plant thermal storage system, with one plant being the supplier of heat (plant A) and the other being the consumer (plant B). A thermal storage tank acts as a buffer between the two plants to facilitate the energy exchange. The tank interacts with the two plants via heat exchangers, as shown in Fig. 3.

Further, the tank can directly be heated up via a local heating source. For example, in the context of industrial clusters, surplus heat from the flue gases that result from various chemical processes can be used to heat up the tank. Thus, the local heating source is considered inexpensive. If the energy in the tank is insufficient to meet the energy requirements on the demand side, the plant has to purchase the excess energy from the market. Energy from the market is usually much more expensive than the local heating source. The objective is to operate the system such that the total cost of energy purchase is minimized.

5.1 Process model

The heat exchangers are modeled as devices with two chambers representing the hot side and the cold side. Both chambers of the heat exchanger, as well as the tank itself, are considered to have the same temperature throughout their volumes. Thus the temperatures exiting

[1] For the data shown in Fig. 2, the first principal component explains 96.4% of the variance.

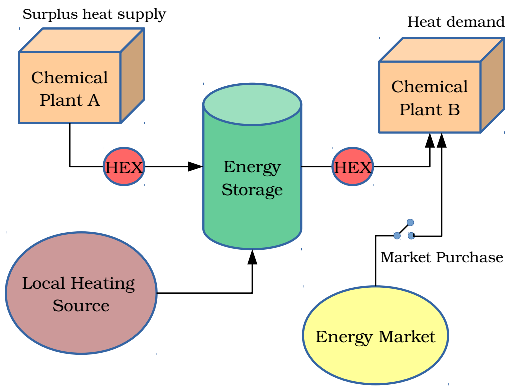

Fig. 3. Illustration of a simple energy storage system.

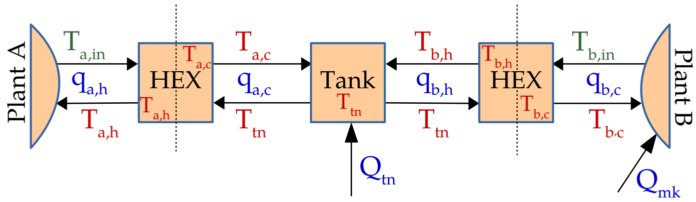

Fig. 4. Schematic of the model. The states, inputs and disturbances are shown in red,blue, and green respectively.

Table 1. Model parameters

Symbol	Value	Symbol	Value
V_{tn}	100 m^3	V	0.5 m^3
A_{tn}	100 m^2	A	300 m^3
U_{tn}	0.5 kW/m^2K	U	0.5 kW/m^2K
$(q_{ah})_{max}$	1 m^3/s	$(q_{ah})_{min}$	0 m^3/s
$(q_{ac})_{max}$	1 m^3/s	$(q_{ac})_{min}$	0 m^3/s
$(q_{bh})_{max}$	1 m^3/s	$(q_{bh})_{min}$	0 m^3/s
$(q_{bc})_{max}$	1 m^3/s	$(q_{bc})_{min}$	0 m^3/s
$(T_{tn})_{max}$	100 $°C$	$(T_{tn})_{min}$	30 $°C$
$(Q_{tn})_{max}$	5000 kW	$(Q_{tn})_{min}$	0 kW
T_{amb}	20 $°C$	Q_{demand}	5000 or 10000 kW
P_{tn}	5 units/kW	P_{mk}	1000 units/kW
ρ	1000 kg/m^3	C_p	4.18 kJ/kgK

these volumes are considered to be same as those inside the volumes. The driving force for the heat exchange between the two chambers is the difference in the temperatures of the two chambers. Further, we consider hot water as the fluid for heat exchange.

Both heat exchangers have an area A, volume V and heat transfer coefficient U. The tank has a volume V_{tn}, a surface area of A_{tn}, and experiences heat loss with a heat transfer coefficient U_{tn}. The density and specific heat capacity of water are denoted by ρ and C_p respectively. Temperatures on the hot and cold sides of heat exchangers on both sides (plant A and plant B) are $T_{a,h}$, $T_{a,c}$, $T_{b,h}$ and $T_{b,c}$ respectively. The tank temperature is T_{tn}, whereas the ambient temperature is T_{amb}. Inlet temperatures from plant A and plant B are $T_{a,in}$ and $T_{b,in}$ respectively The flow rates on either side of the heat exchangers are $q_{a,h}$, $q_{a,c}$, $q_{b,h}$ and $q_{b,c}$ respectively. Local heat supply is denoted by Q_{tn} and the energy purchased from the market is denoted by Q_{mk}. The model parameters are shown in Table 1.

Fig. 4 shows the schematic with the different states, inputs and disturbances in the model. The exit temperatures from the chambers of heat exchangers and the tank are the states $\mathbf{x}$. Flow rates on either sides of the heat exchanger, along with the local and market heat supply, are the inputs $\mathbf{u}$. The inlet temperatures from the two plants are the disturbances. The energy balances on the hot and cold chambers of the heat exchangers, and on the tank, become:

$$\frac{dT_{a,h}}{dt} = \frac{1}{V}\left\{q_{a,h}(T_{a,in} - T_{a,h}) - \frac{UA(T_{a,h} - T_{a,c})}{\rho C_p}\right\} \tag{3a}$$

$$\frac{dT_{a,c}}{dt} = \frac{1}{V}\left\{q_{a,c}(T_{tn} - T_{a,c}) + \frac{UA(T_{a,h} - T_{a,c})}{\rho C_p}\right\} \tag{3b}$$

$$\frac{dT_{b,h}}{dt} = \frac{1}{V}\left\{q_{b,h}(T_{tn} - T_{b,h}) - \frac{UA(T_{b,h} - T_{b,c})}{\rho C_p}\right\} \tag{3c}$$

$$\frac{dT_{b,c}}{dt} = \frac{1}{V}\left\{q_{b,c}(T_{b,in} - T_{b,c}) + \frac{UA(T_{b,h} - T_{b,c})}{\rho C_p}\right\} \tag{3d}$$

$$\frac{dT_{tn}}{dt} = \frac{1}{V_{tn}}\Bigg\{q_{a,c}(T_{a,c} - T_{tn}) - q_{b,h}(T_{tn} - T_{b,h}) + \frac{Q_{tn}}{\rho C_p} - \frac{U_{tn}A_{tn}(T_{tn} - T_{amb})}{\rho C_p}\Bigg\} \tag{3e}$$

5.2 Uncertainty description

We consider uncertainties in the form of disturbances in the system, $T_{a,in}$ and $T_{b,in}$, the inlet temperatures from the two plants. The temperature data from various streams in a process is usually logged. A period of 24 hours is considered for this uncertainty data. Further, we consider that the inlet temperature distributions are different for three distinct phases of the day (12 AM to 8 AM, 8 AM to 4 PM and 4 PM to 12 AM). We assume that historic data is available for the temperature distributions for these phases.

The plants operate at higher temperatures during the day phase and lower temperatures during the evening and night phases. Further, it is reasonable to expect that these plant temperatures are correlated to each other. This is because the periods of high and low activity in plants in an industrial cluster are similar. The scatter plot of the two inlet temperatures is shown in Fig. 5, for a series of historic data.

5.3 Formulating the multistage MPC problem

For a given energy demand profile for plant B, the economic objective is to minimize the cost of energy. In context of the multistage MPC formulation (1), the cost function for ith scenario and jth time step can then be stated as:

$$P_{tn}(Q_{tn})_{i,j} + P_{mk}(Q_{mk})_{i,j}$$

where P_{tn} and P_{mk} are the prices of the local energy supply and the market supply respectively.

The starting values of all temperatures are imposed as the constraints (1b). The model equations (1c) of the multistage MPC formulation are obtained by discretizing

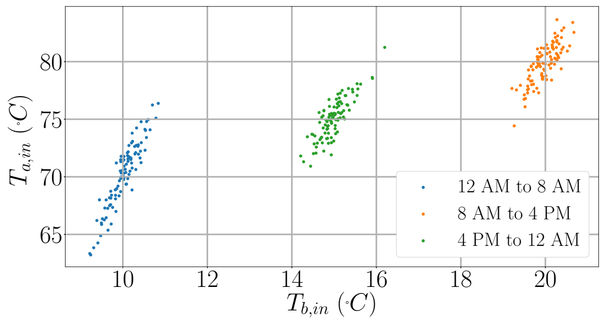

Fig. 5. Uncertainty data - inlet temperatures from the plants for different phases of the day. It can be seen that the temperatures are correlated differently throughout each phase of the day.

the energy storage model equations (3) using collocation on finite elements. These form the equality constraints of the problem. Bounds on the temperatures, flow rates, and energy supplies are imposed for each scenario and time step, and these form the inequality constraints (1d). In addition, the non-anticipativity constraints (1e) are also imposed.

Recall that plant B purchases energy from the market to satisfy the demands in excess of what can be satisfied solely through the storage system. For the ith scenario and jth time step, this can be formulated as the constraint:

$$(Q_{mk})_{i,j} + \rho C_p q_{b,c}((T_{b,c})_{i,j} - (T_{b,in})_{i,j}) \geq Q_{demand}$$

where Q_{demand} is the total energy demand of plant B. This imposes the condition that the energy purchased from the market in addition to the energy from the storage system must at least be equal to energy demand from the plant.

6. RESULTS

The multistage MPC problem is formulated with $N = 24$ hours (finite elements) and $N_r = 1$ hour, with control action changing every hour. `JuMP` (version 0.18.2) (Dunning et al., 2017), a modeling tool within the framework of `Julia` (version 0.6.2) (Bezanson et al., 2017) programming language, is used to implement the multistage MPC problem. The resulting nonlinear optimization problem is solved using `Ipopt` (Wächter and Biegler, 2006). The results are divided into the following two parts:

(1) Comparison of a scenario selection using a data-driven PCA and a conventional "BOX" approach, with a constant $N_r = 1$.
(2) Studying the effect of dynamically adjusting N_r closer to an anticipated change in the uncertainty data, while using PCA for scenario selection.

6.1 Data-driven vs conventional scenario selection

For comparison, the dynamic scenario selection is done with two methods. In the first method, a conservative approach is used, selecting scenarios as the four corner points from the box that encompasses all the uncertainty data over each 8-hour phase during the day. Essentially, these scenarios represent the combinations of the minimum and maximum of the data set along each dimension. A

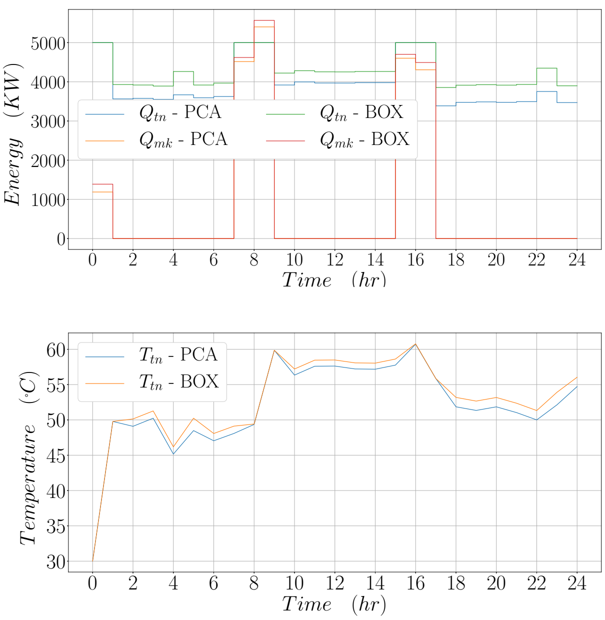

Fig. 6. The energy supplies and the tank temperature across the 24-hour period, for the two methods.

fifth scenario is chosen to represent the mean value along each dimension. This approach is referred to as the "BOX" method. The BOX method does not account for any correlations between the uncertain inlet temperatures.

In the second, "PCA" method, the scenarios are chosen by performing PCA dynamically over the data sets that are relevant for the corresponding phases of the day (shown by different colors in Fig. ??).The PCA results in two principal components for each data set, with the first principal component explaining 96.51% of the total variability for the 12 AM to 8 AM data, 89.72% for the 8 AM to 4 PM data, and 91.84% for the 4 PM to 12 AM data. Since the first principal components explain a large fraction of the variance in the data, two scenarios are selected corresponding to the minimum and maximum scores along this principal component. A third scenario is chosen to represent the mean value along each dimension.

We consider the uncertainty to be time-varying over the MPC horizon, where the "true" realization of the uncertainty in the simulator is chosen randomly for each hour from the corresponding data set. The simulation is considered for the 24 hour horizon from 12 AM to 12 AM. The demand is constant at 5000 kW throughout the day, except for 7 AM to 9AM and 3 PM to 5 PM, when there is a peak demand of 10,000 kW.

The results of the optimization are shown in Fig. 6. It can be clearly seen that for the BOX method, the solution is more conservative. The tank is heated to a higher temperature for satisfying the same demand profile across the day. Similarly, the heat supplied to the tank is also more compared to the PCA method. Also, the data-driven approach leads to lower purchases of the expensive energy from the market during peak demands.

Moreover, the simulations were repeated 30 times for the time-varying uncertainty case. The "true"" set of uncertain parameters in each simulation was a randomly chosen subset of the available data set. The performance for each simulation run was evaluated based on the integrated objective function, which sums up the values of the objective cost for all stages and scenarios. The results are presented in Fig. 7, where it can be seen that the PCA method outperforms the BOX method with a lower cost. Note that the integrated objective costs are divided by the respective number of scenarios chosen for each method for a fairer comparison.

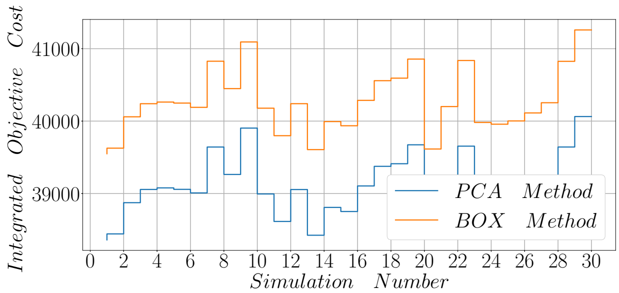

Fig. 7. The averaged integrated cost for 30 different simulation runs.

6.2 Dynamically adjusting the robust horizon

Here, the simulations were run such that the multistage MPC was implemented using two N_r cases. In the first case, the robust horizon was kept constant throughout the simulation at $N_r = 1$. This is referred to as the "NR1" case.

In the second case, the robust horizon was dynamically adjusted (switched) from $N_r = 1$ to $N_r = 2$, *one hour before* the night-to-morning and morning-to-evening phase changes. This change was implemented only for the corresponding next one time step, and subsequently reduced back to $N_r = 1$ for later time steps. Thus, $N_r = 2$ was used for the MPC time steps at 7 AM and 3 PM. This is because it is known that, at these times, the uncertainty data will be updated in one hour due to change in phase; and the robust horizon length of 2 hours reflects this. This is denoted as the "NR2" case.

The scenario selection was done via PCA for both cases. Moreover, the energy demand was considered to be constant at 5000 kW throughout the 24-hour period (i.e. no peak heating). The results are shown in Fig. 8

It can be seen that by dynamically extending the robust horizon (NR2 case), the optimization anticipates the upcoming rise in inlet temperatures by pre-empting the local tank heating at 7 AM (shown by a higher Q_{tn} in the NR2 case than in the NR1 case). This can also be seen from the tank temperature profile, where the temperature of the tank rises higher in the NR2 case at 7 AM than in the NR1 case. Consequently, the market purchase at 8 AM is smaller in the NR2 case, leading to lower cost. During the 4 PM phase change, the temperatures are dropping anyway so market purchase is unnecessary. This leads to the same temperature and heating profiles in the evening phase for both NR1 and NR2 cases. This is because, at this time, the tank has enough energy to satisfy the energy demand of Plant B.

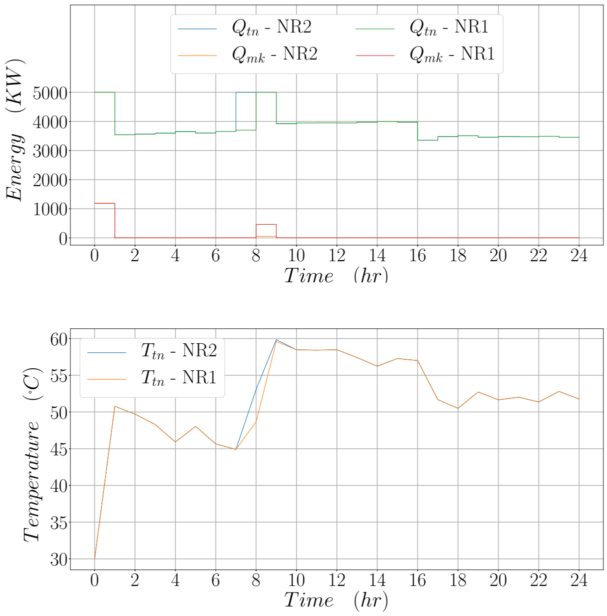

Fig. 8. The energy supplies and the tank temperature across the 24-hour period, for $N_r = 1$ and $N_r = 2$.

7. CONCLUSION AND FURTHER WORK

The case study of the thermal storage tank demonstrates that the same energy demand profile can be satisfied by heating the tank *less* if the scenario selection is data-driven and dynamic. Not only does the tank operate at a lower temperature, but the cost of operation is also significantly lower.

In addition, extending the robust horizon dynamically leads to the consideration of future changes in inlet temperatures by the MPC algorithm. This prompts preemptive control action so that the tank is heated up in anticipation even before the uncertainty data changes. The result is that the market purchase is reduced when the energy is demanded at a higher inlet temperature.

To conclude, we have demonstrated that an online PCA-based, dynamic scenario-tree adaptation approach leads to solutions that are less conservative while still hedging against the uncertainty. Moreover, the approach involves solving an optimization problem of a smaller size since less scenarios, chosen only along the dominant principal component are needed to describe the uncertainty.

With respect to further work in this domain, multistage MPC could be combined with tube-based MPC in a similar fashion as described in Subramanian et al. (2018) to seek robustness against the uncertainty in the direction of the "insignificant" principal components, which were discarded in scenario selection procedure in this work. In terms of modeling, a thermal storage system with multiple suppliers and consumers of energy is more realistic. Further, the effect of *uncertain* time-varying peak loads on the optimal operation (i.e. using scenarios to describe varying magnitudes of peak loads) can be studied.

REFERENCES

AlGhazzawi, A. and Lennox, B. (2009). Model predictive control monitoring using multivariate statistics. *Journal of Process Control*, 19(2), 314–327.

Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1), 65–98.

Birge, J.R. and Louveaux, F. (2011). *Introduction to Stochastic Programming*. Springer Publishing Company, Incorporated, 2nd edition.

Campo, P.J. and Morari, M. (1987). Robust model predictive control. In *1987 American Control Conference*, 1021–1026.

de Oliveira, V., Jäschke, J., and Skogestad, S. (2016). Optimal operation of energy storage in buildings: Use of the hot water system. *Journal of Energy Storage*, 5, 102 – 112.

Dunning, I., Huchette, J., and Lubin, M. (2017). Jump: A modeling language for mathematical optimization. *SIAM Review*, 59(2), 295–320.

Jolliffe, I.T. and Cadima, J. (2016). Principal component analysis: a review and recent developments. *Philosophical transactions Series A, Mathematical, physical, and engineering sciences*.

Krishnamoorthy, D., Foss, B., and Skogestad, S. (2018a). A distributed algorithm for scenario-base model predictive control using primal decomposition. *IFAC ADCHEM 2018 (In-Press)*.

Krishnamoorthy, D., Thombre, M., Skogestad, S., and Jäschke, J. (2018b). Data-driven scenario selection for multistage robust model predictive control. In *6th IFAC Conference on Nonlinear Model Predictive Control (In-Press)*. IFAC.

Loquasto, F. and Seborg, D.E. (2003). Model predictive controller monitoring based on pattern classification and pca. In *American Control Conference, 2003. Proceedings of the 2003*, volume 3, 1968–1973. IEEE.

Lucia, S. and Engell, S. (2013). Robust nonlinear model predictive control of a batch bioreactor using multi-stage stochastic programming. In *Control Conference (ECC), 2013 European*, 4124–4129. IEEE.

Lucia, S., Finkler, T., and Engell, S. (2013). Multi-stage nonlinear model predictive control applied to a semi-batch polymerization reactor under uncertainty. *Journal of Process Control*, 23(9), 1306–1319.

Maurath, P.R., Laub, A.J., Seborg, D.E., and Mellichamp, D.A. (1988). Predictive controller design by principal components analysis. *Industrial & engineering chemistry research*, 27(7), 1204–1212.

Scokaert, P. and Mayne, D. (1998). Min-max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, 43(8), 1136–1142.

Subramanian, S., Lucia, S., and Engell, S. (2018). A synergistic approach to robust output feedback control: Tube-based multi-stage nmpc. In *10th IFAC International Symposium on Advanced Control of Chemical Processes*, 494–499. IFAC.

Wächter, A. and Biegler, L.T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1), 25–57.