Computers and Chemical Engineering 33 (2009) 1735–1746

Contents lists available at ScienceDirect

Computers and Chemical Engineering

journal homepage: www.elsevier.com/locate/compchemeng

Optimization-based strategies for the operation of low-density polyethylene tubular reactors: nonlinear model predictive control

Victor M. Zavala [1], Lorenz T. Biegler *

Department of Chemical Engineering, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

ARTICLE INFO

Article history:
Received 25 September 2008
Received in revised form 15 February 2009
Accepted 28 April 2009
Available online 5 May 2009

Keywords:
Polyethylene
NMPC
MHE
Large-scale
Nonlinear programming
Economics
Fouling

ABSTRACT

In this work, we present a general nonlinear model predictive control (NMPC) framework for low-density polyethylene (LDPE) tubular reactors. The framework is based on a first-principles dynamic model able to capture complex phenomena arising in these units. We first demonstrate the potential of using NMPC to simultaneously regulate and optimize the process economics in the presence of persistent disturbances such as fouling. We then couple the NMPC controller with a compatible moving horizon estimator (MHE) to provide output feedback. Finally, we discuss computational limitations arising in this framework and make use of recently proposed advanced-step MHE and NMPC strategies to provide nearly instantaneous feedback.

1. Motivation and background information

The decision-making hierarchy in many continuous chemical processes has been traditionally based on detailed *first-principles* steady-state models for economic optimization and empirical *data-driven* dynamic models for multivariable control (Marlin & Hrymak, 1996). Recently, it has been recognized that this model inconsistency can lead to infeasibility issues and performance deterioration (Engell, 2007; Yip & Marlin, 2004). In addition, the disturbances affecting the economics of many chemical processes occur at time scales that cannot be captured adequately by a steady-state optimization layer. The incorporation of first-principles dynamic models in the decision-making hierarchy opens the possibility to avoid these limitations and thus achieve unprecedented process efficiency. This can be done through economics-oriented nonlinear model predictive control (NMPC) or dynamic real-time optimization (D-RTO) schemes (Helbig, Abel, & Marquardt, 2000; Kadam et al., 2003). As expected, an important enabler of these control strategies is the development of efficient optimization strategies able to accommodate computational intensive dynamic models in on-line environments (Zavala, Laird, & Biegler, 2008).

In this work, we seek to integrate detailed first-principles dynamic models to optimize the operation of low-density polyethylene (LDPE) tubular reactor processes. LDPE is an important commodity polymer in today's economy due to its high flexibility and relatively low-cost (Knuuttila, Lehtinen, & Nummila-Pakarinen, 2004). LDPE is typically produced in tubular reactors by free-radical polymerization of ethylene at supercritical conditions (2000–3000 atm and 150–350 °C). A typical tubular reactor and corresponding temperature profiles for the reactor core and jackets are sketched in Fig. 1. LDPE reactors consist of long pipes (1–3 km) with small inner diameters (5–10 cm) and thick reactor walls (2–5 cm) which are divided into several reaction and cooling zones. These multi-zone configurations give rise to strong multivariable interactions along the reactor and thus lead to complex operating procedures. In addition, the operation is further complicated due to persistent dynamic disturbances such as fouling and initiator deactivation (Buchelli et al., 2005; Kiparissides et al., 2005; Luft et al., 1977). These disturbances have a strong impact on the process economics.

The potential economic benefits and high operational complexity of LDPE reactors have motivated research efforts in many areas. Extensive experimental studies have been performed in order to understand the principles governing these systems. This increased level of understanding has translated into numerous first-principles models of different fidelity and complexity (Goto et al., 1981; Kim

* Tel.: +1 412 268 2232.
E-mail addresses: vzavala@mcs.anl.gov (V.M. Zavala), lb01@andrew.cmu.edu (L.T. Biegler).
[1] Present address: Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S Cass Avenue, Argonne IL 60517, USA.

0098-1354/$ – see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.compchemeng.2009.04.008

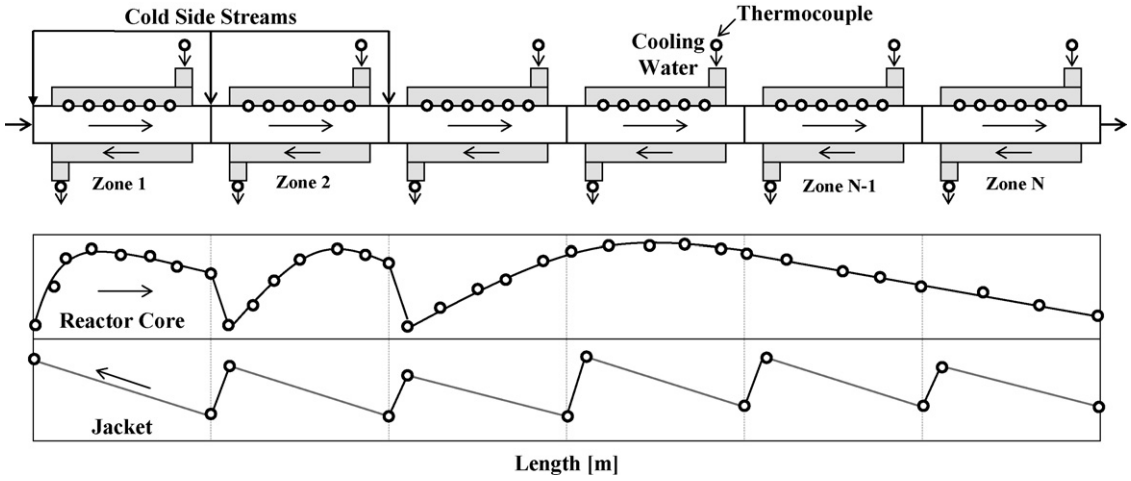

Fig. 1. Schematic representation of multi-zone LDPE tubular reactor (top). Typical reactor core and jacket temperature profiles (bottom).

& Iedema, 2004; Kiparissides, Verros, Pertsinidis, & Goosens, 1996; Kiparissides, Verros, & McGregor, 1993; Zabisky, Chan, Gloor, & Hamielec, 1992; Brandolin, Lacunza, Ugrin, & Capiati, 1996; Bokis, 2001; Häfele, Kienle, Boll, & Schmidt, 2006). The predictive capabilities of these models have also been evaluated and refined through systematic parameter estimation techniques (Kiparissides et al., 2005; Zavala & Biegler, 2006). Some of these models have been used for off-line tasks such as reactor design and dynamic transient analysis (Häfele et al., 2006; Pertsinidis, Papadopoulos, & Kiparissides, 1996). As a natural step, it is desired to use these models to perform on-line tasks such as economic optimization and model-based control. Zavala and Biegler (2008), demonstrated that significant economic benefits can be realized in LDPE reactors through model-based optimization strategies. In that study, a steady-state tubular reactor model described by large sets of differential and algebraic equations (DAEs) was used to evaluate the performance limitations of LDPE reactors in the presence of fouling disturbances. Zavala and Biegler (2008, 2009) demonstrated that detailed LDPE reactor dynamic models can be accommodated in on-line environments with the aid of current optimization capabilities to perform tasks such as moving horizon estimation (MHE) and nonlinear model predictive control (NMPC). These dynamic models are significantly more challenging since they involve computationally intensive partial differential, ordinary differential and algebraic equations (PDAEs) (Häfele et al., 2006). Here, we extend these results and derive a general NMPC framework to optimize the operation of LDPE reactors. We study the performance of both traditional NMPC designs in which the controller is used for regulation around a fixed target and we study the performance of economics-oriented designs in which the controller optimizes the process profitability and simultaneously performs regulation tasks. In addition, we discuss computational issues arising in the integration of NMPC and MHE tasks. Here, we will see that the on-line solution of the associated large-scale PDAE-constrained optimization problems gives rise to long feedback delays. Motivated by this, we make use of recently proposed synchronization or advanced-step strategies for NMPC and MHE based on nonlinear programming (NLP) sensitivity to overcome these limitations and provide nearly instantaneous feedback.

The paper is organized as follows: In the next section, we derive NMPC and MHE formulations for LDPE reactors. In Section 3 we present a short description of the computational strategies used to solve the associated PDAE-constrained optimization problems and to avoid long feedback delays. In Section 4 we discuss the performance of the NMPC framework under diverse scenarios and demonstrate the computational efficiency of the proposed strategies. The paper closes with Section 5 in which we present general conclusions and discuss directions of future work.

2. NMPC and MHE formulations

A typical operational hierarchy in industrial LDPE processes consists of a target-setting layer in which an operator receives the production schedule of different polymer grades. The operator sets the temperature profile of the reactor that is known by experience to give the desired polymer properties (e.g. melt index, density). The temperature targets are communicated to multiple PID controllers distributed along the reactor that try to keep the temperature profile at the desired target. The main tasks of the regulatory (PID) control layer are to reject short-term disturbances and to follow the temperature profile targets provided by the operator during grade changes.

As shown in Fig. 2, the PID controllers are normally grouped by zones in order to regulate the *local* temperature profile at each zone. Input variables such as the initiator flow, the jacket inlet temperature, the jacket inlet flow, and the side stream temperatures can be manipulated independently by each local set of controllers. The fouling onset is, in particular, a difficult disturbance to reject. As the reactor fouls, the controllers need to keep the local temperature profile at target. Because of this, they will tend to compensate by dropping the initiator flows and, implicitly, the production levels. It is important to emphasize that the controllers do not have any knowledge of the production levels of the reactor. Their objective is to keep the temperature profile at target which is set by the operator. Another problem that arises with a regulatory control architecture is that the controllers cannot foresee downstream (cascaded) interactions arising along the reactor. Because of this, the control of polymer properties at the reactor outlet can become cumbersome. It is believed that a *centralized* NMPC strategy able to

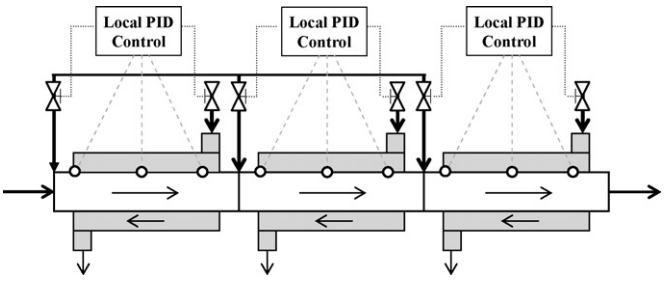

Fig. 2. Regulatory control structure of LDPE tubular reactors.

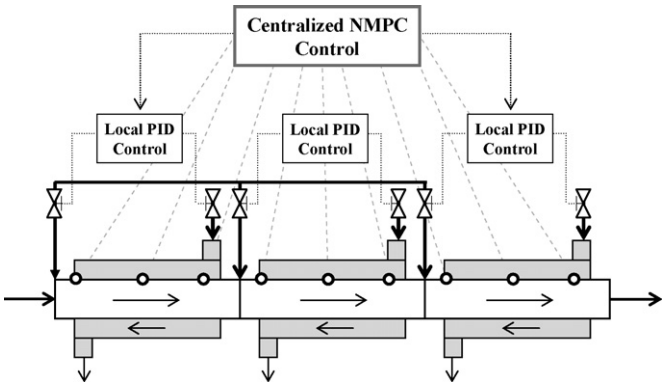

Fig. 3. Centralized NMPC control structure of LDPE tubular reactors.

take all these complex interactions into account would achieve a much better performance; such a strategy is illustrated in Fig. 3.

To formulate the NMPC controller, we consider the scenario in which the LDPE process is located at sampling time t_k. Using the current state as initial conditions, we would like to use a dynamic first-principles model to compute optimal policies for the controls over a future horizon $t \in [t_k,\ t_{k+N}]$ with sampling times of equal length $\delta = t_{k+1} - t_k$ that minimize a given performance index. The NMPC problem formulation under consideration is given by,

$$\min_{u(j\delta)} \sum_{j=0}^{N}\sum_{i=1}^{N_m} (\gamma(j\delta, x_i) - \bar{\gamma}_{k+j,i})^T \mathbf{Q}_\gamma^{-1} (\gamma(j\delta, x_i) - \bar{\gamma}_{k+j,i}) + \sum_{j=0}^{N-1} (u(j\delta) - \bar{u}_{k+j})^T \mathbf{Q}_u^{-1} (u(j\delta) - \bar{u}_{k+j}) \tag{1a}$$

s.t.

$$\frac{\partial z}{\partial \tau} + \nu(\tau, x)\frac{\partial z}{\partial x} + \kappa(\tau, x)\frac{\partial^2 z}{\partial x^2} = f_z(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{1b}$$

$$\frac{\partial w}{\partial x} = f_w(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{1c}$$

$$0 = f_y(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{1d}$$

$$\gamma(\tau, x) = \chi(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{1e}$$

$$0 \geq g(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{1f}$$

$$z(0, x) = z^k(x) \tag{1g}$$

$$0 = \varphi\left(z(\tau, 0), z(\tau, x^L), \frac{\partial z}{\partial x}(\tau, 0), \frac{\partial z}{\partial x}(\tau, x^L), w(\tau, 0), u(\tau)\right)$$

$$\tau \in [0,\ \ N\delta]. \tag{1h}$$

To represent the overall dynamic model of the multi-zone reactor in simple terms, we collapse the PDAEs corresponding to material and energy balances, thermodynamics and transport expressions, and kinetic expressions for all zones into a single set of PDAEs (1b)–(1d). This can be done by grouping the states corresponding to all zones into a single variable vector and defining the continuity equations between zones as algebraic equations. With this, the boundary conditions (1h) need only to be defined at the reactor inlet and outlet points given by $x = 0$ and $x = x^L$, respectively. In addition, we define the internal time dimension for the NMPC problem $\tau := t - t_k$ that uses the current time of the process t_k as reference point. Symbol $z(\tau, x) \in \mathcal{R}^{n_z}$ is used to represent the differential states in space and time. The current state is denoted by $z^k(x)$ which is a fixed quantity. In LDPE reactors, these states correspond to the cooling water temperature in the jacket and the reactor wall temperature at all zones. Symbol $w(\tau, x) \in \mathcal{R}^{n_w}$ denotes differential states in space such as the reacting mixture temperature, the molar flow rates of gaseous components, the chain moments, among others. These states arise from a quasi-steady-state assumption of the reactor core (Kiparissides et al., 2005; Zavala & Biegler, 2009). For instance, symbol $y(\tau, x) \in \mathcal{R}^{n_y}$ denotes the algebraic states corresponding to the rest of the model variables such as the cooling water and reacting mixture velocities, densities, and heat capacities. Symbol $p(\tau) \in \mathcal{R}^{n_p}$ denotes time-varying parameters used to account for unmodeled effects and uncertainty. For instance, these can represent the heat transfer coefficients (HTCs) of the reactor zones which vary with time due to fouling. Symbol $u(\tau) \in \mathcal{R}^{n_u}$ denotes the model inputs corresponding to side-stream inlet temperatures, and flow rates. In the actual reactor, the inputs can only be fed at the beginning and end of each reactor zone. Consequently, no explicit dependence on the *internal* spatial dimension x is considered.

We define a set of output variables $\gamma(\tau, x) \in \mathcal{R}^{n_\gamma}$ in (1e) that map all the model states into a set of output variables *measured* in the actual LDPE process. The performance index of the NMPC problem can be a general function of the model states and inputs but in the above formulation we use a tracking objective in order to simplify the presentation. This tracking objective implicitly minimizes the transition time from the current state to the desired target defined as $\bar{\gamma}$. In an actual industrial reactor, the outputs are normally controlled at discrete points in time $j\delta$ and space x_i with corresponding targets $\bar{\gamma}_{k+j,i}$. The total number of measurement locations in space is denoted by N_m. In addition, we impose zero hold constraints on the inputs (i.e. $u(\tau) := u(j\delta)$, $\tau \in [j\delta, (j+1)\delta]$). The weighting matrices for the outputs and inputs are given by $\mathbf{Q}_\gamma^{-1}$ and $\mathbf{Q}_u^{-1}$, respectively. The dynamic reactor model contains about 3 PDEs, 20 ODEs, and 500 AEs. The equations are defined over long axial horizons (1–2 km) containing all the reactor zones and time horizons of less than one hour. For a more detailed explanation of the model equations we refer the reader to Zavala and Biegler (2006, 2009).

From the solution of the NMPC problem, we extract the current control actions $u^k = u^*(\delta)$. The plant then evolves to the next state $z^{k+1}(x)$ and this is used to compute the next control action. In the nominal NMPC case, we assume that the model is perfect and use the predicted states as the actual states of the reactor. In the presence of model mismatch and uncertainty, it is necessary to use a MHE estimator to infer the current state based on the available measurement information. In this work, we make use of the MHE estimator presented in Zavala and Biegler (2009). To formulate the MHE problem, we consider the scenario in which the LDPE process is currently located at sampling time t_k and we have a *past* output measurement history $\{\vartheta_\gamma^k, \vartheta_\gamma^{k-1}, \ldots, \vartheta_\gamma^{k-N}\}$ and corresponding inputs $\{\vartheta_u^{k-1}, \vartheta_u^{k-2}, \ldots, \vartheta_u^{k-N}\}$ distributed over a time horizon containing N sampling times. Using this information, we seek to obtain an estimate of the current *true* state of the reactor $z^k(x)$ and the parameters p^k through the model. We assume that the model structure is *correct* and that all the uncertainty related to the real process can be encapsulated in the model parameters $p(\tau)$, in the initial conditions $z_0(x)$, and in the measurement errors. Following this reasoning, the MHE estimation problem becomes,

$$\min_{p(\tau), z_0(x), u(\tau)} (z_0(x) - \bar{z}^{k-N}(x))^T \mathbf{\Pi}_{k-N}^{-1} (z_0(x) - \bar{z}^{k-N}(x)) + \sum_{j=0}^{N}\sum_{i=1}^{N_m} (\gamma(j\delta, x_i) - \vartheta_\gamma^{k+j-N,i})^T \mathbf{R}_\gamma^{-1} (\gamma(j\delta, x_i) - \vartheta_\gamma^{k+j-N,i}) + \sum_{j=0}^{N-1} (u(j\delta) - \vartheta_u^{k+j-N})^T \mathbf{R}_u^{-1} (u(j\delta) - \vartheta_u^{k+j-N}) \tag{2a}$$

s.t.

$$\frac{\partial z}{\partial \tau} + \nu(\tau, x)\frac{\partial z}{\partial x} + \kappa(\tau, x)\frac{\partial^2 z}{\partial x^2} = f_z(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{2b}$$

$$\frac{\partial w}{\partial x} = f_w(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{2c}$$

$$0 = f_y(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{2d}$$

$$\gamma(\tau, x) = \chi(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{2e}$$

$$0 \geq g(z(\tau, x), w(\tau, x), y(\tau, x), p(\tau), u(\tau)) \tag{2f}$$

$$z(0, x) = z_0(x) \tag{2g}$$

$$0 = \varphi\left(z(\tau, 0), z(\tau, x^L), \frac{\partial z}{\partial x}(\tau, 0), \frac{\partial z}{\partial x}(\tau, x^L), w(\tau, 0), u(\tau)\right)$$

$$\tau \in [0, \ \delta N]. \tag{2h}$$

In this case, t_{k-N} is used as the zero reference time so that $\tau := t - t_{k-N}$. As before, note that the process output measurements are only defined at discrete points in time k and space x_i denoted by $\vartheta_\gamma^{k,i}$. In addition, we can simultaneously reconcile the inputs to the past measurements ϑ_u^k in order to eliminate measurement noise (Zavala & Biegler, 2009). The matrices $\mathbf{R}_\gamma \in \mathcal{R}^{n_\gamma \times n_\gamma}$ and $\mathbf{R}_u \in \mathcal{R}^{n_u \times n_u}$ are covariance matrices for the expected output and input measurement errors, respectively. The first term in the objective function is the *arrival cost*, which summarizes past measurement information before t_{k-N}. The internal variable $z_0(x)$ represents the unknown initial states at time t_{k-N} while $\bar{z}^{k-N}(x)$ denotes the corresponding *a priori* value with covariance $\mathbf{\Pi}_{k-N} \in \mathcal{R}^{n_z \times n_z}$. From the solution of the MHE problem, we extract the estimate of the true state as $\tilde{z}^k(x) \leftarrow z^*(\delta N, x)$, the current parameters $\tilde{p}^k \leftarrow p^*(\delta N)$. We use this information to obtain the current control action u^k from the solution of the NMPC problem. Having this new input measurement ϑ_u^k we wait for the corresponding outputs ϑ_γ^{k+1} measured at t_{k+1}. We drop the last measurements and incorporate the new ones to obtain the new measurement histories $\{\vartheta_\gamma^{k+1}, \vartheta_\gamma^k, \ldots, \vartheta_\gamma^{k+1-N}\}$ and $\{\vartheta_u^k, \vartheta_u^k, \ldots, \vartheta_u^{k+1-N}\}$. Accordingly, we update the *a priori* estimate of the initial state using the previous solution as $\bar{z}^{k-N+1}(x) \leftarrow z^*(\delta, x)$. In some cases, it is also necessary to update the covariance matrix to $\mathbf{\Pi}_{k-N+1}$ using the covariance of the predicted state $z^*(\delta, x)$ to account for poor estimates of the initial conditions (Rawlings & Bakshi, 2006). For strongly observable systems such as LDPE reactors, this update might not be necessary to ensure convergence of the MHE estimator (Alessandri, Baglietto, & Battistelli, 2008; Zavala & Biegler, 2009). Using the updated information we solve the new MHE problem to obtain the new state estimates. In Fig. 4, we sketch the interaction between MHE and NMPC and their corresponding time horizons.

As can be seen, the MHE and NMPC problems arising in this application are computationally expensive PDAE-constrained optimization problems. This is a critical issue since the feedback delay is proportional to both the solution time of the MHE problem to obtain the state estimates and the solution time of the NMPC problem to obtain the control actions. In the following section, we describe the approach taken to solve these problems. In addition, we present synchronization strategies to minimize the feedback delay. In particular, we will exploit the fact that the NMPC problem is parametric on the initial state and the fact that the MHE problem is parametric on the measurements. This property can be used to provide fast approximate solutions constructed around reference solutions obtained in between sampling times.

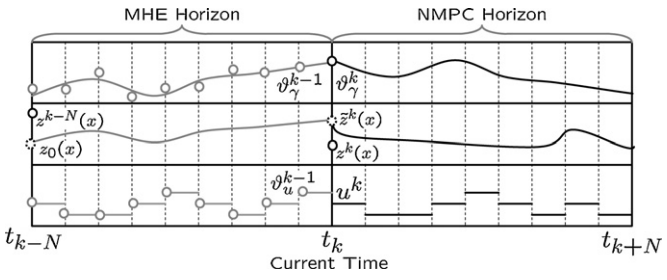

Fig. 4. Schematic representation of MHE and NMPC horizons.

3. Computational strategies

We follow a full discretization approach to solve the PDAE-constrained optimization problems. For the discretization along the axial dimension we follow a Radau collocation on finite elements scheme: an average of 10 finite elements for the reaction zones and 2 finite elements for the cooling zones. Three collocation points are used in each spatial element. Upon spatial discretization, the PDAE reactor model translates into a DAE model containing around 9000–10,000 DAEs in time. For time discretization, we follow a 1 pt. Radau collocation scheme (implicit Euler). We use 5–20 time steps for time discretization.

After discretization, the resulting NLP problems are solved with IPOPT (Wächter & Biegler, 2006). The NLPs have the general form,

$$\min_{\mathbf{x}} F(\mathbf{x}, \eta) \tag{3a}$$

$$\text{s.t. } c(\mathbf{x}, \eta) = 0 \tag{3b}$$

$$\mathbf{x} \geq 0 \tag{3c}$$

where $\mathbf{x} \in \mathcal{R}^{n_\mathbf{x}}$ is variable vector containing all the discretized states, controls and parameters, and $\eta \in \mathcal{R}^{n_\eta}$ is a *fixed* data vector that can be used to represent the moving initial conditions in the NMPC problem and the measurements in the MHE problem.

In interior-point solvers, the inequality constraints of problem (3) are handled *implicitly* by adding barrier terms to the objective function,

$$\min_{\mathbf{x}} F(\mathbf{x}, \eta) - \mu \sum_{j=1}^{n_\mathbf{x}} \ln(\mathbf{x}_{(j)}) \tag{4a}$$

$$\text{s.t. } \mathbf{c}(\mathbf{x}, \eta) = 0 \tag{4b}$$

where $\mathbf{x}_{(j)}$ denotes the j-th component of vector $\mathbf{x}$. Solving (4) for a decaying sequence of $\mu \to 0$ results in an efficient strategy to solve the original NLP (3). Using an initial barrier parameter μ, IPOPT tries to solve the Karush-Kuhn-Tucker (KKT) conditions of a sequence of barrier problems (4),

$$\nabla_\mathbf{x} F(\mathbf{x}, \eta) + \nabla_\mathbf{x} \mathbf{c}(\mathbf{x}, \eta)\boldsymbol{\lambda} - \boldsymbol{\nu} = 0 \tag{5a}$$

$$\mathbf{c}(\mathbf{x}, \eta) = 0 \tag{5b}$$

$$\mathbf{X} \cdot \mathbf{V}\, \mathbf{e} = \mu \mathbf{e} \tag{5c}$$

where $\mathbf{X} = \text{diag}(\mathbf{x})$, $\mathbf{V} = \text{diag}(\nu)$ and $e \in \mathcal{R}^{n_\mathbf{x}}$ is a vector of ones. Note that the KKT conditions of the barrier problem match those of the original NLP for $\mu = 0$. Symbols $\lambda \in \mathcal{R}^{n_\lambda}$ and $\nu \in \mathcal{R}^{n_\mathbf{x}}$ are Lagrange multipliers for the equality constraints and bounds, respectively. The gradient of the objective function is $\nabla_\mathbf{x} F(\mathbf{x}, \eta) \in \mathcal{R}^{n_\mathbf{x}}$ while $\nabla_\mathbf{x} \mathbf{c}(\mathbf{x}, \eta) \in \mathcal{R}^{n_\mathbf{x} \times n_\lambda}$ is the constraint Jacobian. To solve this system of nonlinear equations IPOPT uses an exact Newton method. We initialize the iteration sequence at point $s_o^T := [\mathbf{x}_o^T \ \lambda_o^T \ \nu_o^T]$. At the i th iteration, the search direction $\Delta s_i = s_{i+1} - s_i$ is computed by

linearization of the KKT conditions (5),

$$\begin{bmatrix} \mathbf{H}_i & \mathbf{A}_i & -\mathbb{I}_{n_\mathbf{x}} \\ \mathbf{A}_i{}^T & 0 & 0 \\ \mathbf{V}_i & 0 & \mathbf{X}_i \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x}_i \\ \Delta\lambda_i \\ \Delta\nu_i \end{bmatrix} = - \begin{bmatrix} \nabla_\mathbf{x} F(\mathbf{x}_i, \eta) + \mathbf{A}_i\lambda_i - \nu_i \\ \mathbf{c}(\mathbf{x}_i, \eta) \\ \mathbf{X}_i\mathbf{V}_i e - \mu e \end{bmatrix} \quad (6)$$

where $\mathbf{A}_i := \nabla_\mathbf{x}\mathbf{c}(\mathbf{x}_i, \eta)$. Matrix $\mathbf{H}_i \in \mathcal{R}^{n_\mathbf{x}\times n_\mathbf{x}}$ is the Hessian of the Lagrange function $\mathcal{L} = F(\mathbf{x}_i, \eta) + \mathbf{c}(\mathbf{x}_i, \eta)\lambda_i - \nu_i{}^T\mathbf{x}_i$ and $\mathbf{\Sigma}_i := \mathbf{X}_i^{-1}\mathbf{V}$. Symbol $\mathbb{I}_{n_\mathbf{x}}$ denotes the identity matrix. We provide exact Hessian and Jacobian information through the modeling platform AMPL (Fourer, Gay, & Kernighan, 1992). After solving a sequence of barrier problems for $\mu \to 0$, the solver returns the optimal solution triplet $s_*^T(\eta) = [\mathbf{x}_*^T\ \lambda_*^T\ \nu_*^T]$ which is an implicit function of the fixed data vector. To reduce the number of iterations required to solve the neighboring MHE and NMPC problems from one sampling time to the next, we warm-start the problems using the shifted trajectories for the states and multipliers and keep the barrier parameter μ to a small value (Zavala & Biegler, 2008).

Solving the KKT system (6) is the most computationally intensive step in the solution of the NLP. However, the KKT matrix arising from the PDAE-constrained optimization problems is highly sparse and presents multiple nested structures that can potentially be exploited. In IPOPT, after eliminating the bound multipliers from the KKT system, we apply a symmetric indefinite factorization of the resulting KKT matrix. The computational complexity of this strategy is in general very favorable. However, significant fill-in and computer memory bottlenecks might arise during the factorization step if the sparsity pattern is not properly exploited. In order to factorize the KKT matrix, we use the general-purpose linear solver MA57 from the Harwell library (Duff, 2004). Recently, we have explored the impact of different reordering strategies on the factorization time of the KKT matrix. The default reordering strategy used in MA57 is an approximate minimum degree (AMD) ordering algorithm. Another strategy available is a nested dissection ordering based on the multi-level graph partitioning algorithm implemented in Metis (Karypis & Kumar, 1999). For large-scale problems such as those arising in PDAE-constrained optimization, we have observed that nested dissection techniques are significantly more efficient and permit much faster factorizations (Zavala & Biegler, 2009; Zavala & Biegler, 2008).

In the context of MHE-NMPC, it is important to realize that the factorization of the KKT matrix limits the ability to provide fast feedback. This is due to the fact that, in standard MHE and NMPC strategies, we must wait for the measurement information to start the solution of the optimization problems (Diehl, Bock, & Schlöder, 2005). To avoid these limitations, we make use of the so-called advanced-step strategies (Zavala, Laird, & Biegler, 2008; Zavala & Biegler, 2009). The main idea here is to use the model and the current control action u^k to predict the future state $\bar{z}^{k+1}(x)$ and associated measurements $\bar{\vartheta}^{k+1}$. Using this information, we solve reference MHE and NMPC problems simultaneously *in between* sampling times and hold the corresponding KKT matrices at the solutions. Once the true measurements ϑ^{k+1} become available, we perturb the right-hand side of the KKT system (6) of the MHE problem (2) and execute a fast backsolve to obtain an approximate state estimate $z^{k+1}(x)$. This strategy is known as advanced-step MHE (as-MHE). We then use the as-MHE state estimate to perturb the KKT system of the NMPC problem (1) to obtain the approximate control action u^{k+1}. This strategy is known as advanced-step NMPC (as-NMPC). Numerical properties of advanced-step strategies have been presented in (Zavala & Biegler, 2008). In that work it is demonstrated that, under mild assumptions, the feedback delay can be reduced by two orders of magnitude without significant losses in performance.

4. Case studies and results

In this section, we analyze the performance of the proposed NMPC framework under a particular scenario arising in the operation of LDPE reactors. The study is based on confidential industrial data so we normalize all the variables using the values at the initial point. These values are also used as targets throughout the horizon. We first discuss the performance of a nominal NMPC controller and we then discuss the performance of an output feedback MHE-NMPC controller. Finally, we analyze the computational performance of the proposed strategies.

4.1. Tracking NMPC

We consider a nominal NMPC controller in which the performance index is a tracking objective. In a typical LDPE process, the controlled outputs are the temperature profile, the polymer melt index and the polymer density. Here, we assume that the outlet temperatures of the jackets are not controlled variables but we impose physical bounds on them (e.g. reasonable temperature levels for cooling water going back to cooling tower). It is also important to emphasize that the temperature profile of the core is normally controlled at discrete positions along the reactor (i.e. where the thermocouples are present). Nevertheless, having a first-principles model allows to control or impose constraints on temperatures at *unmeasured* positions. In addition, we can control and/or impose constraints on unmeasured properties such as the polydispersity or molecular weights. The available manipulated variables include the input variables of Table 1.

To test the performance of the tracking NMPC controller we consider a typical fouling–defouling scenario arising in the operation of LDPE reactors. LDPE tubular reactors undergo periods where the polymer layer is *defouled* by means of pressure or thermal shocks (Buchelli et al., 2005). After this, the reactor fouls again. Stabilizing the reactor under these disturbance cycles is important in order to avoid reactor runaway and to keep the polymer properties at target. To simulate the NMPC cleaning-fouling cycle, we ramp the HTCs for all the reactor zones from their nominal value to +50% and then back to −20%. Since the controller cannot predict the behavior of the HTCs, they act as exogenous disturbances (represented in the model as parameters). This is done in order to test the inherent robustness properties of the controller. The plant response is obtained by simulating the reactor model using the current control action and the *true* HTCs. Each time step corresponds to 2 min of operation. We use a prediction horizon of 10 time steps. The closed-loop simulations are run for a hundred time steps. The objective of the controller in this scenario is to keep the temperature profile as close as possible to a given target profile as the reactor fouls and defouls. In addition, the polymer melt index deviations should be kept close to the target.

In the top graph of Fig. 5 we present the heat transfer profiles for zones 2 and 3. Note that in the first period the reactor defouls so the values of the heat transfer coefficient go up, while in the second the reactor fouls and the heat transfer coefficients go down. Note also that the reactor zones have different cooling capabilities due to design considerations or uneven fouling along the reactor. In the middle graph we present the closed-loop response of the reactor temperature profile. The dark line is the target profile and the gray lines represent the controller responses. As we can see, the controller is able to stabilize the system and keep the temperature profile close to the reference. The impact of the fouling–defouling disturbances can be appreciated in the bottom graph. Here, we present the dynamic responses for the wall and jacket temperatures for the first 35 time steps (defouling). The temperature levels rise significantly (more than 30% from reference value) due to the decrease of the fouling layer thickness (increasing heat transfer

Table 1
Summary of available measurements in industrial LDPE reactors.

Measurement	Input	Output	On-Line Sensor	Laboratory
Inlet Pressure	x		x	
Jacket Inlet Temperatures	x		x	
Sidestreams Inlet Temperatures	x		x	
Ethylene Sidestream Flow Rates	x		x	
Comonomer Sidestream Flow Rates	x		x	
CTA Sidestream Flow Rates	x		x	
Initiator Flow Rates	x		x	
Temperature Profile		x	x	
Overall Conversion		x	x	
Jacket Outlet Temperatures		x	x	
Melt Index		x	x	x
Weight-Average MW (Mww)		x		x
Number-Average MW (Mwn)		x		x
Polymer Density		x		x
Long-Chain Branching		x		x

coefficient). In other words, more heat can be dissipated to the jacket and walls.

In Fig. 6 we present the reactor responses for the polymer melt index and the reactor overall production. In the top graph, we can observe that the controller achieves tight control of the melt index. The melt index variability is less than 0.1% with respect to the reference value. Although a proper control of the reactor temperature profile is important to keep the melt index at target, the controller also makes use of the chain-transfer flow rate to help with the control. In the bottom graph we present the response for the reactor overall production. In the first transition, note that the normalized production level recovers as the reactor is cleaned or defouled. This results from the increase of the closed-loop initiator flows that keep the temperature profile peaks at target. On the other hand, when the reactor fouls again, the controller needs to drop the initiator flows and this in turn results in decreased production levels. It is important to emphasize that the controller objective is to keep the reactor at the desired temperature profile and not to control the production levels. In other words, the fluctuating production levels are an implicit consequence of the control actions.

In Fig. 7 we present the closed-loop responses for some of the controller inputs. In the top graph we can see the response of the jacket inlet temperatures for zones 2 and 3. Note that the controller responses do not follow obvious trends. This implies that there exists a relatively high degree of nonlinearity and coupling in the system. In particular, the jacket inlet temperatures have a strong impact on the outlet temperatures of the reaction zones (where the initiators are totally consumed) and on the temperature profiles of the cooling zones. In the middle graph we can see that the initiator flows follow a monotonic trend as we ramp the HTCs. The controller decreases the amount of initiator during the fouling periods and increases it during the defouling periods. This is done in order to compensate for the fluctuating cooling capacity. Since the reactor temperature levels are largely determined by the heat of reaction, the initiator flows are used to control the local conversion levels in each zone. In the bottom graph we appreciate

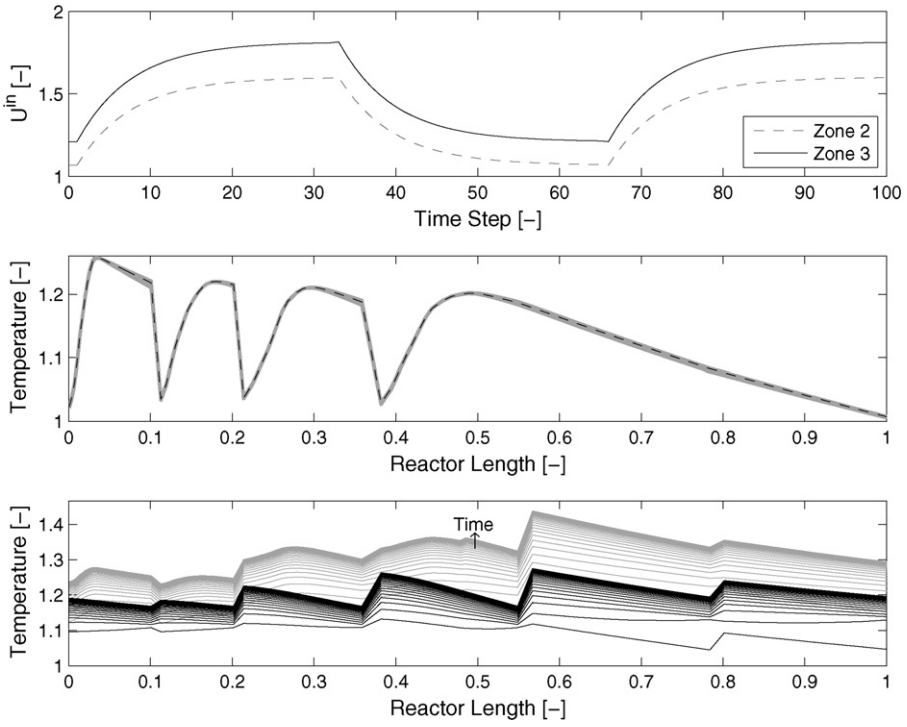

Fig. 5. NMPC controller behavior. Ramping of heat transfer coefficients (top). Controller response for temperature profile (middle)—gray lines are the responses and dashed dark line is the reference temperature profile. Responses of wall (gray) and jacket (black) temperatures (bottom).

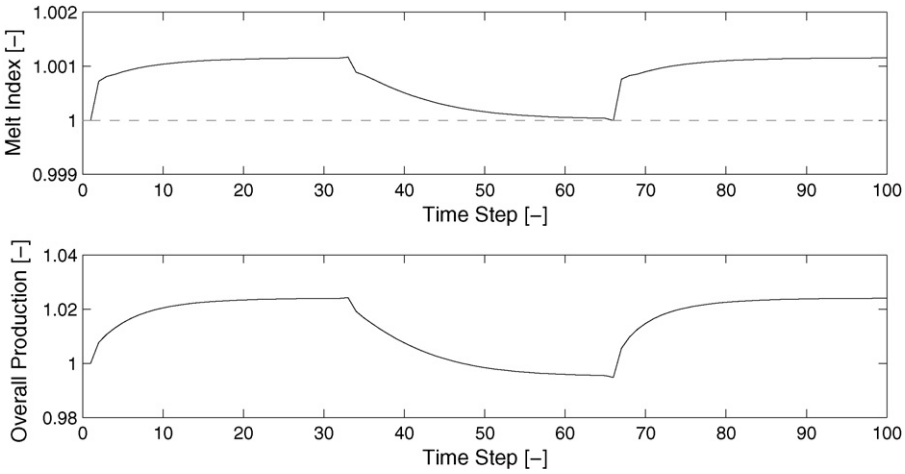

Fig. 6. NMPC controller behavior. Closed-loop response of outlet melt-index (top), target is dashed line and control profile is solid line. Response of reactor overall production (bottom).

the closed-loop response of the the side-stream temperatures. The side-stream temperatures tend to have the strongest impact on the inlet temperatures at each zone. In addition to the initiator flows, the inlet side-stream temperatures play an important role in shaping the temperature profile at each zone. However, in this particular scenario, the controller does not make use of the side-stream temperature of Zone 3. This illustrates that the controller is implicitly taking into account interactions between zones.

4.2. Economics-oriented NMPC

One would expect that the incorporation of a centralized NMPC controller would result in increased robustness and performance since it can better handle the multivariable interactions along the LDPE reactor. However, it is often difficult to appreciate the economic benefits of incorporating such a sophisticated controller. In principle, we could appreciate these benefits more easily if we would use a real-time optimization (RTO) layer in order to compute steady-state targets (Zavala, & Biegler, 2008) that optimize the process profitability and then use the NMPC controller to obtain fast transitions between targets. However, this approach is difficult to apply to LDPE processes since they are seldom at steady-state due to persistent dynamic disturbances such as fouling. In this work, we propose the incorporation of an economic objective to the NMPC controller. This economics-oriented controller can be seen as an all-at-once RTO-NMPC strategy that directly optimizes the process performance. We refer to this strategy as D-RTO. Such a strategy can account for the dynamic disturbances that affect the process profitability and can exploit more efficiently the degrees of freedom available in distributed LDPE reactors (Engell, 2007).

For the design of the D-RTO or economics-oriented NMPC controller, we will incorporate the accumulated production rate as a measure of process profitability. This is justified by the fact that, in high-throughput LDPE processes, the operating costs are

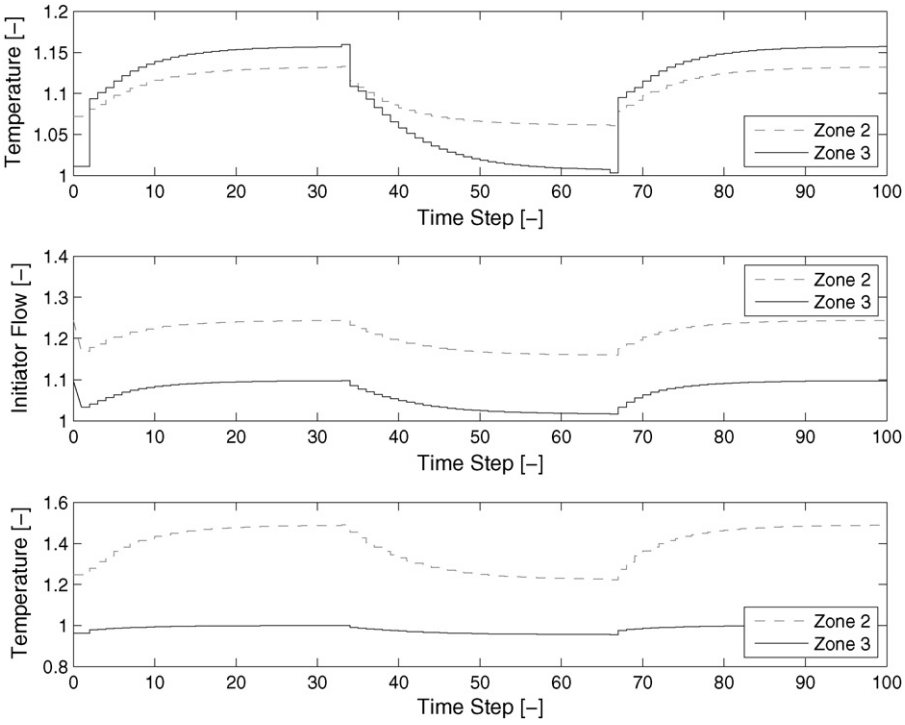

Fig. 7. NMPC controller behavior. Jacket inlet temperatures (top). Initiator flows (middle). Side-streams temperatures (bottom).

largely affected by production losses associated with fouling. However, more detailed economic objectives can also be considered. The objective function of the controller will take the following form,

$$\Phi^k = -\alpha_e \sum_{j=0}^{N} Production(jN) + \alpha_t \sum_{j=0}^{N} \sum_{i=1}^{N_m} (\gamma(j\delta, x_i) - \bar{\gamma}_{k+j,i})^T \mathbf{Q}_\gamma^{-1} (\gamma(j\delta, x_i) - \bar{\gamma}_{k+j,i}) + \alpha_t \sum_{j=0}^{N-1} (u(j\delta) - \bar{u}_{k+j})^T \mathbf{Q}_u^{-1} (u(j\delta) - \bar{u}_{k+j}) \quad (7)$$

Here, the controller tries to minimize the transition time from the current state to the targets and simultaneously tries to maximize the reactor production. We have found that a *purely economic objective* (first term in (7)) leads to *ill-posed* optimization formulations due to the large number of degrees of freedom and the distributed nature of LDPE reactors. With a purely economic objective we cannot guarantee the uniqueness of the solution (e.g. second order optimality conditions do not hold) and this leads to ill-posed search directions in IPOPT that make the NMPC problems difficult to solve. These observations agree with those of (Huesman, Bosgra, and Van den Hof, 2007) and with those of Skogestad (2000) in the context of steady-state RTO. In this work, we add a tracking term to *regularize* the problem and thus obtain smoother solutions. Notice the addition of positive weighting factors α_t and α_e, which are controller design parameters. In the context of LDPE processes, the tuning factor α_t reflects how much the controller is allowed to deviate from the target temperature profile in order to maximize the production rate. In this study, we needed to set $\alpha_t = 10$ and $\alpha_e = 1$ in order obtain unique solutions. Note also that the production term is a highly complex nonlinear mapping of the model states, given by the total molar flow of polymer at the reactor exit. Here, we emphasize the importance of using exact second-order derivative information in the NLP solver to handle these complex objectives (as opposed to traditional Quasi-Newton approximations).

We have tested the D-RTO controller using the same fouling–defouling scenario analyzed in the previous section. In Fig. 8 we compare the performance of the NMPC (tracking) and the D-RTO (economic) controllers. In the top graph we see that the tracking controller keeps the reactor temperature profile close to the given target. On the other hand, in the middle graph we see that the D-RTO controller responses tend to deviate from the target. In the bottom graph we see that the adjustment of the temperature profile results in a direct increase on the reactor production levels during both the fouling and the defouling periods. The D-RTO controller recognizes that the given target profile *is not the optimum profile* in terms of reactor productivity and tends to correct it. This results in an increase in the reactor production rate, as seen in the bottom graph. Since the controller also manipulates the chain-transfer agent flow rate, these changes can be made without having to sacrifice polymer quality (e.g. melt index).

In general, we can conclude that the main benefit of the D-RTO controller is the ability to find better strategies to distribute the polymer production across multiple zones. Since this in turn depends on the time-varying fouling levels at each zone, the controller is an efficient alternative to manage fouling issues. Nevertheless, from the bottom graph of Fig. 8 we can also see that the D-RTO controller *still needs to drop production* during the fouling period in order to remain feasible and stable. This implies that the reactor performance is still limited by the fouling onset.

4.3. MHE-NMPC coupling

We now couple the NMPC controller to an MHE estimator to account for potential plant-model mismatch and estimator errors. In addition, we discuss computational scale-up issues, and use advanced-step strategies to overcome feedback delays.

4.3.1. Output feedback performance

We simulate an output feedback controller in the fouling–defouling scenario described in the previous section. Again, this is done by ramping the reactor HTCs. Since the LDPE reactor model cannot predict the fouling disturbance, we use the MHE estimator to estimate the HTCs p^k and the unmeasured model states $\tilde{z}^k(x)$ (e.g. wall temperature profile) at each time step. The objective of the NMPC controller is to use the estimated reactor state $\tilde{z}^k(x)$ to drive the reactor outputs to the desired target by computing the optimal control action u^k. In this simulated scenario, we generate the true plant response $z^k(x)$ from the model with the *true* HTCs. Since the MHE estimator starts with *wrong* initial guesses of the state and the parameters, and since we add Gaussian measurement noise, it will introduce an estimation error that acts as an additional disturbance to the controller. In addition, note that since the fouling phenomenon cannot be predicted, the estimator can only converge to the true value of the HTCs one step behind. Once the HTC disturbance vanishes, the estimator converges to the true values and the NMPC controller recovers its nominal stability properties.

In the top graph of Fig. 9 we compare the predicted temperature profile of the NMPC controller at time step 25 and the corresponding profile corrupted with $\sigma = 3\%$ Gaussian noise. For a reference temperature of 100 °C this represents a variance of 9 °C. In the middle graph we see that the MHE estimation converges to the true wall profile in less than 10 time steps despite the noise and the wrong initial guess. In the bottom graph we see the effect of noise on the control actions by the NMPC controller. In the top graph of Fig. 10 we illustrate the convergence of the MHE estimator to the jacket water temperature at a particular point along the reactor. In the middle graph, we illustrate the convergence of wall temperature at the same axial position. In the bottom graph we can see that the NMPC controller is able to stabilize the system despite multiple disturbances, and is able to keep the reactor temperature profile close to the target. It is important to emphasize that the strong observability properties of the LDPE reactor result in good robustness and stability properties of the NMPC controller (Zavala & Biegler, 2009). In particular, the number and location of the temperature measurements available along the reactor are important for this.

4.3.2. Computational performance

In this section, we discuss computational issues associated with the proposed MHE and NMPC strategies for LDPE reactors. In order to motivate the discussion, we present results on the total wall-clock time required to solve the tracking NMPC and the MHE problems in the output feedback case study of the previous section. All calculations were obtained using a quad-core Intel processor running Linux at 2.4 GHz. The solution times also include some overhead (around 10–20 s) coming from I/O communication tasks arising in the implementation and from AMPL, which requires some time to generate the derivative information before calling the NLP solver. The prediction and estimation horizons N were set to 10 time steps (20 min) and sampling times of 2 min were used. The NMPC problem consists of an NLP with 80,950 constraints and 370 degrees of freedom. The MHE problem consists of an NLP with 80,300 constraints and 648 degrees of freedom. As can be seen in the top graph of Fig. 11, the overall solution time for the NMPC problem is around 60 s. The NLP solver requires 3–4 iterations to converge the problems. In the bottom graph we present total solution times for the

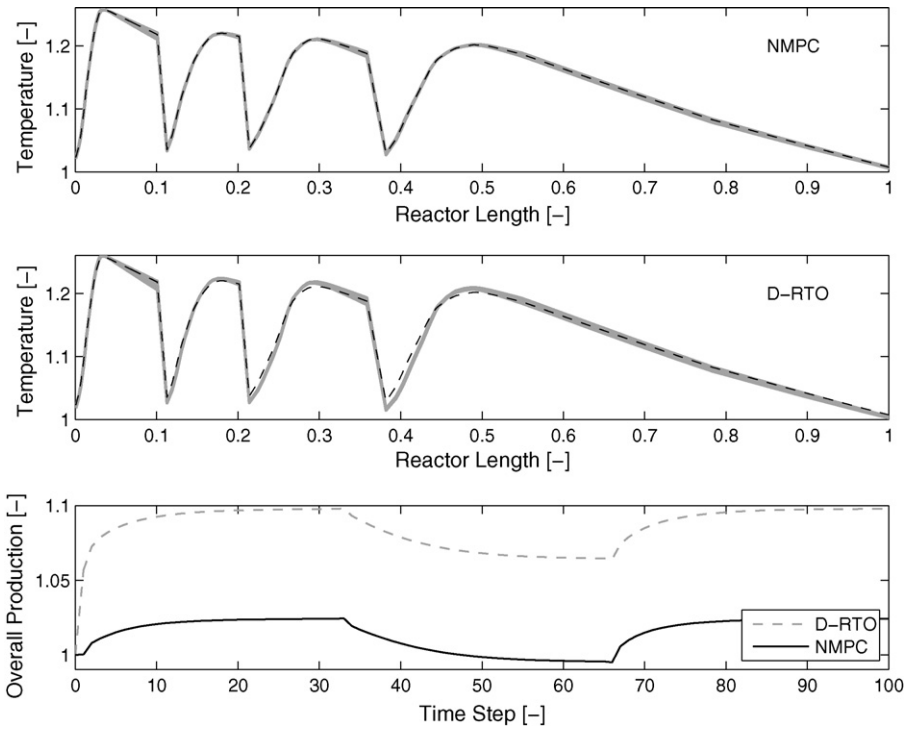

Fig. 8. Comparison of NMPC and D-RTO controller responses. Closed-loop response of temperature profile for NMPC (top) - dashed dark line is target and gray lines are the controller responses. Closed-loop response of temperature profile for D-RTO (middle) - dashed dark line is target and gray lines are the controller responses. Reactor overall production responses (bottom).

MHE estimator. In this case, the estimator is initialized in batch mode (accumulate measurements until an estimation horizon of N time steps is full). Once the estimation horizon is complete, IPOPT takes around 70 s to solve the problem (3–4 iterations are required).

As part of scale-up results, we have found that a nested dissection reordering strategy in MA57 allows the Newton step (6) to have linear scale-up of the solution time as the time horizons of the NMPC and MHE problems are increased (Zavala & Biegler, 2008). This is due to the fact that the nested dissection reordering applied to (6) can identify coarse-grained structures present in the problems more easily (LDPE multi-zone model, finite element structure, etc.), while traditional minimum degree ordering strategies tend to

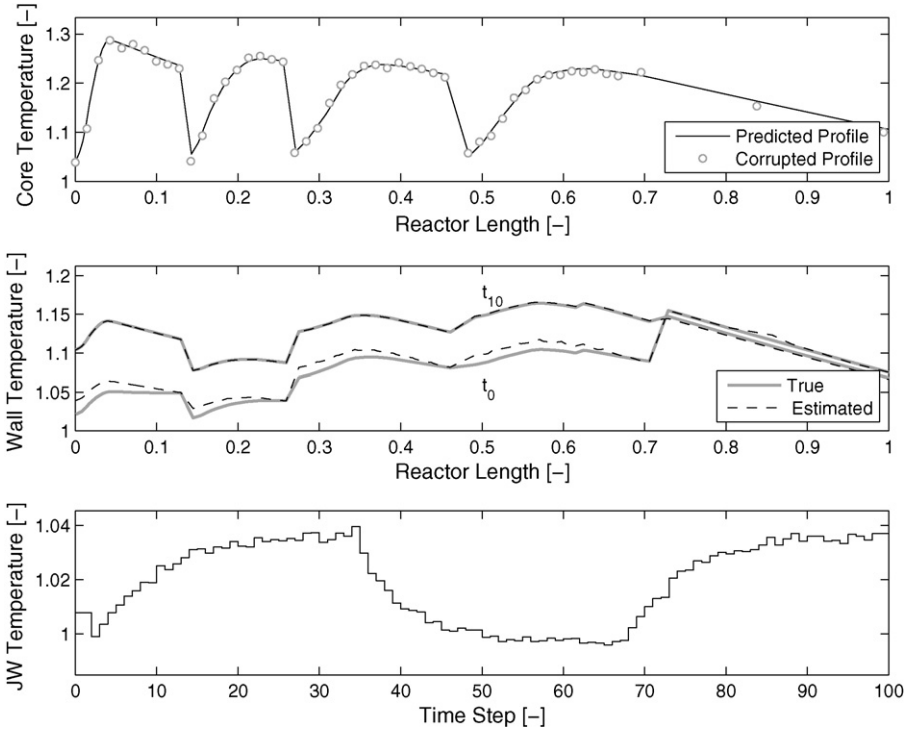

Fig. 9. Performance of coupled MHE-NMPC for output feedback. Predicted temperature profile at time step 25 and corresponding profile corrupted with Gaussian noise (top). Convergence of MHE estimator to true wall profile (middle). Jacket water inlet temperature of Zone 2 computed by NMPC controller (bottom).

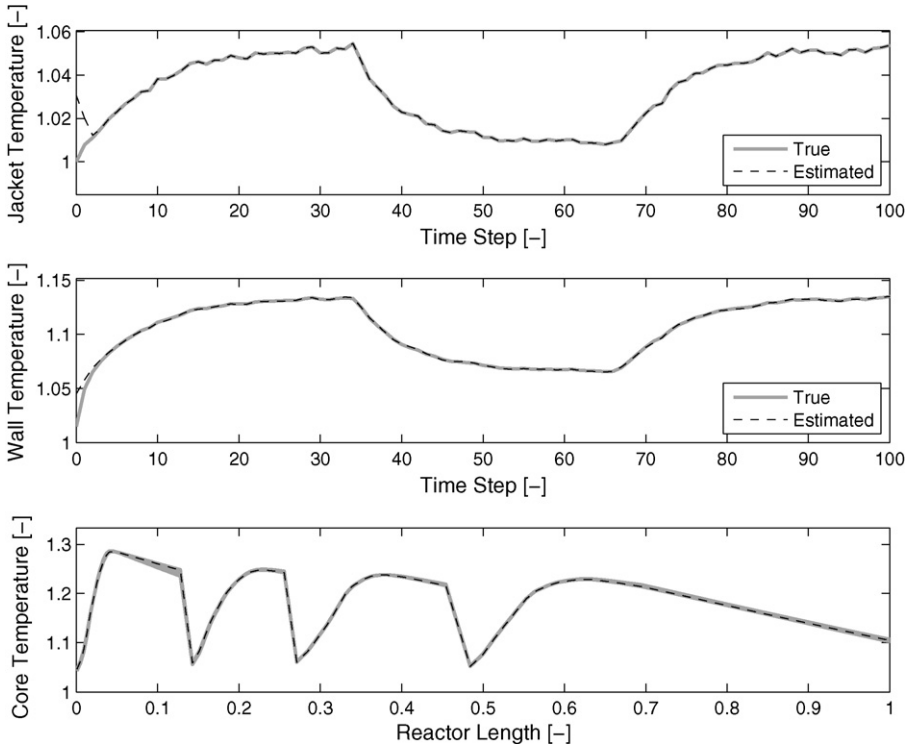

Fig. 10. Performance of coupled MHE-NMPC for output feedback. Convergence of MHE estimator to jacket temperature at a particular location of the reactor (top). Convergence of MHE estimator to wall temperature (middle). Closed-loop responses of NMPC for temperature profile (bottom).

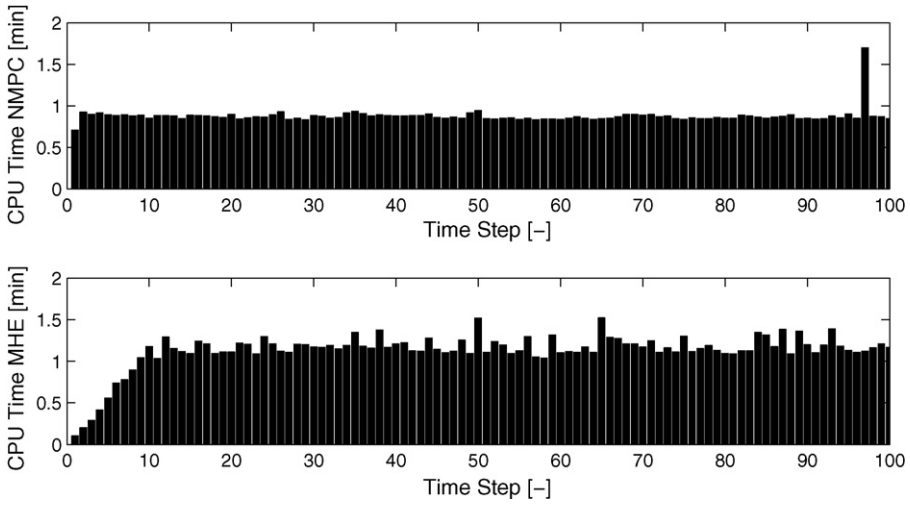

Fig. 11. Solution times for NMPC (top) and MHE problems (bottom) with horizons of $N = 10$ time steps.

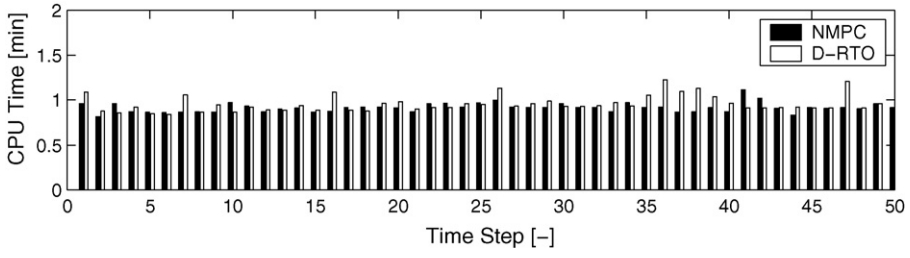

Fig. 12. Comparison of solution times for NMPC problem with tracking and economic objectives.

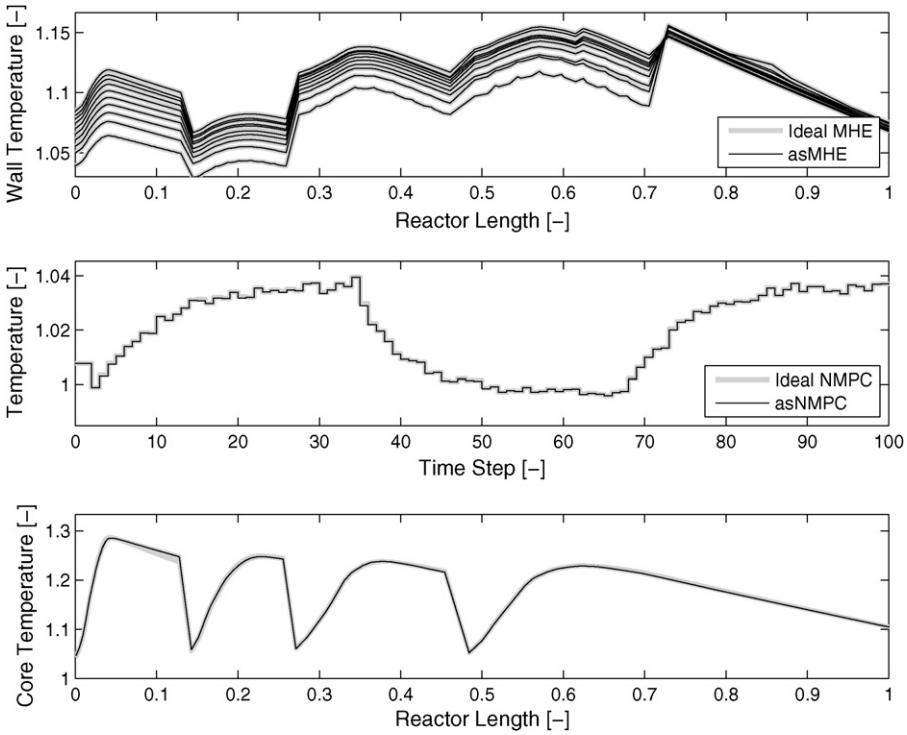

Fig. 13. Effect of NLP sensitivity errors on performance of advanced-step MHE and NMPC strategies in LDPE case study.

focus on fine-grained structures. The NMPC problem with $N = 30$ time steps, resulting in an NLP with 242,850 constraints and 1,110 degrees of freedom, is solved in around 2 min. The MHE problem with $N = 30$ time steps results in an NLP with 244,260 constraints and 1,878 degrees of freedom that can be solved in around 2.5 min. In Fig. 12 we illustrate the solution times for NMPC problems with tracking and economic objectives. We can see that if exact derivative information is used to solve the problems, the solution times (e.g. factorization time and number of iterations) are not altered by the choice of objective function.

4.3.3. Advanced-step strategies

From the previous computational results we can conclude that a full-discretization approach coupled with an sparsity-exploiting NLP solver results in an efficient computational strategy to solve the NMPC and MHE problems. In particular, we have demonstrated that this approach scales well with problem size and number of degrees of freedom. Nevertheless, it is clear that even if we have a fast strategy to solve the problems, the solution time will always become a bottleneck as we consider larger and larger applications. For instance, in the LDPE case study, the time required to solve the MHE problem to obtain the state estimate plus the time required to solve the NMPC problem to obtain the control actions is more than 2 minutes. This solution time is in fact longer than the assumed sampling time. This problem becomes more obvious as we consider longer time horizons. Here, we demonstrate that these limitations can be avoided with advanced-step strategies.

We demonstrate the performance of a coupled advanced-step MHE-NMPC strategy on the same output feedback scenario presented in the previous section. Again, a prediction horizon N of 10 time steps and sampling times of 2 min have been used. Here, we demonstrate that the approximation errors introduced by NLP sensitivity do not destabilize the controller. In this scenario, we recognize that since the plant response differs from that of the NMPC controller prediction and we introduce noise, the as-MHE estimator will see a difference between the measured and the predicted outputs and will correct on-line using NLP sensitivity. We have found that the approximation errors are negligible and the as-MHE estimator has almost identical convergence properties to that of the ideal MHE estimator. In the top graph of Fig. 13, we can see that the as-MHE estimates are identical to those of the ideal or optimal MHE estimator. Using the estimated states and HTCs, the as-NMPC controller then corrects the predicted state on-line. In the middle graph of Fig. 13 we present the closed-loop response of one of the jacket water inlet temperatures for the as-NMPC controller and of its ideal NMPC counterpart. As can be seen, both control actions are identical. In the bottom graph of this figure we can see that the as-NMPC controller is able to stabilize the temperature profile around the target. The sensitivity calculations for both the NMPC and MHE problems require less that 0.1 CPU seconds. The total background times required to converge the predicted MHE and NMPC problems and update the KKT matrices are very similar to those presented in Fig. 11. On the other hand, note that the advanced-step strategy allows to solve the background MHE and NMPC problems independently. In the case study analyzed, this allows the total background computation ($\leq$ 70 CPUs) to be less than the sampling time of 2 min. With this, we can conclude that the advanced-step MHE-NMPC strategies allow us to consider implementations with long time horizons, reasonable sampling times, and negligible feedback delays.

5. Conclusions and future work

In this work, we incorporate a detailed LDPE tubular reactor model in a general NMPC framework. Using a particular fouling–defouling scenario, we demonstrate that a tracking NMPC controller is able to stabilize the reactor in the face of persistent fouling disturbances. We also demonstrate that if an economic objective function is added, the NMPC controller can exploit the multiple degrees of freedom of the process in order to simultaneously optimize the process profitability. We have coupled an MHE estimator to the NMPC controller to provide output feedback to the

process. We have demonstrated that a full-discretization approach coupled to a sparsity-exploiting interior point solver results in an fast strategy to solve both the NMPC and MHE problems. Finally, we implement advanced-step NMPC and MHE strategies to avoid feedback delays. We demonstrate that a highly sophisticated dynamic model can be accommodated on-line to provide nearly instantaneous feedback to the LDPE process with negligible approximation errors.

As part of future work, different extensions to the NMPC framework can be considered. For instance, structural model mismatch and gross errors can be handled within the MHE formulation. The controller performance can also be tested in complex grade transition scenarios in which the entire temperature profile needs to be moved. More general objective functions can also be considered in the controller. In particular, energy costs can become relevant in LDPE processes if the side-stream temperatures need to be cooled down significantly. We also believe that the use of an economics-oriented controller can help identify design bottlenecks and guide retrofitting tasks. Finally, the fouling problem has been treated so far as an uncertain disturbance. Instead, the fouling rate could potentially be manipulated through appropriate control actions (e.g. thermal shocks, pressure shocks). In order to do this, predictive fouling models are required. With this, the NMPC controller can also be used as a long-term planner of fouling/defouling operations that can be interfaced with a lower-level controller.

References

Alessandri, A., Baglietto, M., & Battistelli, G. (2008). Moving-horizon state estimation for nonlinear discrete-time systems: new stability results and approximation schemes. *Automatica*, *44*, 1753–1765.

Bokis, C. P. (2001). Physical properties, reactor modeling, and polymerization kinetics in the low-density polyethylene tubular reactor process. *Industrial & Engineering Chemistry Research*, *41*, 1017–1030.

Brandolin, A., Lacunza, P., Ugrin, L., & Capiati, N. J. (1996). High-pressure polymerization of ethylene and improved mathematical model for industrial tubular reactors. *Polymer Reaction Engineering*, *4*, 193–241.

Buchelli, A. (2005). Modeling fouling effects in LDPE tubular polymerization reactors. 1: Fouling thickness determination. *Industrial & Engineering Chemistry Research*, *44*, 1474–1479.

Diehl, M., Bock, H. G., & Schlöder, J. P. (2005). A real-time iteration scheme for nonlinear optimization in optimal feedback control. *SIAM Journal of Control and Optimization*, *43*, 1714–1736.

Duff, I. S. (2004). MA57—A code for the solution of sparse symmetric definite and indefinite systems. *ACM Transactions on Mathematical Software*, *30*, 118–144.

Engell, S. (2007). Feedback control for optimal process operation. *Journal of Process Control*, *17*, 203–219.

Fourer, R., Gay, D. M., & Kernighan, B. W. (1992). *AMPL: A Modeling Language for Mathematical Programming*. Belomont, CA: Duxbury Press.

Goto, S. (1981). Computer model for commercial high-pressure polyethylene reactor based on elementary reaction rates obtained experimentally. *Journal of Applied Polymer Science*, *36*, 21–40.

Häfele, M., Kienle, A., Boll, M., & Schmidt, C. U. (2006). Modeling and analysis of a plant for the production of low density polyethylene. *Computers and Chemical Engineering*, *31*, 51–65.

Helbig, A., Abel, O., & Marquardt, W. (2000). Structural concepts for optimization based control of transient processes. In *Nonlinear Model Predictive Control*. Basel: Birkhaüser-Verlag., pp. 295–312.

Huesman, A. E. M., Bosgra, O. H., & Van den Hof, P. M. J. (2007). Degrees of freedom analysis of economic dynamic optimal plantwide operation. In *Proceedings of 8th International Symposium on Dynamics and Control of Process Systems* Cancun, Mexico.

Kadam, J., Marquardt, W., Schlegel, M., Backx, T., Bosgra, O., & Brouwer, P. J. (2003). Towards integrated dynamic real-time optimization and control of industrial processes. In *Proceedings Foundations of Computer-Aided Process Operations* Cambridge, USA.

Karypis, G., & Kumar, V. (1999). A fast and high quality multilevel scheme for partitioning irregular graphs. *The SIAM Journal on Scientific Computing*, *20*, 359–392.

Kim, D. M., & Iedema, P. D. (2004). Molecular weight distribution in low-density polyethylene polymerization; impact of scission mechanisms in the case of a tubular reactor. *Chemical Engineering Science*, *59*, 2039–2052.

Kiparissides, C., Verros, G., Pertsinidis, A., & Goosens, I. (1996). On-line parameter estimation in a high-pressure low-density polyethylene tubular reactor. *AICHE J.*, *42*, 440.

Kiparissides, C., Verros, G., & McGregor, J. (1993). Mathematical modeling, optimization and quality control of high-pressure ethylene polymerization reactors. *Journal of Macromolecular Science—Reviews in Macromolecular Chemistry and Physics*, *C33*, 437–527.

Kiparissides, C., et al. (2005). Mathematical modeling of free-radical ethylene copolymerization in high-pressure tubular reactors. *Industrial & Engineering Chemistry Research*, *44*, 2592–2605.

Knuuttila, H., Lehtinen, A., & Nummila-Pakarinen, A. (2004). Advanced polyethylene technologies: Controlled material properties. *Advanced Polymer Science*, *169*, 13–27.

Luft, G. (1977). Effectiveness of organic peroxide initiators in the high-pressure polymerization of ethylene. *Journal of Macromolecular Science: Chemistry*, *A11*, 1089–1112.

Marlin, T. E., & Hrymak, A. N. (1996). Real-time operations optimization of continuous processes. In *Fifth International Conference on Chemical Process Control (CPC-5)*.

Pertsinidis, A., Papadopoulos, E., & Kiparissides, C. (1996). Modeling and analysis of a plant for the production of low density polyethylene. *Computers and Chemical Engineering*, *20*, 449–S454.

Rawlings, J. B., & Bakshi, B. R. (2006). Particle filtering and moving horizon estimation. *Computers and Chemical Engineering*, *30*, 1529–1541.

Skogestad, S. (2000). Self-optimizing control: The missing link between steady-state optimization and control. *Computers and Chemical Engineering*, *24*, 569–575.

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

Yip, W. S., & Marlin, T. E. (2004). The effect of model fidelity on real-time optimization performance. *Computers and Chemical Engineering*, *28*, 267–280.

Zabisky, R. C. M. l., Chan, W. M., Gloor, P. E., & Hamielec, A. E. (1992). A kinetic model for olefin polymerization in high-pressure tubular reactors. A review and update. *Polymer*, *33*, 2243–2262.

Zavala, V. M., Laird, C. D., & Biegler, L. T. (2008). Fast implementations and rigorous models: Can both be accommodated in NMPC? *International Journal of Robust and Nonlinear Control*, *18*, 800–815.

Zavala, V. M., & Biegler, L. T. (2006). Large-scale parameter estimation in low-density polyethylene tubular reactors. *Industrial & Engineering Chemistry Research*, *25*, 7867–7881.

Zavala, V. M., & Biegler, L. T. (2008). Large-scale nonlinear programming strategies for the operation of low-density polyethylene tubular reactors. In *Proceedings of ESCAPE 18* Lyon, France.

Zavala, V. M., & Biegler, L. T. (2008). Nonlinear programming strategies for state estimation and model predictive control. In *Proceedings of Assessment and Future Directions of Nonlinear Model Predictive Control* Pavia, Italy.

Zavala, V. M., & Biegler, L. T. (2009). Optimization-based strategies for the operation of low-density polyethylene tubular reactors: moving horizon estimation. *Computers and Chemical Engineering*, *33*, 379–390.

Zavala, V. M., Laird, C. D., & Biegler, L. T. (2008). A fast moving horizon estimation algorithm based on nonlinear programming sensitivity. *Journal of Process Control*, *18*, 876–884.

Zavala, V. M., & Biegler, L. T. (2009). The advanced step NMPC controller: stability, optimality and robustness. *Automatica*, *45*, 86–93.