MULTISCALE PRODUCTION ROUTING IN INDUSTRIAL GAS SUPPLY CHAINS

Qi Zhang[1], Arul Sundaramoorthy[2], Ignacio E. Grossmann*[1] and Jose M. Pinto[3]

[1]Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA
[2]Business and Supply Chain Optimization R&D, Praxair, Inc., Tonawanda, NY 14150, USA
[3]Business and Supply Chain Optimization R&D, Praxair, Inc., Danbury, CT 06810, USA

Abstract

The production routing problem (PRP) considers the coordination of production, inventory, distribution, and routing decisions in a supply chain. In this work, we focus on production routing in the industrial gas business, where the challenge lies in the simultaneous optimization of complex and power-intensive production plants and highly integrated supply chains with vendor-managed inventory (VMI). We present a multiscale PRP (MPRP) model with two time grids, to which an MILP-based heuristic solution approach is applied to solve large-scale instances. The proposed MPRP framework is applied to a real-world industrial gas supply chain with 2 plants, approximately 240 customers, 20 vehicles, and a planning horizon of 4 weeks. The results show that the proposed solution method clearly outperforms available alternative approaches in terms of solution quality.

Keywords

Production routing, supply chain planning, industrial gas, MILP-based heuristic.

Introduction

Industrial gas supply chains are among the most complex in the process industry. In the so-called merchant liquid business, industrial gas companies distribute liquid products (liquid oxygen, nitrogen, argon, hydrogen, etc.) in bulk to the customers using tractor-trailers. These supply chains are typically very large, with multiple plants and hundreds of customers. The industrial gas industry is one of the first to adopt the concept of vendor-managed inventory (VMI), which allows the direct control of customers' inventories. VMI significantly increases the flexibility of the supply chain, but also increases the complexity in decision-making. Moreover, the production process (cryogenic air separation) is highly power-intensive; hence, rapid operational changes in response to time-sensitive electricity prices have to be considered.

Because of the strong interdenpendencies in the supply chain, production and distribution operations have to be coordinated, which has become a major goal in integrated supply chain management in recent years (Laínez and Puigjaner, 2012). Glankwamdee et al. (2008) formulate a simplified production-distribution linear programming (LP) model in which the distribution part is approximated by resource constraints on truck and driver hours required for the planned deliveries. Marchetti et al. (2014) propose a production routing framework in which a heuristic is applied to generate a number of routes a priori, where a route is defined as a set of customers that can be visited in one trip. These routes are then included in the integrated model such that the assignment of routes to available vehicles can be optimized. In their proposed frameworks, Glankwamdee et al. (2008) and Marchetti et al. (2014) apply rather simplistic models of the production processes, which can be a serious drawback as process dynamics are not accurately represented, and hence, solutions may be suboptimal or even infeasible when implemented in practice. Zamarripa et al. (2016) apply a rolling horizon heuristic to large-scale instances of the model proposed by Marchetti et al. (2014), obtaining

*To whom all correspondence should be addressed *grossmann@cmu.edu*

near-optimal solutions in shorter computation times.

The PRP in its classical form (Adulyasak et al., 2015) is formulated as a mixed-integer linear program (MILP) and integrates the lot-sizing problem (LSP) and the inventory routing problem (IRP). Zhang et al. (2016b) extend the classical PRP by replacing the LSP with a more detailed scheduling formulation designed for complex continuous production processes. In addition, a second, finer time grid is created to accommodate the production part while the IRP is formulated using the coarse time grid. Furthermore, an MILP-based heuristic approach has been developed in order to solve instances of industrially relevant sizes. This paper complements the work by Zhang et al. (2016b) with a thorough analysis of a real-world industrial case study, where we focus on the comparison between the proposed approach and alternative solution strategies that are more commonly used in practice.

Problem Statement

We consider an industrial gas business that produces and sells a set of products $i \in I$, namely liquid oxygen (LO2), liquid nitrogen (LN2), gaseous oxygen (GO2), and gaseous nitrogen (GN2). While LO2 and LN2 (product subset $\bar{I}$) can be stored and transported to customer sites using tractor-trailers, GO2 and GN2 (product subset $\hat{I}$) are nonstorable and have to be distributed via pipelines immediately after their production; hence, routing decisions only involve liquid product customers. The supply chain consists of a set of continuous air separation plants $p \in P$ and a number of product-specific customers, of which each customer $c \in C_i$ has a given demand and storage capacity for product i.

We assume that each production plant can operate in a set of discrete operating modes $m \in M_p$, where each mode is defined by its production capacity and cost function. The complexity in the production process arises from the fact that generally, the products cannot be produced independently from each other; hence, correlations in production rates have to be considered. Furthermore, the dynamic behavior of the plant is constrained by restrictions on the rate of change and transitions between operating modes. The plants have inventory capacities for the liquid products.

Product-specific vehicles are used to transport products from the plants to the customers. Each vehicle is assigned to one particular plant and is defined by its capacity, speed, and cost, which may include fuel and labor costs. For every trip, a vehicle leaves the plant, visits one or multiple customers, and returns to the plant at the end of the trip. The length of a trip is limited.

The goal of the MPRP is to optimize the production schedule and routing decisions for a given scheduling horizon. For each time period, the production schedule should provide the following information: the operating mode, the production rate for each product, and the amounts of products stored. Routing decisions are made on the assignment of vehicles to trips and the allocation of customers to each trip. We assume that products can be purchased externally at given costs if customers do not have sufficient inventory to satisfy demand.

Multiscale Model

We apply a discrete-time formulation with two time grids, one with a fine and the other with a coarse time discretization, where the scheduling horizon is divided into time periods of the lengths Δt^{f} and Δt^{c}, respectively. For the sake of clarity, we refer to a time period in the fine time grid as a level-1 time period and to a time period in the coarse time grid as a level-2 time period whenever this distinction is necessary. The sets of time periods in the fine and the coarse time grids are denoted by $\overline{T}^{\mathrm{f}}$ and $\overline{T}^{\mathrm{c}}$, respectively.

The concept of the multiscale time discretization is shown in Figure 1, which also illustrates the two main routing assumptions: (1) every trip is completed within a level-2 time period; (2) shipments are loaded into the vehicles in the first level-1 time period of the corresponding level-2 time period.

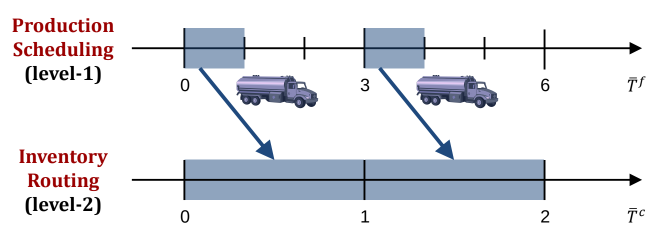

Figure 1. Multiscale time discretization.

The production scheduling part of the MPRP model is based on a mode-based formulation developed in previous works; for more details, see Zhang et al. (2016a,b). The IRP part is modeled with a set-partitioning formulation, which selects the optimal routes among a set of prespecified feasible routes. The resulting integrated model is an MILP. In the following, the MPRP model is presented in a compact form, which highlights the route assignment part of the formulation:

$$
\begin{aligned}
\min \quad & \sum_{i\in\bar{I}}\sum_{p}\sum_{t\in\overline{T}^{c}}\sum_{s\in\overline{S}_{ipt}} \beta_{ips}\, x_{ipts} + C(y) \\
\text{s.t.} \quad & \overline{DL}_{ipts} \le V_i\, x_{ipts} \quad \forall\, i \in \bar{I},\, p,\, t\in\overline{T}^{c},\, s\in\overline{S}_{ipt} \\
& \widehat{DL}_{iptsc} \le \widehat{DL}_{ict}^{\max}\, x_{ipts} \quad \forall\, i \in \bar{I},\, p,\, t\in\overline{T}^{c}, \\
& \qquad\qquad s\in\overline{S}_{ipt},\, c\in\overline{C}_{ips} \\
& \sum_{s\in\overline{S}_{ipt}} x_{ipts} \le L_{ipt} \quad \forall\, i\in\bar{I},\, p,\, t\in\overline{T}^{c} \\
& y \in Y
\end{aligned}
\qquad \text{(MPRP)}
$$

where $\overline{S}_{ipt}$ is the set of routes that can be used by vehicles distributing product i and assigned to plant p in level-2 time period t, $\overline{C}_{ips}$ is the set of customers that can be visited on route $s \in S_{ip}$, and x_{ipts} is a binary variable that equals 1 if route $s \in \overline{S}_{ipts}$ is selected. Associated with x_{ipts} is the fixed transportation cost β_{ips}. The first constraint limits the delivery quantity, $\overline{DL}_{ipts}$, in a single trip with the vehicle capacity, V_i. The delivery to customer $c \in \overline{C}_{ips}$ is denoted by $\widehat{DL}_{iptsc}$, which is limited by $\widehat{DL}_{ict}^{\max}$. The number of available vehicles, L_{ipt}, bounds the number of selected routes. The remaining variables are aggregated in the vector y, $C(y)$ denotes the corresponding linear cost function, and $y \in Y$ represents all remaining constraints.

Solution Method

The set-partitioning formulation is known to exhibit a relatively tight LP relaxation, but it can require an exponential number of routes to fully describe the problem. However, at a feasible solution, only a very small fraction of all possible routes are selected. Hence, instead of working with the full route set, we propose to only consider a small subset of routes when solving (MPRP) and dynamically update the route set such that only candidate routes that can potentially lead to reduced costs are included.

The flowchart in Figure 2 shows the main steps in the proposed algorithm. We start with an initial set of candidate routes, which could be all single-stop routes or a subset of them obtained through a customer inventory analysis. The MPRP is then solved with the current set of candidate routes. Based on the solution of (MPRP), the route set is updated. To solve large-scale instances, it is crucial to keep this route set small. The algorithm stops if no new routes are added, the solution has not improved for a number of consecutive iterations, or the time limit is reached.

Algorithm 1 shows the general scheme for generating routes based on the current solution of (MPRP), which provides an estimate of the amount of product that needs to be delivered to each customer in each time period. Using this information, the algorithm identifies inefficiencies in the current selection of routes and proposes new candidate routes that may improve the solution.

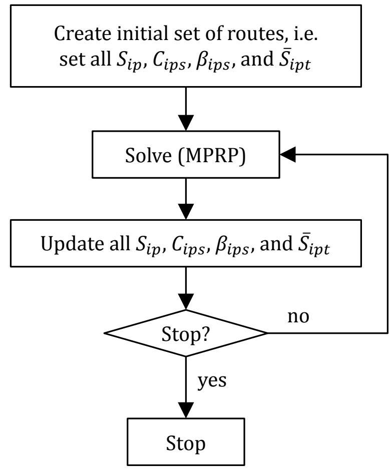

Figure 2. Flowchart for the proposed solution algorithm.

At each iteration, the algorithm is applied to every product $i \in \bar{I}$, plant p, and time period $t \in \overline{T}^{c}$. First, the procedure RemoveRoutes(i,p,t,Ω) removes routes that have not been selected for Ω consecutive iterations from the set S_{ipt}. Next, we examine every selected route s for which the delivery quantity is less than the vehicle capacity, i.e. $\overline{DL}_{ipts} < V_i$. The procedure CreateRoutesA(i,p,t,s) generates new routes, if possible, by inserting additional customers into the current route s. A selection of these new routes are added to the route set S_{ipt} based on a ranking of the potential savings. Besides underutilized vehicles, another indicator for distribution inefficiency is the purchase of products at high costs, which usually occurs due to the lack of efficient multistop routes. Hence, in the next step, we consider customers whose demands are met by purchasing additional products, i.e. all $c \in \widehat{C}_{ip}$ for which $PC_{ict} > 0$, where $\widehat{C}_{ip}$ is a subset of C_i and denotes the the set of customers that can be reached from plant p. Similar to CreateRoutesA(i,p,t,s), the procedure CreateRoutesB(i,p,t,c) generates multistop routes involving customer c and adds them to S_{ipt} based on a ranking of the potential savings.

Algorithm 1 General scheme for route generation based on current solution.

```
for all i ∈ Ī, p, t ∈ T̄^c do
    RemoveRoutes(i, p, t, Ω)
    for all s for which x_ipts = 1 and DL̄_ipts < V_i do
        CreateRoutesA(i, p, t, s)
    end for
    for all c ∈ Ĉ_ip for which PC_ict > 0 do
        CreateRoutesB(i, p, t, c)
    end for
end for
```

The proposed solution algorithm is inspired by the concept of column generation, with the main difference being that here, new columns are generated by using a heuristic rather than by solving a rigorous pricing problem. As a result, convergence to the optimal solution cannot be guaranteed, which is the main limitation of the proposed algorithm. For more details on the algorithm, we refer to Zhang et al. (2016b).

Industrial Case Study

The proposed MPRP framework is now applied to a real-world industrial test case provided by Praxair. We consider a supply chain consisting of 2 plants, P1 and P2, and approximately 240 customers. The two plants have a combined fleet of 10 LO2 and 10 LN2 tractor-trailers. While Plant P1 has to satisfy demand for both liquid and gaseous products, Plant P2 only serves liquid product customers.

Cryogenic air separation is highly power-intensive such that the vast majority of the variable production cost is the cost of electricity. Electricity prices can vary significantly across different locations. In this case, Plant P1 participates in the day-ahead market in which the price varies over time, whereas Plant P2 purchases power at a constant unit price. A forecast of the day-ahead prices is available for the given planning horizon.

The MPRP is solved for a planning horizon of 4 weeks, where we choose Δt^{f} and Δt^{c} to be 4 h and 12 h, respectively, resulting in 168 level-1 and 56 level-2 time periods. We apply the proposed algorithm to this large-scale MPRP and present the solution obtained after one hour runtime. Note that due to confidentiality reasons, we cannot disclose detailed information about the supply chain network, plant specifications, and actual product demands. Therefore, all results are given as dimensionless quantities, and numerical values are normalized if necessary.

Figure 3 shows the electricity consumption and price profiles for both plants over the entire planning horizon. One can see that the electricity price at Plant P2 is significantly higher than the average electricity price at Plant P1. As a result, in order to reduce energy cost, Plant P2 is shut down three times for extensive periods of time and also at the end of the planning horizon. One can further see that the solution suggests load shifting at Plant P1 in order to take advantage of low-price hours.

There is a trade-off between production and distribution costs that is not apparent from Figure 3. Although the electricity price is almost always lower at Plant P1, it does not utilize its full production capacity, i.e. more production could be shifted from Plant P2 to Plant P1. However, the higher production cost is offset by the reduction in distribution cost because more customers are located closer to Plant P2 than to Plant P1.

Figures 4 and 5 show the product flows and inventory

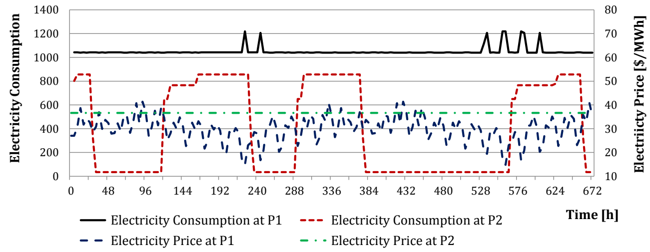

Figure 3. Electricity consumption and electricity price profiles for each plant.

profiles for the liquid products at Plants P1 and P2, respectively. In Figure 4, one can clearly see the effect of load shifting at Plant P1. At Plant P2, inventory is accumulated during hours of production such that products can be drawn from the inventory and distributed to the customers when the plant is shut down, as depicted in Figure 5.

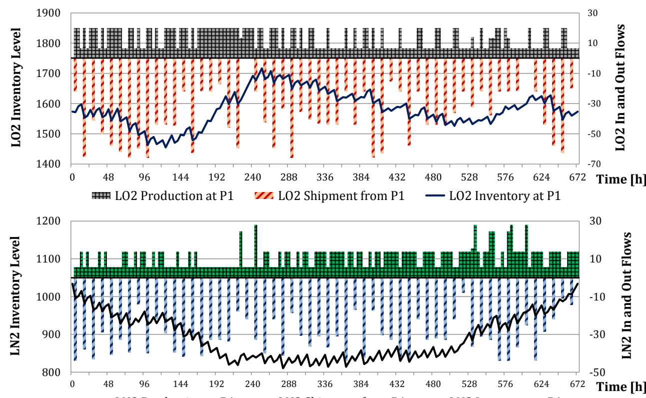

Figure 4. Production quantities, shipments, and inventory levels of LO2 and LN2 at Plant P1.

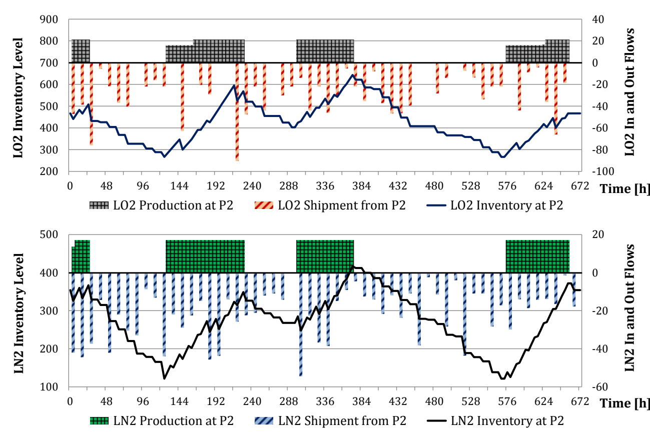

Figure 5. Production quantities, shipments, and inventory levels of LO2 and LN2 at Plant P2.

We compare our solution with the ones obtained from two alternative solution methods. The first method is a typical two-phase heuristic. In Phase 1, the MPRP is solved with a simplified distribution model only considering direct shipments. The delivery quantities obtained from Phase 1 are then used as fixed orders in Phase 2, where routing decisions are made. Since the orders are fixed, the routing problem decomposes into independent subproblems, one for each product i, plant p, and time period $t \in \overline{T}^c$. In the following, we refer to this approach as Heuristic PH1. The second approach is an extension of Heuristic PH1, referred to as Heuristic PH2, which further incorporates fixed costs for customer visits. The fixed distribution costs in Heuristic PH2 prevent the model from suggesting a large number of deliveries with small quantities; however, they also introduce additional binary variables that considerably increase the computational complexity.

Heuristics PH1, PH2, and H3, with the latter being our proposed algorithm with dynamic route generation, apply equivalent representations of the production side; however, the distribution side is modeled with different levels of accuracy. For this comparative study, we first apply Heuristics PH1, PH2, and H3 to obtain the production plan and the plant-to-customer allocation decisions for each of the three solution approaches. Then, the same routing tool is applied to the three sets of plant-to-customer allocation decisions to determine optimal (or near-optimal) routes and accurate routing costs.

Table 1 compares the solutions obtained from Heuristics PH1, PH2, and H3. The table shows the breakdown of the total costs (TC) into the production costs (CPD) and distribution costs (CDI) for each plant. In this test case, no additional product purchase is required, and inventory costs are negligible; hence, these costs are omitted. Furthermore, the table shows the computation time for each solution method. In terms of total cost, Heuristic H3 outperforms both Heuristics PH1 and PH2, with relative cost savings of 8.7 and 2.4 %, respectively, which can be attributed to the more rigorous modeling of routing decisions. One can see that compared to Heuristics PH1 and PH2, Heuristic H3 suggests producing less at Plant P1 and more at Plant P2. This production plan results in higher total production cost, but in overall proves to be the better choice since the routing cost can be significantly reduced by distributing more from Plant P2.

Table 1. Comparison of costs and solution times for the industrial test case.

	Heuristic PH1	Heuristic PH2	Heuristic H3
TC	100.00	93.46	91.26
CPD_{P1}	32.67	32.66	31.88
CPD_{P2}	13.05	13.12	15.01
CDI_{P1}	42.53	36.61	32.40
CDI_{P2}	11.75	11.07	11.97
ST [s]	218	900	3600

Figure 6 shows for each day of the planning horizon the number of customers to visit as suggested by each of the three solutions. While Heuristic PH1 proposes to visit on average 66 customers per day, the average numbers of visited customers per day are 30 and 25 for Heuristics PH2 and H3, respectively. Heuristic PH1 creates many deliveries with small quantities, which leads to inefficient routes. This effect is mitigated in Heuristic PH2 by introducing fixed distribution costs, ultimately resulting in lower routing costs. However, the improved solution quality comes at the cost of higher computational expense. While Heuristic PH1 solves in 218 s, the solution from Heuristic PH2 is obtained after 900 s. Among the three solution approaches, Heuristic H3 obtains the best solution, but only after 3600 s.

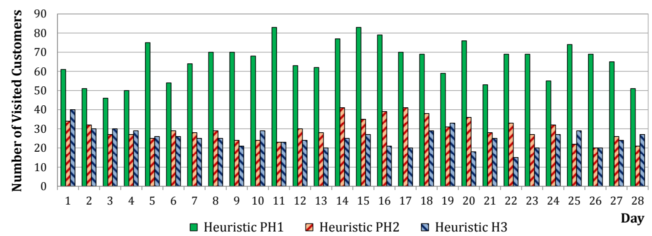

Figure 6. Comparison of the numbers of customers to be visited on each day of the planning horizon as suggested by Heuristics PH1, PH2, and H3.

In practice, under normal circumstances, the plant-to-customer allocation is fixed, i.e. each customer is assigned to a particular plant and only receives delivery from this plant, which may limit the flexibility in the supply chain operations. To compare the differences between the proposed solutions and the current practice, we show in Figure 7 for each of the three solutions the changes in plant-to-customer allocation compared to the current plant-to-customer allocation. Here, an allocation change is defined as one customer that is to be visited in the corresponding solution from a plant different from the one to which it is currently assigned. The number of allocation changes can be interpreted as a measure for the amount of disruption in the default assignment required to obtain the suggested solution. In practice, small changes are desired; a large number of allocation changes may suggest that the current plant-to-customer allocation or the current assignment of vehicles to plants is inadequate. In this case, significantly fewer allocation changes, on average 7 per day, are required for Heuristic H3 than for Heuristics PH1 and PH2, which require on average 24 and 11 allocation changes per day, respectively.

Another advantage of Heuristic H3 is that it only considers feasible routes; hence, the proposed deliveries are guaranteed to be feasible. In contrast, Heuristics PH1 and PH2 may make plant-to-customer allocation decisions that are infeasible in the subsequent routing step, in the sense that not all proposed deliveries can actually be made. In this particu-

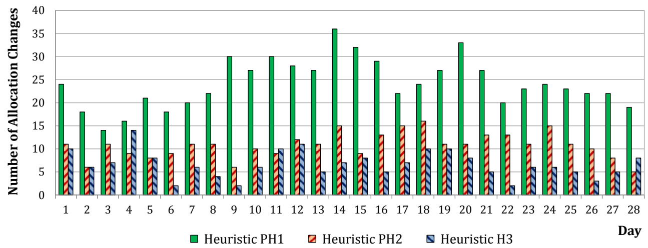

Figure 7. Comparison of the numbers of plant-to-customer allocation changes from the current assignment required for Heuristics PH1, PH2, and H3.

lar test case, routing infeasibility does not occur because the customers are located relatively close to each other such that the limit on the travel distance is not an issue. However, in other supply chain networks with longer inter-customer distances, the situation of routing infeasibility may very well arise when Heuristics PH1 and PH2 are applied.

Conclusions

In this work, we have applied a previously developed MPRP framework to a large-scale real-world industrial gas supply chain. The MILP MPRP model involves two different time grids. While a detailed mode-based production scheduling model captures all critical operational constraints on the fine time grid, vehicle routing is considered in each time period of the coarse time grid. In order to solve the industrial-scale MPRP, an iterative MILP-based heuristic approach has been applied, which solves the MILP model with a restricted set of candidate routes at each iteration and dynamically updates the set of candidate routes for the next iteration.

The main features of the model have been demonstrated in the industrial case study. In particular, the results show the level of detail at which production scheduling is considered, which is necessary because of the time-sensitive electricity prices. Also, the trade-off between production and distribution costs is captured well in the integrated model. Moreover, the computational results show that the proposed algorithm outperforms available alternative solutions in terms of solution quality, although longer computation times are required. With a runtime of one hour, a solution was achieved that improves the solution obtained from a standard heuristic approach by approximately 9 %.

Acknowledgments

The financial support from the National Science Foundation under Grant No. 1159443 and from Praxair is gratefully acknowledged.

References

Adulyasak, Y., Cordeau, J.-F., and Jans, R. (2015). The production routing problem: A review of formulations and solution algorithms. *Computers & Operations Research*, 55:141–152.

Glankwamdee, W., Linderoth, J., Shen, J., Connard, P., and Hutton, J. (2008). Combining optimization and simulation for strategic and operational industrial gas production and distribution. *Computers & Chemical Engineering*, 32(11):2536–2546.

Laínez, J. M. and Puigjaner, L. (2012). Prospective and perspective review in integrated supply chain modelling for the chemical process industry. *Current Opinion in Chemical Engineering*, 1(4):430–445.

Marchetti, P. A., Gupta, V., Grossmann, I. E., Cook, L., Valton, P.-M., Singh, T., Li, T., and André, J. (2014). Simultaneous production and distribution of industrial gas supply-chains. *Computers & Chemical Engineering*, 69:39–58.

Zamarripa, M., Marchetti, P. A., Grossmann, I. E., Singh, T., Lotero, I., Gopalakrishnan, A., Besancon, B., and André, J. (2016). Rolling Horizon Approach for Production-Distribution Coordination of Industrial Gases Supply Chains. *Industrial & Engineering Chemistry Research*, 55:2646–2660.

Zhang, Q., Sundaramoorthy, A., Grossmann, I. E., and Pinto, J. M. (2016a). A discrete-time scheduling model for continuous power-intensive process networks with various power contracts. *Computers & Chemical Engineering*, 84:382–393.

Zhang, Q., Sundaramoorthy, A., Grossmann, I. E., and Pinto, J. M. (2016b). Multiscale production routing in multicommodity supply chains with complex production facilities. *Submitted for publication.*