12th IFAC International Symposium on
Advanced Control of Chemical Processes
July 14-17, 2024. Toronto, Canada

Economic Model Predictive Control for Cryogenic Air Separation Unit Startup

Anthony W. K. Quarshie, José Matias, Christopher L. E. Swartz [1]

[*] *Department of Chemical Engineering, McMaster University, 1280 Main Street West, Hamilton, Ontario, Canada, L8S 4L7*

Abstract: Current energy market trends incentivize frequent optimal load changes, including the startup of cryogenic air separation units (ASUs), which are large electricity consumers. In this study, we assess the potential benefits of using an economic nonlinear model predictive control (ENMPC) framework for the optimal startup of ASUs in the presence of a measured process disturbance. We also considered strategies for improving the solution computational speed of the ENMPC problem. Our case study shows substantial profit recovery by the control strategy relative to offline pre-computed optimal inputs in response to the disturbance.

Keywords: Disturbance rejection, optimization-based control, discontinuity, smoothing

1. INTRODUCTION

Cryogenic air separation units (ASUs) are energy-intensive and produce some combination of high-purity liquid and gaseous products of nitrogen, oxygen, and argon from air using distillation processes. The products from ASUs find uses in various sectors, including steel manufacturing, semiconductor production, refining, pharmaceuticals, and aeronautical applications. Air being the only feed to the process, the operating cost of ASUs is mostly electricity due to multi-stage compression.

In the current volatile electricity market, an economically advantageous strategy for ASUs may be to shut down during periods of high electricity pricing, especially if the startup operation is fast or economically optimal (Miller et al., 2008a). Such a strategy would result in more startups than before. Meanwhile, the status quo of industrial startup of ASUs is limited automation (Caspari et al., 2020), long startup time (a few hours to a day (Miller et al., 2008b)), and limited revenue generation, motivating the study of ASU startups. The startup phase presents complex challenges due to the intricate interplay of thermodynamics, process dynamics, and discontinuous behavior. Although ASUs vary in size: from the single-column nitrogen plant to the multiproduct ASU with four distillation columns, they all share the aforementioned complexities. As such, much effort in the literature on ASU startup has been directed toward modeling and dynamic simulation of the startup operation. Miller et al. (2008b) present a dynamic simulation study of the multiproduct ASU, in which startup discontinuities are captured using nonsmooth formulations. The resulting high-fidelity model demonstrates the startup time reduction potential of reintroducing liquid collected from shutdown. Contributions by Kender et al. (2019, 2021) focused on startup model development for double-column and multiproduct ASUs using a pressure-driven approach. The digital twin obtained simulates the entire operating range of the ASU, from warm startup (startup operation initialized at ambient temperatures instead of cryogenic temperatures) to shutdown, and is used to perform hazard analysis. More recently, Caspari et al. (2020) and Quarshie et al. (2023b) presented dynamic optimization (DO) studies on a nitrogen plant with startup discontinuities modeled using a smoothed Fischer-Burmister formulation, and the multiproduct ASU with discontinuities modeled using smooth sigmoidal functions, respectively. Both studies use a quadratic target-tracking objective function, which minimizes deviation from targets, such as product flows from their steady-state values. In a subsequent contribution by Quarshie et al. (2023a), readily interpretable metrics such as startup time and profit are used as objective functions in the DO of the multiproduct ASU startup.

This study builds on the work of Quarshie et al. (2023a). While quantifiable improvements in time and profits have been demonstrated using offline open-loop control strategies, these gains may not be fully realized due to model uncertainty and disturbances. As shown by Schäfer et al. (2019), feedback control through a sub-optimal ENMPC offers superior performance in uncertainty mitigation relative to an offline open-loop control strategy for a nitrogen plant. As such, the main objective of this study is to evaluate the potential benefit of ENMPC for ASU startups in the presence of disturbances. The evaluation uses a startup model of the nitrogen plant developed using formulations presented in our earlier contribution (Quarshie et al., 2023b). We use an ENMPC formulation for which we apply the so-called shift initialization strategy (Diehl et al., 2009) for rolling horizon formulations for improved computational performance.

The remainder of this paper is organized as follows. Section 2 describes the nitrogen plant and summarizes the corresponding startup model used in this study. We then provide details on our formulation and implementation in Section 3. Assessment of the formulation and implementa-

[1] Corresponding author (e-mail: swartzc@mcmaster.ca)

tion through the case study is presented in Section 4, and we recap key aspects of our study and provide directions for future research in Section 5.

2. STARTUP MODEL DESCRIPTION

2.1 Nitrogen plant process description

Figure 1 represents the nitrogen plant. The process begins with the compression of filtered air before the cooling and removal of water-soluble impurities and other impurities, such as hydrocarbons, in the molecular sieve units.

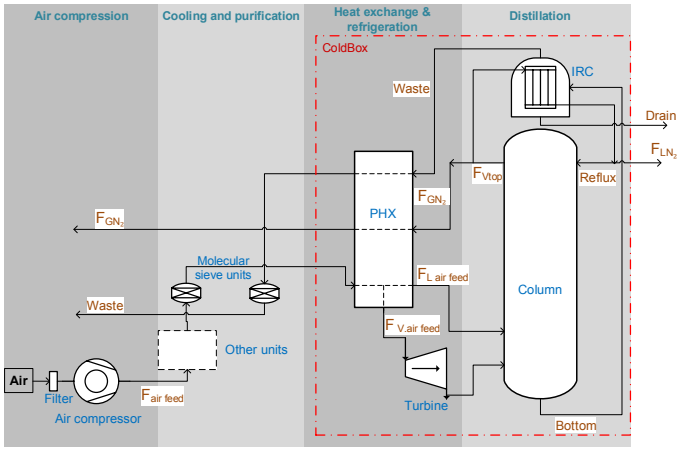

Fig. 1. Process diagram for a cryogenic air separation plant that produces gaseous and liquid nitrogen (Adapted from Cao et al. (2016b) and Linde Engineering (2019)). PHX = Primary Heat Exchanger; IRC = Integrated Reboiler Condenser. (We only modeled units in the cold-box, the section of the plant demarcated with red dashes, in our study)

The purified air ($F_{\text{air feed}}$) enters the PHX warm end, exchanging heat with the countercurrent cryogenic product nitrogen and waste streams. At an intermediate position through the PHX, a large portion of the air stream is withdrawn as turbine air ($F_{\text{V air feed}}$), which is expanded in the turbine to provide refrigeration to the plant. The expanded air is sent to the bottom of the distillation column as a vapor stream. The other portion of air goes through the entire length of the PHX, exiting as a liquid ($F_{\text{L air feed}}$), and is fed to the column above the vapor feed.

Rectification of air occurs in the distillation column around 6.6 bar (Caspari et al., 2018) and -180 degrees Celsius (Cao et al., 2016a). An oxygen-rich liquid stream is routed from the column sump (Bottom) to the reboiler side of the integrated reboiler condenser (IRC), while the high-purity nitrogen vapor stream from the top of the column (V_{top}) is split, with one portion going towards the PHX as gaseous nitrogen product (F_{GN_2}) and the remainder to the condenser side of the IRC. In the IRC, the nitrogen vapor condenses against the boiling oxygen-rich liquid on the reboiler side. A fraction of the condensed liquid nitrogen can be drawn off as a liquid nitrogen product(F_{LN_2}) while the rest of the liquid is sent to the column as reflux. The boil-off generated on the reboiler side is routed toward the PHX as the waste stream, which, together with the gaseous nitrogen product, provides cooling in the PHX. The waste stream is also used to regenerate the molecular sieve units.

2.2 Modeling process units

The distillation column modeling approach used in this study is consistent with past studies on cryogenic ASUs (Miller et al., 2008b; Cao et al., 2015, 2016a), where tray-by-tray modeling is considered. The approach mirrors our earlier work on the multiproduct ASU (Quarshie et al., 2023b), which has multiple columns. The distinguishing features of the column in this study are that it is entirely made up of trays and does not use the collocation model reduction technique discussed in our earlier work. Noteworthy characteristics of the column include dynamic material balance, pseudo steady-state energy balance, variable stage pressure drop, non-ideal thermodynamics, and heat leakage terms.

The IRC is a composite model of the reboiler and condenser sides, modeled using an approach that mirrors that of Cao et al. (2015). The reboiler side is characterized by dynamic material balance and pseudo-steady-state energy balance and is modeled after the column trays. The liquid level in the reboiler is controlled using a PI controller.

The PHX is a multi-stream heat exchanger in which phase change occurs. We use the approach in Cao et al. (2015) by which zones are predetermined for various phases. Specifically, the approach uses two zones, with vapor air, the $F_{Vairfeed}$ in Figure 1, being drawn at the boundary between the two zones and phase change occurring in the second zone with a liquid air stream at the outlet.

We use the approach presented in Cao et al. (2015) for the turbine. In Cao et al. (2015), the turbine is modeled as a system of algebraic equations, which define polytropic efficiency and correlations for discharge pressure, turbine speed, and stream flow rate obtained from plant data.

2.3 Startup discontinuity modelling

Simulation of startup behavior is accomplished by capturing discontinuities using smoothing functions such as the hyperbolic tangent function. These discontinuities are associated with three of the four categories of discontinuities presented in our earlier work (Quarshie et al., 2023b), namely:

- Liquid flow discontinuity: This involves transitioning from zero flow, when tray liquid height is below the top of the weir, to nonzero flow computed using the Francis weir equation when the liquid level is above the weir.
- Vapor flow discontinuity: This entails a change in vapor flow dynamics from when the vapor from a lower stage escapes through the unblocked downcomer to reach the given tray to when the downcomer is blocked, and vapor flow is determined from an energy balance around the given stage.
- Flow out of sump and reboiler: This describes the switching between manual control, for which the flow rate is directly specified, and automatic control, for which the flow rate is determined by a controller, for flow out of these receptacles.

2.4 Other modeling aspects

Modeling aspects not mentioned above but critical to the study are presented below:

- Product flow computation: products from the process, liquid, and gaseous nitrogen, have purity thresholds that are met later in the startup. As such, product flow is considered zero early on when the threshold is not satisfied, even though the flow in the product stream may be nonzero. To model this "switch-on" behavior, we use the smoothed switching function formulation proposed in our earlier work (Quarshie et al., 2023a) to which the reader is referred.
- Logarithmic transformation of components: logarithmic transformation is applied to mole fractions of the components to improve numerical robustness. We use the transformations proposed in our earlier work (Quarshie et al., 2023a).

3. FORMULATION AND IMPLEMENTATION

The resulting model is an index-1 differential algebraic equation (DAE) system with 330 differential and 3,440 algebraic variables. We carry out the implementation in Python using the CasADi (Andersson et al., 2018) front-end package on a DELL XPS 8940 with 16GB RAM and an Intel Core i9-10900 processor running MS Windows. We use the optimization-based economic NMPC as our approach. Following are details on the ASU startup optimization, the control strategy setup, the ENMPC optimization formulation, and the ENMPC implementation, under which we discuss strategies applied for improved solution speed.

3.1 ASU startup optimization

ASU startup optimization involves taking the ASU at a defined initial point, manipulating process inputs to take the process to a defined steady state, improving a metric (objective function) along the way, and ensuring process constraints are not violated. In this study, we consider an economic objective function for the underlying dynamic optimization. We base our study on cold startup, when the plant is initially at cryogenic temperatures.

The economic objective function is cumulative profit over the startup horizon and is obtained by integrating the difference between the total instantaneous revenues and costs. This is illustrated by the function below:

$$\mathcal{J} = \int_0^{t_f} \left(\sum_{x \in X} P_x.F_{prod\,x} - C_{elec} \right) \mathrm{dt}, \tag{1}$$

where P_x is the unit price for a given product x, $F_{prod\,x}$ is the flow rate of the on-spec product computed using the formulation described under Section 2.4, C_{elec} is the electricity cost associated with air compression (Cao et al., 2017), X represents the set of products, GN_2 and LN_2 which are gaseous nitrogen and liquid nitrogen products respectively. The manipulated inputs of our system are listed in Table 1. The first four inputs are self-explanatory and can be identified in Figure 1. The fifth input: $R_{gasdraw}$ is the fraction of the vapor leaving the top of the column sent towards the PHX; an identical input is used in the nitrogen plant study in Cao et al. (2015). This vapor flow traverses the same path as the F_{GN_2}. Only one of F_{GN_2} or $R_{gasdraw}$ is active at a time. We found that the system was better controlled when the $R_{gasdraw}$ was used at the early stages of startup and the F_{GN_2} more suitable as steady state is approached. The choice between the two inputs is based on the composition of the top-stage vapor, where the logical condition is modeled using a hyperbolic tangent function. Mathematically, this is given by:

$$F_{GN_2}^{PL} = 0.5\Big(1 + \tanh\big[\gamma(y_{top}^{th} - y_{top})\big]\Big) R_{g.draw}.F_{V.top} + 0.5\Big(1 - \tanh\big[\gamma(y_{top}^{th} - y_{top})\big]\Big) F_{\mathrm{GN}_2} \tag{2}$$

where $F_{GN_2}^{PL}$ is the vapor flow through the GN_2 product line, $F_{V.top}$ is the vapor flow from the top of the column, γ is a parameter that controls the steepness of the hyperbolic tangent function, y_{top} and y_{top}^{th} denote the composition of the column's top tray vapor and the threshold composition value at which active input selection switches from $R_{g.draw}$ to F_{GN_2} respectively. According to Equation 2, when the composition, y_{top}, is below the threshold value, y_{top}^{th}, $F_{GN_2}^{PL}$ is the product $R_{g.draw}.F_{V.top}$, whereas when y_{top} is greater than y_{top}^{th}, $F_{GN_2}^{PL}$ equals F_{GN_2}. The final values of γ and y_{top}^{th} were obtained after numerous simulation experiments.

Table 1. Input variables

Variable	Symbol
Air feed to compressor	$F_{airfeed}$
Liquid air feed	$F_{L.airfeed}$
GN_2 rate	F_{GN_2}
LN_2 rate	F_{LN_2}
Top tray gas draw fraction	$R_{g.draw}$

Table 2 lists the constraint variables that must be tracked in the optimization problem. All these constraints are path constraints. Except for the product stream composition, all the variables are constrained in the optimization formulation. With startup, the product purity threshold, measured in parts per million of oxygen (PPMO), is attained towards the end of startup, corresponding to the steady state. Given that product generation is contingent on meeting product purity, showing product flow towards the end of startup implies that the product stream composition satisfies the product purity threshold towards the end of startup.

Table 2. Constraint related variables

Constraint	Type
Top tray flood percent	path
Reboiler liquid level	path
Sump liquid level	path
Turbine outlet dew point	path
Turbine inlet dew point	path
Liquid air feed bubble point	path
Product stream composition	path

3.2 Control setup

In this subsection, the control setup for measured disturbance rejection is discussed. This setup is illustrated in Figure 2. We do not consider plant-model mismatch, and assume full state feedback. The ENMPC serves the purpose of counteracting the impact of the measured distur-

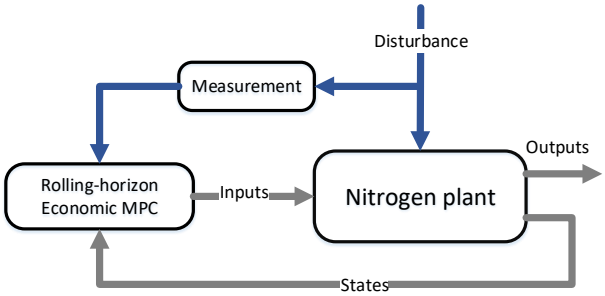

Fig. 2. Control setup for measured disturbance rejection

bance on the plant. With the current states and measured disturbance, the rolling horizon ENMPC computes over the control horizon, and the inputs corresponding to the first input interval are sent to the plant for implementation.

3.3 Optimization formulation and implementation

We apply the simultaneous solution approach for DO by which both states and inputs are discretized, effectively converting the problem into a nonlinear programming (NLP) problem. For the discretizations, we use zero-order hold for the inputs with the finite element (FE) length of 0.36 hrs and backward Euler for the states by equally subdividing the FE into 19 parts. The NLP is solved using IPOPT (Wächter and Biegler, 2006), an interior point solver. Given that the performance of IPOPT is dependent on the linear solver used (Tasseff et al., 2019), we experimented with various linear solvers from the Harwell subroutine library (HSL, 2023) and selected MA97 in the process. One of the challenges of the simultaneous solution DO approach is the requirement of a good initial guess. To overcome this hurdle, we generate the initial guess using backward Euler integration over the startup horizon. However, the backward Euler integration also requires an initial guess, which we obtain by integrating the startup model over the startup horizon using the SUNDIALS IDAS integrator (Hindmarsh et al., 2005).

3.4 Rolling horizon economic model predictive control

Startup of the ASU is a highly nonlinear process, spans a wide operating range, and has multiple constraints (Table 2). These attributes make the economic NMPC formulation a suitable choice. Startup is often carried out to drive the process to a pre-determined steady state, implying a defined endpoint. As we have done in our past studies on startup (Quarshie et al., 2023b,a), we fix the inputs of the last control interval to reach the pre-selected steady state in this study. This feature gives ASU startup a batch process character. Conventionally, economic NMPC for batch processes is implemented using a shrinking horizon formulation (Nagy and Braatz, 2003) for which both the control and prediction horizons shrink as the controller marches toward the endpoint. Nonetheless, we implement a rolling horizon approach for which the control and prediction horizon lengths are maintained as the controller marches along.

We select the rolling horizon approach over the shrinking horizon approach to eliminate pre-optimization problem formulation steps that are time-consuming. Solving an NLP in CasADi first requires building an NLP formulation, which, given our problem size, takes a considerable amount of time. With the shrinking horizon formulation, building a different NLP is required each time inputs have to be computed due to the problem size changing from one step to the next. However, the rolling horizon approach requires just a single instance of NLP creation, and the NLP is reused when inputs have to be computed. In essence, the rolling horizon and shrinking horizon approaches represent tradeoffs between a single, although large, NLP build; and an NLP rebuild, shrinking in dimension, at each EMPC iteration, respectively. In our case, it is more advantageous to build the NLP once and solve a large optimization problem at each iteration.

The optimization horizon is 20 input intervals long. For the first ENMPC execution, the last interval has inputs set to the steady-state values. This means that for the first ENMPC execution, the determination of optimal inputs is limited to only the first 19 input intervals. As the controller marches on, for each additional step, the number of input intervals for which inputs are set to their steady-state values increases by one, implying the number of intervals for which inputs have to be determined decreases by one (this is illustrated by Figure 3). Although the shrinking of the input space resembles the shrinking horizon formulation, it must be noted that the number of states to be determined is maintained from one step to the next. To solve the

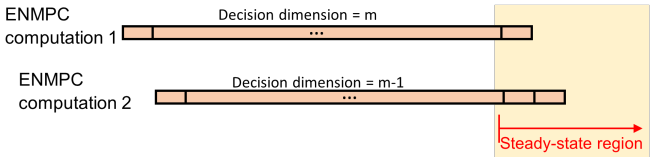

Fig. 3. Control horizon for ENMPC computations, illustrating change in input decisions demension

ENMPC optimization problem, we employ the simultaneous DO solution approach as described under Section 3.3. With input and state discretization as described early on, the resulting NLP has more than 1.8 million variables: states and inputs combined. Given the nonlinearity of the startup model and dimensionality of the NLP, this makes for a complex problem to solve with a nontrivial consequence to the computational time. Thus, we use the so-called shift initialization, a warm start strategy (Diehl et al., 2009). The warm start strategy leverages the similarity of neighboring NLPs for efficient initialization of a subsequent problem using primal-dual information from the preceding NLP solve. The shift initialization for the rolling horizon approach is illustrated in Figure 4. To use the primal-dual information from a preceding NLP solve, the information corresponding to the first input interval is discarded while that for the last interval is duplicated. As will be seen in the next section, using the shift initialization strategy resulted in a good computational performance. This is attributed, in part, to the long horizon used, as using this strategy with a short horizon length did not lead to a superior performance relative to the conventional warm start strategy (Diehl et al., 2009).

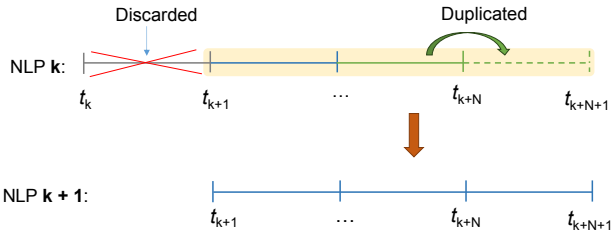

Fig. 4. Shift initialization for rolling horizon formulation

4. CASE STUDY: DISTURBANCE REJECTION

In this subsection, we address the problem of measured disturbance rejection using the ENMPC described in Section 3.4. The disturbance being considered is motivated by the scenario in which the air vent valve at the compressor discharge is inadvertently opened. The impact of this switch is that the cold-box does not receive the full flow of air from the compressor. We limit the impact of this disturbance to the change in the total air flow rate only. This disturbance is modeled as the fraction of the supposed air flow rate from the compressor that enters the PHX. This is illustrated in Figure 5, where the total fraction of air leaving the compressor enters the PHX except at the time period between 4hr and 6hr, during which the fraction is 0.94. However, the ENMPC "sees" these changes through the measurement in real-time and is not "aware" of the time and duration of the change ahead of time. The ENMPC control interval in this study is 0.36 hrs. To assess the performance of the ENMPC in rejecting

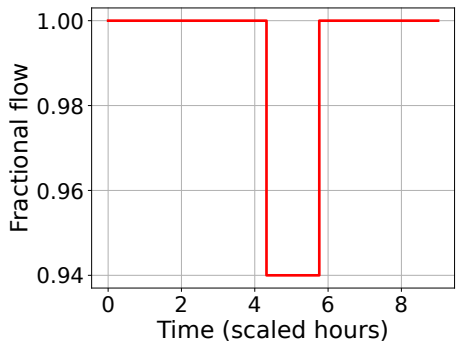

Fig. 5. Fractional change in air flow due to disturbance

this disturbance, we apply offline pre-computed open-loop optimal base case (BC) inputs to the plant along with the disturbance (BC u+dist). As such, while the air flow to the PHX dips between 4hr and 6hr, the cost of electricity due to compression should remain the same relative to the BC because the same amount of air is compressed. However, revenues, and in turn profit, should be affected given that less feed is going into the process.

Table 3. Cumulative profit performance assessment of ENMPC for disturbance rejection.

Execution type	Scaled profits	% change relative to BC
BC	51.906	-
BC u +dist	44.536	-14.2
ENMPC	48.972	-5.6

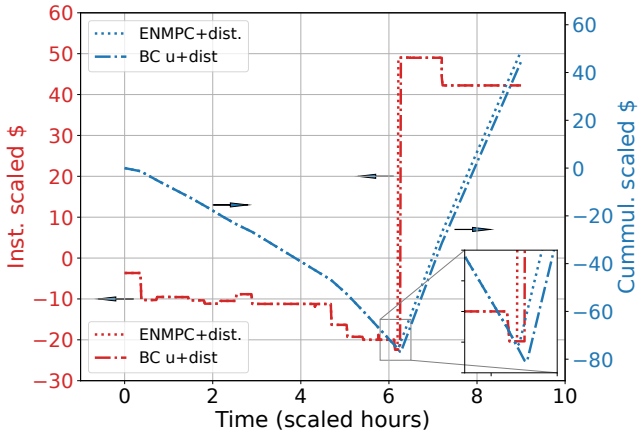

Fig. 6. Instantaneous and cumulative profits from using base case inputs with disturbance (BC u+dist) and ENMPC computed inputs in response to disturbance. (The arrows point to the reference axis)

Table 3 shows a summary of the performance of the ENMPC. Whereas the disturbance results in losses in both instances, the ENMPC significantly minimizes the profit loss relative to using the BC inputs. The ENMPC reduces profit loss by 9%, from 14% per the table. Figure 6 sheds more light on the comparison by showing cumulative and instantaneous profit trajectories. Up to just after hour 6, the two trajectories coincide perfectly, with losses being incurred as costs are more dominant relative to revenues. Beyond that point, the ENMPC trajectory begins to recover first from the loss, even as instantaneous profit is positive and the cumulative profit begins an upward trend.

Computational statistics are presented in Table 4. The cycles correspond with control horizon optimizations for which the CPU secs and the number of NLP iterations are reported. The measured disturbance occurs across cycles 12 to 15, thus the highlight. Outside the highlighted parts of the table, the number of NLP iterations is mostly 1, lasting about 200 CPU secs. This low number of iterations indicates the efficacy of the shift initialization warm start strategy. However, cycle 14 shows high CPU seconds and NLP iteration counts. This long solution time is attributed, in part, to the disturbance changing model states such that the preceding solutions are no longer suitable, requiring the controller to perform more iterations, which takes a longer time.

5. CONCLUSION

In this work, we set out to assess the potential benefit of a control strategy beyond an offline optimal open-loop strategy for disturbance rejection in ASU startup. For this purpose, we proposed a control framework, leveraging an ENMPC formulation to which we applied the shift-initialization strategy for improved computational speed. The proposed framework showed significant profit recovery relative to the use of pre-computed optimal inputs, indicating that ASU startup operations stand to benefit economically from the application of ENMPC. The computational statistics demonstrate the efficacy of the strategies adopted for improved computational speed, although the framework, as is, has limitations for online implementation due to the long computational time reported for one of

Table 4. Computational statics of the ENMPC

Cycle	CPU sec	NLP Iter.
0	1460	8
1	220	1
2	242	1
3	222	1
4	219	1
5	330	2
6	318	2
7	216	1
8	223	1
9	221	1
10	219	1
11	178	1
12	228	1
13	207	1
14	7990	30
15	248	1
16	231	1
17	177	1
18	516	2
19	148	0

the cycles. As such, an obvious future research direction is the exploration of strategies to further improve solution time. Other potential directions include the consideration of plant-model mismatch, due to parameter uncertainty, which may require computationally efficient parameter estimation, and the consideration of uncertainty external to the process, such as electricity price uncertainty.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge support from the Natural Sciences and Engineering Research Council of Canada Grant RGPIN-2017-05627, the McMaster Advanced Control Consortium (MACC), and Linde plc.

REFERENCES

Andersson, J.A.E., Gillis, J., Horn, G., Rawlings, J.B., and Diehl, M. (2018). CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*.

Cao, Y., Swartz, C.L.E., Baldea, M., and Blouin, S. (2015). Optimization-based assessment of design limitations to air separation plant agility in demand response scenarios. *Journal of Process Control*, 33, 37–48.

Cao, Y., Swartz, C.L.E., and Flores-Cerrillo, J. (2017). Preemptive dynamic operation of cryogenic air separation units. *AIChE Journal*, 63(9), 3845–3859.

Cao, Y., Swartz, C.L.E., Flores-Cerrillo, J., and Ma, J. (2016a). Dynamic modeling and collocation-based model reduction of cryogenic air separation units. *AIChE Journal*, 62(5), 1602–1615.

Cao, Y., Swartz, C.L.E., and Flores-Cerrillo, J. (2016b). Optimal dynamic operation of a high purity air separation plant under varying market conditions. *Industrial & Engineering Chemistry Research*, 55(37), 9956–9970.

Caspari, A., Fahr, S.R., Offermanns, C., Mhamdi, A., Biegler, L.T., and Mitsos, A. (2020). Optimal start-up of air separation processes using dynamic optimization with complementarity constraints. *Computer Aided Chemical Engineering*, 48, 1–8.

Caspari, A., Faust, J.M.M., Schäfer, P., Mhamdi, A., and Mitsos, A. (2018). Economic nonlinear model predictive control for flexible operation of air separation units. *IFAC-PapersOnLine*, 51(20), 295–300.

Diehl, M., Ferreau, H.J., and Haverbeke, N. (2009). *Efficient Numerical Methods for Nonlinear MPC and Moving Horizon Estimation*, 391–417. Springer Berlin Heidelberg, Berlin, Heidelberg.

Hindmarsh, A.C., Brown, P.N., Grant, K.E., Lee, S.L., Serban, R., Shumaker, D.E., and Woodward, C.S. (2005). SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Transactions on Mathematical Software (TOMS)*, 31(3), 363–396.

HSL (2023). A collection of Fortran codes for large-scale scientific computation. URL `http://www.hsl.rl.ac.uk/`.

Kender, R., Kaufmann, F., Rößler, F., Wunderlich, B., Golubev, D., Thomas, I., Ecker, A.M., Rehfeldt, S., and Klein, H. (2021). Development of a digital twin for a flexible air separation unit using a pressure-driven simulation approach. *Computers & Chemical Engineering*, 151, 107349.

Kender, R., Wunderlich, B., Thomas, I., Peschel, A., Rehfeldt, S., and Klein, H. (2019). Pressure-driven dynamic simulation of start up and shutdown procedures of distillation columns in air separation units. *Chemical Engineering Research and Design*, 147, 98–112.

Linde Engineering (2019). Air separation plants: History and technological progress in the course of time.

Miller, J., Luyben, W.L., and Blouin, S. (2008a). Economic incentive for intermittent operation of air separation plants with variable power costs. *Industrial & engineering chemistry research*, 47(4), 1132–1139.

Miller, J.J., Luyben, W.L., Belanger, P., Blouin, S., and Megan, L. (2008b). Improving agility of cryogenic air separation plants. *Industrial & Engineering Chemistry Research*, 47(2), 394–404.

Nagy, Z.K. and Braatz, R.D. (2003). Robust nonlinear model predictive control of batch processes. *AIChE Journal*, 49(7), 1776–1786.

Quarshie, A.W.K., Swartz, C.L.E., Cao, Y., Wang, Y., and Flores-Cerrillo, J. (2023a). Dynamic optimization of multiproduct cryogenic air separation unit startup. *Industrial & Engineering Chemistry Research*, 62(27), 10542–10558.

Quarshie, A.W.K., Swartz, C.L.E., Madabhushi, P.B., Cao, Y., Wang, Y., and Flores-Cerrillo, J. (2023b). Modeling, simulation, and optimization of multiproduct cryogenic air separation unit startup. *AIChE Journal*, 69(2), e17953.

Schäfer, P., Caspari, A., Kleinhans, K., Mhamdi, A., and Mitsos, A. (2019). Reduced dynamic modeling approach for rectification columns based on compartmentalization and artificial neural networks. *AIChE Journal*, 65(5), e16568.

Tasseff, B., Coffrin, C., Wächter, A., and Laird, C. (2019). Exploring benefits of linear solver parallelism on modern nonlinear optimization applications. *arXiv preprint arXiv:1909.08104*.

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, 25–57.