Simultaneous Production and Distribution of Industrial Gas Supply-Chains

Pablo A. Marchetti[1], Vijay Gupta[1], Ignacio E. Grossmann[1†],
Lauren Cook[2], Pierre-Marie Valton[3], Tejinder Singh[2], Tong Li[2], and Jean André[3]

[1]Dept. of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213
[2]American Air Liquide Inc., Delaware Research and Technology Center, Newark, DE 19702
[3]Air Liquide, Paris Saclay R&D Center, 78350 Jouy-en-Josas, France

Abstract

In this paper, we propose a multi-period mixed-integer linear programming model for optimal enterprise-level planning of industrial gas operations. The objective is to minimize the total cost of production and distribution of liquid products by coordinating production decisions at multiple plants and distribution decisions at multiple depots. Production decisions include production modes and rates that determine power consumption. Distribution decisions involve source, destination, quantity, route, and time of each truck delivery. The selection of routes is a critical factor of the distribution cost. The main goal of this contribution is to assess the benefits of optimal coordination of production and distribution. The proposed methodology has been tested on small, medium, and large size examples. The results show that significant benefits can be obtained with higher coordination among plants/depots in order to fulfill a common set of shared customer demands. The application to real industrial size test cases is also discussed.

Keywords

Supply-chain optimization, Industrial gases, Production planning, Inventory routing problem, Multi-period model, Mixed-integer linear programming

[†] Corresponding author. *E-mail address:* grossmann@cmu.edu

1. Introduction

This paper addresses the problem of determining optimal operational level decisions for the coordinated production and distribution of industrial gas supply-chains. In this industry, cryogenic air separation processes are used to produce oxygen, nitrogen, and argon both as gaseous and liquid products. Air separation units consume large amounts of electricity, mainly due to the operation of the compressors used at different stages of the process. Depending on the equipment configuration selected, alternative operation modes with different production capacities and energy efficiencies are available at each plant. The electricity market has greatly evolved over the last decades and electricity prices fluctuate during the day depending on market conditions. Moreover, power providers offer different pricing schemes, where electricity cost variations can occur every hour, every minute, or on a peak/off-peak basis depending on the scheme adopted. Because the cost of electricity is the main component of the production cost, production level decisions can be optimized by following the electricity market conditions.

On the distribution side, gaseous and liquid customers of industrial gases are usually served by pipeline and bulk truck delivery, respectively. Gaseous products are supplied into the pipeline directly from the air separation unit. Customers of gaseous products are usually located near the plants and referred as “on-site” or “over-the-fence” customers. Their demands are tied by strict contractual obligations and must always be met. Therefore, when an event can impact the gaseous production, inventory of liquid product may be gasified and sent to the pipeline to ensure that over-the-fence customer demands are satisfied. Moreover, product must be imported from other sources if the available inventory is not enough to meet the gaseous demand.

Liquid products are stored on-site in cryogenic storage tanks. From there they are loaded into trailers and carried to customer sites by truck. The transportation cost for bulk truck delivery is the main component of the distribution cost. Both the frequency of deliveries to a given customer and the selection of routes supplying product to multiple customers are critical in order to reduce the transportation cost. Cryogenic storage tanks are available at customer sites and gauge readings received from remote telemetry units are used to keep track of the inventory levels. The vendor is responsible by contract to ensure that customers do not run out of product. Instead of receiving “call in” orders for replenishment, in this industry, vendor managed inventory (VMI) systems are usually used. Based on customer consumption profiles and market conditions, the vendor decides not only how much product to deliver but also when the delivery will take place. The logistics problem that simultaneously considers vehicle routing and inventory management at customer sites is called inventory routing problem (IRP). The distribution schedule depends not only on the availability of trucks, trailers, and drivers, but also on the inventory levels of liquid products at plants and customers. The replenishment of storage tanks at customer locations must be secured by an appropriate distribution schedule, which ideally should feature a minimum distribution cost.

The main goal of this contribution is to assess the benefits of the optimal coordination of production and distribution decisions in an industrial gases supply-chain. A mixed-integer linear programming (MILP) formulation minimizing the overall cost of production and distribution over a limited time horizon (7 to 14 days) is presented. Figure 1 depicts the main processes and decision problems involved in the supply-chain under consideration. On the production side, multiple plants and products are considered, and the optimal operation modes and production rates for every plant taking into account fluctuating electricity prices are sought. On the distribution side, multiple depots are included, and trucks at a given depot can deliver product from multiple plants. Furthermore, in order to ensure customer storage replenishments, products can be purchased from alternative sources. As the number of sources, depots, and customers increase, the selection of the alternative routes becomes a critical issue. The connection between production and truck-distribution is given by the amount of liquid product stored at the plants at any given time. The main focus of this paper is the production and distribution of liquefied product. However, the demand for gaseous product is considered if a plant is forced to decrease its gaseous production (e.g., during a plant shutdown). In this case, as a back-up solution, the liquid product must be vaporized to meet the gaseous customer pipeline demand.

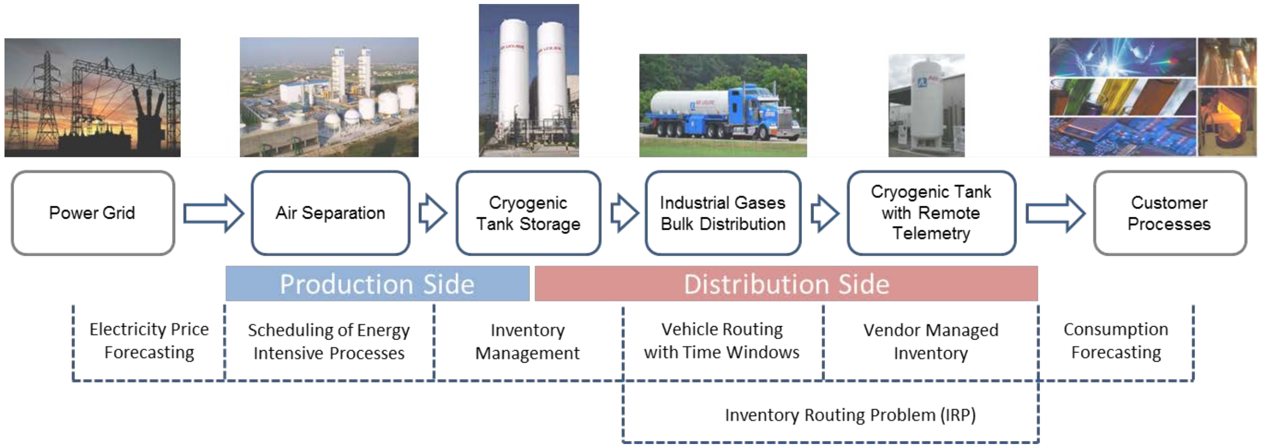

Figure 1. Processes and decision problems involved in the supply-chain of liquefied industrial gases.

The rest of this paper is organized as follows: Section 2 includes a review of previous works on production and distribution of industrial gases and the related energy intensive scheduling and vehicle logistics problems. A formal description of the problem statement and main assumptions are presented in Section 3. Section 4 introduces the mathematical formulation for coordinated multi-plant production and distribution. Different levels of coordination for the supply-chain decisions are also described to compare the proposed simultaneous method with sequential and single-plant alternatives. Two illustrative but realistic examples are presented in Section 5, and the application of the proposed methodology to real industrial-size test cases is also discussed. Finally, conclusions are presented in Section 6.

2. Literature Review

Smith and Klosek (2001) provide an overview of air separation technologies used to obtain nitrogen, oxygen, argon, and other atmospheric or specialty gases. Also, a review of relatively recent developments in cryogenic air separation processes and prospective analysis of future technologies can be found in Castle (2002). An analysis of potential savings on electricity cost under a real time pricing (RTP) scheme for industrial end users through improved demand management is presented by Ross et al. (1998). Several contributions tackle the problem of deciding optimal operational level decisions for energy intensive processes such as air separation, where the cost of electricity is sought to be minimized (Daryanian et al., 1989; Ierapetritou et al., 2002; Karwan and Keblis, 2007; Mitra et al., 2012). By using formulations with multiple time periods and assuming steady state operation at each time, they seek to reduce the overall cost of production over a given planning horizon. Daryanian et al. (1989) studied the application of an optimal algorithm for single storage electricity consuming processes with electricity spot prices. They present a case study of an air separation facility and analyze potential savings comparing flat rate electricity costs with spot priced electricity. The results show that rescheduling electricity consumption provides opportunities for substantial savings in electricity costs. Recent works incorporate uncertainty in the electricity prices. Ierapetritou et al. (2002) developed a two-stage stochastic programming formulation where uncertainty in the power prices is considered within a given portion of the optimization horizon. In turn, Karwan and Keblis (2007) developed a mixed-integer programming formulation embedded in a rolling horizon procedure to minimize the cost of running an air separation unit under real time pricing (RTP). They also conducted simulation studies to assess the robustness of the production plans obtained and investigated the conditions under which a RTP scheme is more attractive than time of use (TOU), which refers to fixed electricity prices for daily, weekly, or seasonal blocks of electricity. They found out that RTP is preferred over TOU when there is more production flexibility, i.e. conditions such as lightly loaded plants or short ramp-up times. In turn, Mitra et al. (2012a) developed a mixed-integer programming model for optimal production planning of processes such as air separation. While considering known electricity prices for a time horizon of one week, they include the modeling of transition times and costs between production modes and improve the tightness of the MILP model. As an additional contribution, long-term strategic investment decisions were considered using cyclic short-term production schedules that take into account seasonal electricity cost fluctuations (Mitra et al., 2012b).

On the distribution side, a description of the inventory routing problem, its main characteristics, and a survey of relevant literature can be found at Campbell et al. (1998) and Bertazzi et al. (2008). Kleywegt et al. (2002) provide a categorization of the variants of the inventory routing problem (IRP) that have been studied by different researchers. Also, refer to the recent paper of Coelho et al. (2013) for the history of IRP and a review of different exact and heuristic approaches considered to solve wide variety of IRP problems. Taking into account customer demands, relevant IRP formulations include either deterministic (Dror et al., 1985; Chien et al., 1989; Jaillet et al. 2002; Campbell et al. 2004; Benoist et al., 2011) or stochastic (Dror and Ball, 1987; Çetinkaya and Lee, 2000; Kleywegt et al., 2002) approaches. The most simplified IRP is NP-hard as it contains the classical vehicle routing problem (VRP). Therefore, to solve the IRP, most research works have focused on heuristic solution approaches given its complexity. In many contributions the IRP problem is decomposed into sub-problems, e.g. Campbell et al. (2004), which are solved by approximate or exact methods (i.e. Branch and Cut or Column Generation). In some cases, heuristic methods are applied to the sub-problems in order to identify upper and lower bounds. Some of the studies provided integrated and iterative approaches and evaluated the effectiveness of integrating routing and inventory decisions in their models. Others have proposed heuristic methods to be compared with approaches used in industrial-gas industry (Dror and Ball, 1987; Campbell et al., 2002). We should also note that several papers dealing with the infinite horizon problem use a distribution policy that is similar to the fixed partition policy (first introduced by Anily and Federgruen, 1993), direct deliveries, order-up-to level policy and zero-inventory ordering (Bertazzi et al., 2002; Chan et al., 1998). Fixed partition policy specifies regions (subset of customers) covering all customers and always replenishes the customers in the same region together. Distribution policy of order-up-to always fills a customer up-to its inventory capacity, whereas in zero-inventory ordering an order is placed only when its inventory drops to zero. Examples of applications combining vehicle routing with inventory management at customer sites for industrial gas distribution are presented in the seminal work of Bell et al. (1983), and more recently by Campbell et al. (2002).

A comprehensive review of the literature addressing supply-chain coordination either at the operational or the strategic planning level is presented by Thomas and Griffin (1996). The potential of a better coordination between production, inventory, and distribution activities has been initially explored by Chandra and Fisher (1994). They presented a computational study to examine the value of a better coordination between production and distribution. The study considered a single plant and multiple products, which are delivered to multiple retail outlets by a fleet of vehicles. They developed both an integrated formulation and a decoupled production and distribution model, and analyzed multiple test cases with alternative values for the main problem parameters (length of horizon, number of products and retail stores, setup, inventory holding, and vehicle travel costs). Reductions ranging from 3 to 20% of the total cost were reported by comparing the solutions obtained. This work has been followed up by several authors including Fumero and Vercellis (1999) and Park (2005). Focusing on industrial gas supply-chains, Glankwamdee et al. (2008) developed a simplified production and distribution planning linear model. In order to account for uncertainty, they extended this formulation both via a minmax model and a two-stage stochastic program, and tested the effectiveness of the proposed methods using simulation. However, only time-aggregated planning decisions were considered and neither plant mode selection nor vehicle routing details are included in the model. Also, You et al. (2011) developed a mixed-integer linear programming model to integrate long term planning decisions of sizing storage tanks at customer locations with truck routing decisions at the operational level. They also propose two efficient computational methods in order to solve large-scale instances, one based on a two-level decomposition strategy and the other on a continuous approximation approach for the routing decisions.

3. Problem Statement and Main Assumptions

3.1 Overall Supply-Chain Problem

The problem of production-distribution coordination of an industrial gases supply-chain can be stated as follows. Given are the following items:

(i) a set of industrial gases production plants $p \in P$,

(ii) a set of production modes or unit configurations $m \in M_p$ in which plant p can operate at any given time,

(iii) a set of liquid products $i \in I$ to be considered in the supply-chain, and the specific products $i \in I_{pm}$ that are produced while plant p operates in mode m,

(iv) the production rate limits ($w_{pmi}^{\min}$, $w_{pmi}^{\max}$) and the energy consumption per unit of product (usp_{pmi}) for each product i, plant p, and mode m,

(v) the maximum storage capacity $Q_{pi}^{\max}$ for product i at plant p, and the initial inventory L_{pi}^{ini} of each product,

(vi) a set of customers $c \in C_i$ for each product type i, the maximum capacity $Q_c^{\max}$ of the storage tank at customer c, and the initial inventory level L_c^{ini},

(vii) a set of depots $d \in D$, where trucks $k \in K_{di}$ are available for the delivery of each product i,

(viii) the maximum capacity U_k^{truck} and the travel cost per unit distance c_k of each truck k,

(ix) the locations of plants, depots, and customers, allowing to calculate route distances,

(x) the time horizon H divided in consecutive time periods $t \in T$, each one with duration Δ_t, where at least two time periods per day are considered (i.e., half day peak and off-peak time periods),

(xi) the electricity price forecast u_{pt} for each plant at each time $t \in T$,

(xii) the forecast of the product consumed, $R_{c,t}$, and the estimated required safety stock, $Q_{ct}^{\min}$, at time t for each customer $c \in C_i$.

The goal is to determine operational level decisions for each time period t including the following: the mode of operation and production rate of the final products at each plant, the amount of inventory maintained for each product at each source and customer location, and the amount of each product to be delivered to customers through the routes to be selected. The objective function is to minimize the total cost of production and distribution for the entire supply-chain.

Figure 2 shows an example of an industrial gases supply-chain as addressed in this paper, which consists of a set of plants (P), depots (D), and customers (C). As mentioned before, while we do not focus on the production of gaseous products, we do take into account the situation when a given plant has a limited production capability and liquid inventory must be gasified to fulfill the gaseous customer demand. We assume here that at each time t the forecast of the liquid volume to be gasified and sent by pipeline, $R_{pi,t}^{site}$, is known. It is also assumed that the following additional information is given: (a) the initial operational state of each plant p (i.e., b_p^{ini} = 1 when plant p is running at the beginning of the time horizon), (b) the fixed start-up cost F_{pt}^{start} if plant p needs to be powered up at time t, and (c) the minimum inventory level $Q_{pit}^{\min}$ (redline) allowed for product i at plant p at any given time.

If multiple product grades $j \in J_i$ can be manufactured for a given product i, the product grade $j = grade(p, i)$ associated with product i at plant p is also available. The set of plants $p \in P_{ci}$ from which product i can be delivered to customer c, or the set of product grades $j \in J_c$ that can be delivered to customer c must be specified. Furthermore, while deliveries are primarily made by sourcing the product from plants of the company, in situations in which there are shortages the product must be purchased from an alternative source $p \in P^{\text{alt}}$. In this case the price per unit volume ($C_{pi,t}^{purchase}$) and the maximum amount of product that can be purchased ($Q_{pi,t}^{purchase}$) for each product *i* and time *t* are assumed to be known.

The savings through full production-distribution coordination are quantified using a model that, while being approximate, has a sufficient level of details to be realistic. On the production side, the changeover times and costs required to switch between production modes are assumed to be negligible. On the distribution side, the detailed hourly scheduling of the drivers and the assignment of the trailers attached to each truck are not considered. Instead, a unique combination of truck/trailer/driver called "truck" is used, disregarding the potential unavailability of drivers or trailers. The distribution costs are exclusively based on distances, not on the time spent to deliver to the customers, which means not considering the exact calculation of loading/unloading, traveling, and waiting times. Besides, complex schedules allowing multiple trips per shift or layovers are not possible. Some of these features can be handled by adding average transition costs between production modes, limiting the number of vehicles for a given time period, or forbidding the selection of routes that do not satisfy specific timing or distance constraints.

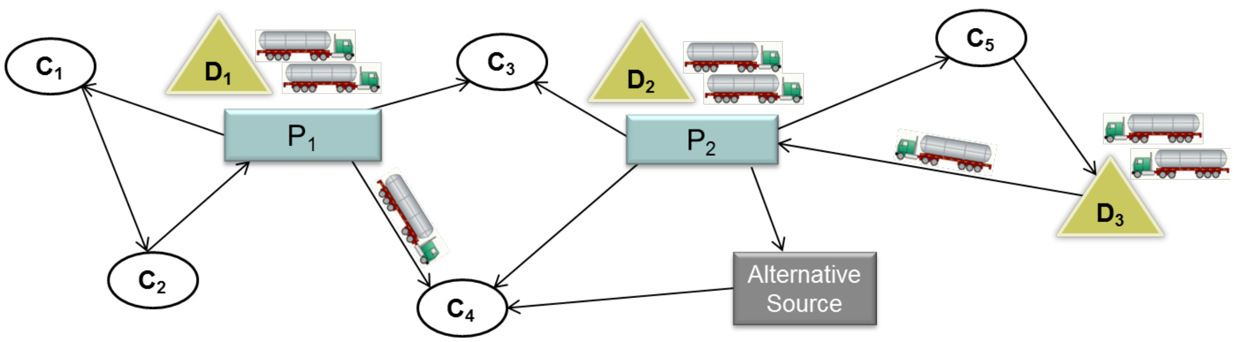

Figure 2: Industrial gas supply-chain illustrative example.

3.2 Coordination Levels for Production and Distribution

The production-distribution coordination problem can be studied at various levels of coordination. In this paper, we introduce definitions for the various levels of coordination of an industrial gas supply chain (see Table 1).

The sequential coordination strategy refers to the approach of first generating a production schedule based on the historical data on the behavior of distribution (e.g. statistics on truck withdrawals) or other sources of information, and then generating an optimal distribution schedule based on the inventory levels (as a result of the production schedule calculated before) and associated customer demands. In contrast, the simultaneous coordination strategy determines production and distribution schedules through a simultaneous optimization approach assuming knowledge of the information input on the production (electricity prices) and distribution sides (customers consumption/demand). Both strategies may be applied across multiple sources (dynamic sourcing) or limited to a single source (fixed sourcing).

	Sequential (Production then Distribution)	**Simultaneous** (Production and Distribution)
Single plant/depot (Fixed Sourcing)	No Coordination b/w plants and production-distribution	Coordination b/w production-distribution but No coordination b/w plants
Multi-plant/depot (Dynamic Sourcing)	Coordination b/w plants but No Coordination b/w production-distribution	Coordination b/w production-distribution as well as plants (fully coordinated)

Table 1. Production-Distribution Coordination Levels

Only operational decisions concerning the existing supply-chain are considered. Design decisions concerning investments in new installations or expansions are not included. Furthermore, while in a real scenario electricity costs and customer demands are subject to uncertainty, in this contribution we do not take into account any uncertainty in the forecasted data. Besides, the production and distribution decisions are limited to the time horizon given. Consequently, the tradeoff between short term savings and reducing the long term overall cost is not explored.

The proposed "fully coordinated" MILP model for the multi-source simultaneous case is described in the next section. The models corresponding to the remaining levels of coordination described in Table 1 are obtained from the fully coordinated formulation as a special case.

4. Simultaneous Production-Distribution Model

In order to develop a model that simultaneously optimizes production and distribution decisions over a finite time horizon, the first step is to define an adequate time representation. In this contribution a uniform discretization of the time horizon H is used. Thus, a finite number of time periods $t \in T$ is given during which both production and distribution events take place. The usual scheme is, for instance, to consider two time periods per day following peak and off-peak electricity price intervals. The model constraints for production and distribution decisions are described next, followed by the objective function to be used.

4.1 Production Side

Constraints (1)-(9) model the production side of the supply-chain. For a given time period *t*, the main model decisions are the operating modes and production rates at each plant, from which the power consumptions and product inventory levels are derived.

4.1.1 Selection of production modes at each plant

The binary variable B_{pmt} is introduced to represent that plant *p* operates in mode *m* at time period *t*. Each plant can operate at most in a single mode during time *t*, this condition being enforced by constraint (1). When $B_{pmt} = 0$ $\forall m \in M_{pt}$ then plant *p* is not in operation (shut-down mode) during time period *t*.

$$\sum_{m \in M_{pt}} B_{pmt} \leq 1 \qquad \forall p \in P, t \in T \tag{1}$$

When a plant starts operating, there is a cost of transitioning from shut-down to any valid mode $m \in M_{pt}$. This start-up cost usually corresponds to the cost incurred while running the plant until the required operating conditions are reached. For example, during this start-up phase, electricity may be consumed while the output of the air separation units does not meet product grade specifications.

Constraints (2) and (3) are included to detect the transition from shut-down mode to any other mode *m*. The binary variable b_{pt}^{start} is 1 if plant *p* is shut-down in the previous time period $t - 1$ and is turned on when time period *t* begins. In particular, constraint (2) represents this condition for the first time period t_0, while constraints (3) correspond to the rest of the time periods. Eqn (2) is only needed if the plant is initially in shut down mode (i.e., $b_p^{init} = 0$). Notice that with the constraints (2) and (3) it is possible to define b_{pt}^{start} as a continuous variable in the interval [0, 1].

$$\sum_{m \in M_{p,t_0}} B_{p,m,t_0} \leq b_{p,t_0}^{start} \qquad \forall p \in P: \ (b_p^{init} = 0) \tag{2}$$

$$\sum_{m \in M_{pt}} B_{pmt} \leq \sum_{m \in M_{p(t-1)}} B_{pm(t-1)} + b_{pt}^{start} \qquad \forall p \in P, t \in T \ : \ t > t_0 \tag{3}$$

Different costs may be considered for each possible transition from one production mode to another; however, in the current model the only cost considered is the transition cost to start-up the plant.

4.1.2 Production rate limits and power consumption

On each operating mode, production capacity constraints that limit the rate of production of each product i must be considered. Let the continuous variable W_{pmit} represent the production rate of product i at plant p while running mode m in time period t. Constraint (4) establishes both the lower ($w_{pmi}^{\min}$) and upper ($w_{pmi}^{\max}$) bounds for the production rate of every product i that can be produced in mode m, given that mode m is on at time period t (B_{pmt} = 1). If a given mode m is not selected (B_{pmt} = 0), then all production rates for that mode are driven to zero. For some configurations the minimum production rates can be defined by the relation $w_{pmi}^{\min} = \eta_p w_{pmi}^{\max}$, where η_p (e.g. 70%) is the turn-down ratio defined for plant p.

$$B_{pmt}\, w_{pmi}^{\min} \le W_{pmi,t} \le B_{pmt}\, w_{pmi}^{\max} \qquad \forall\, i \in I_{pm}, m \in M_{pt}, p \in P, t \in T \tag{4}$$

Moreover, additional constraints limiting the total liquid production for a given production mode m at plant p can be specified by Eqn (5), where the parameters $\alpha_{pm,i\lambda}$ and $\pi_{pm,\lambda}$ are the coefficients and upper bound, respectively, for a linear combination of the production rates of every product i. The set LIM_m stands for the limits of the feasible region of production mode m, where each λ is associated to a limiting hyperplane. Notice that in each mode m we assume that the plants are flexible enough to operate anywhere within the limits given by Eqns (4) and (5).

$$\sum_{i \in I_{pm}} \alpha_{pm,i\lambda} W_{pmi,t} \le B_{pmt}\, \pi_{pm,\lambda} \qquad \forall \lambda \in LIM_m, m \in M_{pt}, p \in P, t \in T \tag{5}$$

The power consumption of plant p in time period t is given by Eqn (6), where the parameter usp_{pmi} is the energy requirement per unit of product i (unit specific power) when plant p operates in production mode m.

$$PW_{p,t} = \sum_{m \in M_{pt}} \sum_{i \in I_{pm}} \left(usp_{pmi} \cdot W_{pmi,t} \right) \qquad \forall\, p \in P, t \in T \tag{6}$$

4.1.3 Inventory constraints at plants

Storage is assumed to be available at the plants to keep the inventory of every product $i \in I_p$. The continuous variable L_{pit} stands for the inventory level of liquid product i at plant p at the end of time period t. Equation (7) establishes the lower and upper bounds for the level of product i, which must lie between the minimum level (redline) and the maximum storage capacity of the facility for that product. The minimum inventory level ensures that excess demand of over-the-fence/on-site gaseous customers can be met using this inventory as a back-up source. Moreover, this redline ($Q_{pit}^{\min}$) is a given parameter that may vary over the planning horizon (not constant) based on the gaseous customer demand profile, while the maximum limit $Q_{pi}^{\max}$ is related to the physical capacity of the storage facility (a constant value) for product i.

$$Q_{pit}^{\min} \le L_{pit} \le Q_{pi}^{\max} \qquad \forall\, i \in I_p, p \in P, t \in T \qquad (7)$$

The material balance constraints (8) are required to keep track of the inventory level of product i at each time period t. In particular, the amount of product in storage at plant p is equal to the inventory of the product at the previous time period, plus the production over time period t, minus both the total amount of product supplied on-site ($D_{pi,t}^{site}$) and the total product distributed by trucks ($D_{pi,t}^{truck}$) at time t. The variables $D_{pi,t}^{site}$ and $D_{pi,t}^{truck}$ are introduced in the next sections. Also, for the first time period the value of $L_{pi,t-1}$ is given by the inventory level of the plant when the time horizon begins (L_{pi}^{ini}).

$$L_{pi,t} = L_{pi,t-1} + \Delta_t \sum_{m \in M_{pt,i}} W_{pmi,t} - D_{pi,t}^{site} - D_{pi,t}^{truck} \qquad \forall\, i \in I_p, p \in P, t \in T \qquad (8)$$

Material balance constraints (8) are the main constraints that connect the production and distribution sides of the supply chain.

4.1.4 Gaseous customer supply

As indicated by Equation (9), the amount of product distributed on-site for each time period t is defined as the gaseous volume supplied by vaporization through the pipeline to an over-the-fence customer sitting near the plant. This vaporization of liquid product is needed only when the gas can not be supplied from the separation column because the plant is shut down. The parameter $R_{pi,t}^{site}$ is the demand forecast of the over-the-fence customer for product *i*.

$$D_{pi,t}^{site} = R_{pi,t}^{site}\left(1 - \sum_{m \in M_{pt}} B_{pmt}\right) \qquad \forall\, i \in I_p,\, p \in P, t \in T \tag{9}$$

4.2 Distribution Side

The main distribution decisions include the amount of product being delivered from a given source, the truck being used, and the set of customers being visited within a given time period. We assume here that a truck performs a round-trip on each time period *t*. While this assumption is valid most of the time, in the general case a driver can eventually complete two or three trips in a single shift before finishing his working hours. To allow several short trips in a single shift, the duration of each time period may be reduced to obtain a more accurate discretization of the time horizon. However, this increases the model size and the computational effort required to find solutions that are accurate enough (just dividing each time period by two duplicates the number of binary variables and constraints of the model). Besides, it is also assumed that a single trip starts and ends at the same depot from which the truck departs. While the inventory capacity constraints are verified only at the time interval limits, it is assumed that truck loading and unloading tasks may occur anytime within these limits and that there is enough capacity available to accommodate the production-distribution schedule if needed.

If multiple combined trips are needed, the problem becomes a multiple source pick-up and delivery problem, a level of detail that is not tackled in this contribution. Multiple trucks, multiple sources and multiple alternative routes generate a combinatorial explosion of the number of alternatives to be explored on the distribution side of the supply-chain.

4.2.1 Selection of routes

Each truck is assigned to a fixed depot and dedicated to transport a unique kind of product. Thus, the set of trucks $k \in K_{di}$ is defined for every depot *d* and product *i*. As shown in Figure 3, distribution by trucks is accomplished by the following steps: (a) a truck *k* travels from its depot *d* to a valid source location (i.e., a related plant *p*), (b) product *i* is loaded at plant *p* such that the truck capacity U_k^{truck} is not exceeded, (c) truck *k* visits a set of customers *s* (one or more) and delivers the product, which is distributed in any required proportion among them, and (d) truck *k* travels back to its depot. Therefore, given the depot *d*, the plant p, and a set of customers s, the route with the shortest distance (dis_{dps}) to complete the delivery can be calculated a-priori (pre-processed).

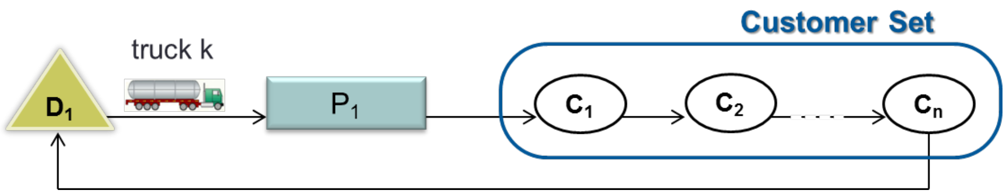

Figure 3. Routes are determined by combining a depot, a plant, and a customer set.

The binary variables Y_{kpt} and y_{kst} are introduced to indicate whether or not truck k is associated to plant p and customer set s, respectively, at time period t. Since each truck is associated to a known depot, there is no need to decide on the depot to be used on a given trip. Constraints (10) and (11) together indicate that each truck k can only be assigned to a single route on each time interval t. Equation (10) represents the fact that a truck k can be assigned to a single set of customers per time period. If the LHS is one, then the truck k is delivering product at time t.

$$\sum_{s \in S_{di}} y_{kst} \leq 1 \qquad \forall k \in K_{di}, i \in I, d \in D, t \in T \tag{10}$$

The set S_{di} stands for all the customer sets associated with routes for product i that start at depot d. Because of Equation (10), at most one customer set s is selected for truck k to visit at each time period t. In turn, each customer set s can include one or more customers. Since the number of possible sets s grows very fast with the number of customers of product i (i.e., $|\ C_i\ |$), an effective route selection method is required to keep the model size reasonable. Appendix A describes the route selection method used herein, which is based on the idea of enumerating all feasible routes, sorting them using an economic criterion, and selecting the most appropriate ones while guaranteeing a minimum number of routes for each customer. Practical sorting criteria are either the route distance or an estimation of the cost per volume sourced for the route.

Given Equation (10), Equation (11) establishes that a sourcing plant is required if and only if truck k is delivering product at time t.

$$\sum_{p \in P_{di}} Y_{kpt} = \sum_{s \in S_{di}} y_{kst} \qquad \forall k \in K_{di}, i \in I, d \in D, t \in T \tag{11}$$

The set P_{di} includes all the plants that are authorized to source product *i* by loading a truck from depot *d*. Section 4.2.7 will further discuss possible delivery restrictions that apply when taking into account different product grades.

4.2.2 Truck load constraints

Continuous non-negative variables E_{kpt} and e_{kst} are introduced to handle the quantity of product delivered by truck *k*. The variable E_{kpt} represents the amount of product loaded by truck *k* at plant *p* in time period *t*, while the variable e_{kst} is the amount delivered by truck *k* to customer set *s* in the same time period. Since only one source is allowed for a given truck, constraint (12) guarantees that only the appropriate variable E_{kpt} is nonzero for some $p \in P_{di}$.

$$E_{kpt} \leq Y_{kpt} U_k^{truck} \qquad \forall k \in K_{di}, p \in P_{di}, i \in I, d \in D, t \in T \tag{12}$$

Also, constraint (13) states that the variable e_{kst} can be nonzero only if truck *k* delivers to the customer set *s* ($y_{kst} = 1$).

$$e_{kst} \leq y_{kst} U_k^{truck} \qquad \forall k \in K_{di}, s \in S_{di}, i \in I, d \in D, t \in T \tag{13}$$

Finally, given the aforementioned bounds for variables E_{kpt} and e_{kst}, Equation (14) is needed to ensure that the amount of product picked up at a given plant is the same one being delivered to the selected customers, for each truck *k* and time period *t*.

$$\sum_{s \in S_{di}} e_{kst} = \sum_{p \in P_{di}} E_{kpt} \qquad \forall k \in K_{di}, i \in I, d \in D, t \in T \tag{14}$$

4.2.3 Plant pick-up and customer delivery amounts

Given equations (12)-(14), three additional constraints are needed to connect both sides of the supply-chain.

On one hand, Equation (15) defines the amount of product *i* delivered by truck from plant *p* at each time *t* (i.e., $D_{pi,t}^{truck}$) as the summation of the product loaded by every truck that stops at *p* at that time period. Delivery limitations established for the depots are taken into account by including the condition $p \in P_{di}$.

$$D_{pi,t}^{truck} = \sum_{d \in D:(p \in P_{di})} \sum_{k \in K_{di}} E_{kpt} \qquad \forall i \in I_p, p \in P, t \in T \tag{15}$$

On the other hand, Eqns (16) and (17) are used to determine the total amount of product delivered to a given customer *c* at time *t* ($D_{c,t}$). Constraint (16) ensures that the product being delivered to each customer set *s* is split among the customers $c \in s$. To this end, the continuous variable d_{sct} is introduced to indicate the amount of product that customer *c* receives at time period *t* from all trucks that deliver to customer set *s* at that time. Notice that the LHS of Eqn (16) is the amount of product carried by all trucks that visit customer set *s*, and the RHS is the amount delivered to the customers in *s*. Moreover, the set S_i includes all customer sets for a given product *i*.

$$\sum_{d \in D:(s \in S_{di})} \sum_{k \in K_{di}} e_{kst} = \sum_{c \in s} d_{sct} \qquad \forall s \in S_i, i \in I, t \in T \tag{16}$$

Finally, constraint (17) calculates $D_{c,t}$ as the summation of all the deliveries being made to *c* through all relevant sets *s*.

$$D_{c,t} = \sum_{s \in S_i:(c \in s)} d_{sct} \qquad \forall c \in C_i, i \in I, t \in T \tag{17}$$

Figure 4 depicts the material flow represented by the material balance constraints (15), (14), (16) and (17) and defined with the continuous variables $D_{pi,t}^{truck}$, E_{kpt}, e_{kst}, d_{sct}, and $D_{c,t}$. Also, Figure 5 shows in more detail the interpretation of the material balance constraint (15) when multiple trucks from different depots load product at plant P_1.

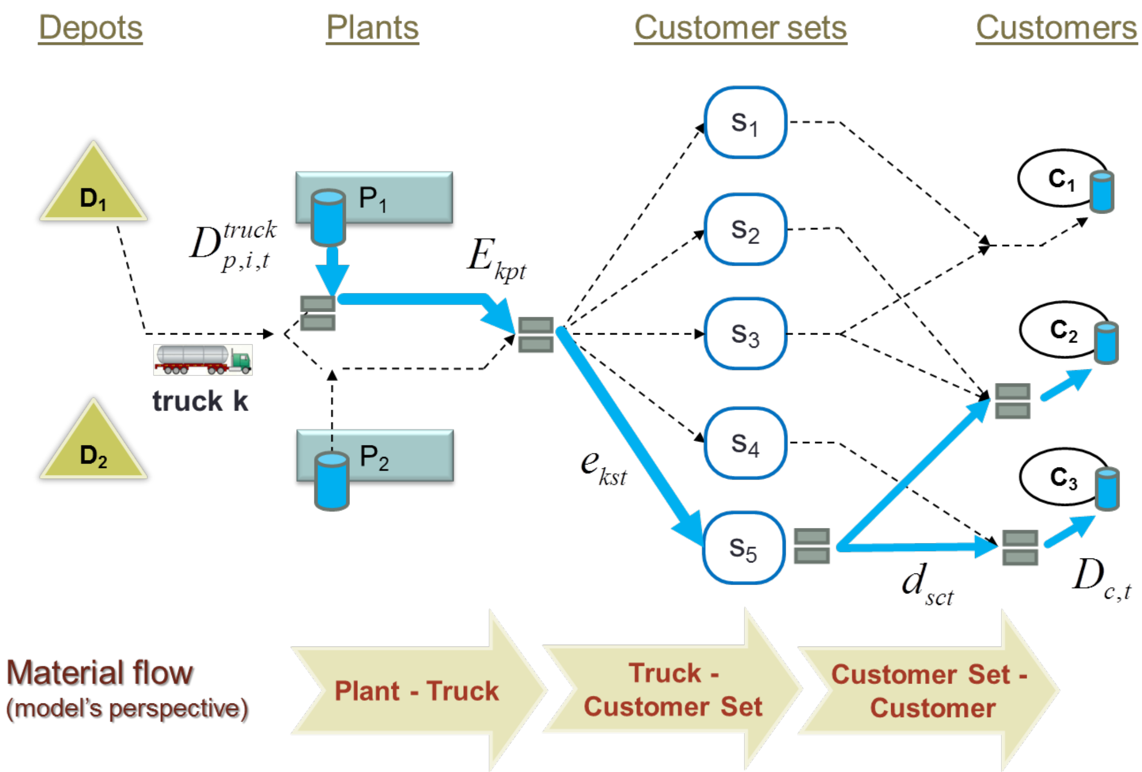

Figure 4. Distribution side continuous variables used to represent the delivery of liquid products from plants to customers.

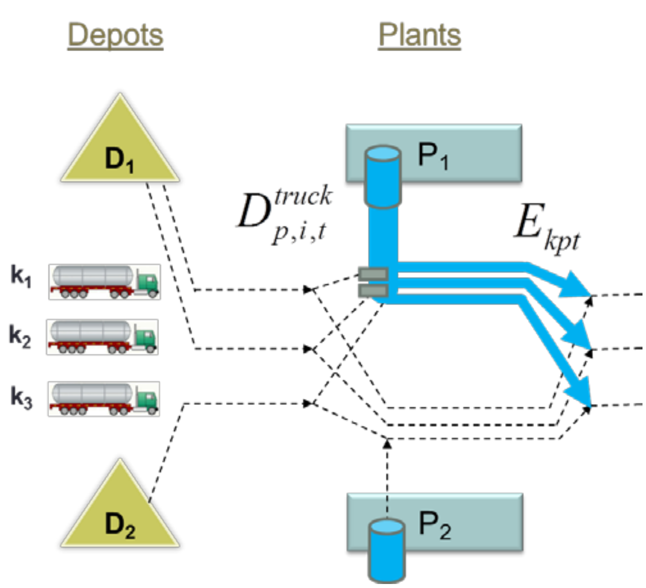

Figure 5. Loading of multiple trucks at a given plant.

4.2.4 Route distances

The next set of constraints is needed to determine the distance traveled by a truck k when a delivery is made at time period t. As mentioned before, given the depot d, source p, and destination s (i.e. a set of customers) associated with each possible trip, the shortest traveling distance (dis_{dps}) can be calculated *a priori*. This can be done either by enumerating all possible alternatives or using a specific TSP algorithm, mainly because the number of customers in every customer set s is relatively small.

For each truck $k \in K_{di}$ departing from depot d at time period t, its selected route will be determined by the specific binary variables Y_{kpt} ($p \in P_{di}$) and y_{kst} ($s \in S_{di}$) that are equal to one. However, since the information on the route is disaggregated on these binary variables, it is not straightforward to calculate the distance traveled by truck k.

Using the parameter dis_{dps}, Eqn (18) defines the minimum distance ($dis_{ds}^{\min}$) required to deliver product i to the set of customers $s \subset C_i$ using any truck from depot d. In other words, the parameter $dis_{ds}^{\min}$ is the traveling distance for the closest plant, taking into account a route with a fixed depot d and customer set s.

$$dis_{ds}^{\min} = \min_{p \in P_{di,s}} [dis_{dps}] \qquad \forall s \in S_{di}, i \in I, d \in D \tag{18}$$

Given the parameter $dis_{ds}^{\min}$, if a source different than the closest one is used, then an additional distance must be added in order to account for the correct delivery cost.

To this end, a non-negative continuous variable β_{kt} is introduced, representing the distance added to $dis_{ds}^{\min}$ to account for a source different than the closest one (usually the default source). Constraint (19) sets the lower bound for variable β_{kt} based on the source and customer-set decision variables, where the parameters δ_{dps} and $\delta_{dp,i}^{\max}$ are defined in equations (20) and (21). The parameter δ_{dps} represents the additional distance needed between the minimum ($dis_{ds}^{\min}$) and the complete distance (dis_{dps}), when plant p is selected.

Furthermore, $\delta_{dp,i}^{\max}$ is the maximum distance δ_{dps} taking into account all routes associated to depot d and product i. When $Y_{kpt} = 1$ (i.e. the plant p has been selected for the truck k), the RHS of constraint (19) becomes equivalent to β_{kt}, and the summation on the LHS provides the adequate lower bound for the additional distance to be considered. Otherwise, $Y_{kpt} = 0$ and variable β_{kt} can always be driven to zero while Eqn (19) is still satisfied.

$$\sum_{s \in S_{di}} \delta_{dps} \, y_{kst} \le \delta_{dp,i}^{\max} \left(1 - Y_{kpt}\right) + \beta_{kt} \qquad \forall p \in P_{di}, k \in K_{di}, i \in I, d \in D, t \in T \tag{19}$$

with the definition of the following parameters:

$$\delta_{dps} = dis_{dps} - dis_{ds}^{\min} \qquad \forall s \in S_{di}, \, p \in P_{di}, i \in I, \, d \in D \tag{20}$$

$$\delta_{dp,i}^{\max} = \max_{s \in S_{di}} [\delta_{dps}] \qquad \forall p \in P_{di}, i \in I, \, d \in D \tag{21}$$

Finally, the distance traveled by truck *k* in time period *t* is given by the continuous variable DIS_{kt}, which is defined in Eqn (22). The RHS includes: (a) the minimum distance required to deliver product *i* to customer set *s* from the plant that is more conveniently located and (b) the additional distance β_{kt} that is needed if a different plant is selected.

$$DIS_{kt} = \sum_{s \in S_{di}} dis_{ds}^{\min} \, y_{kst} + \beta_{kt} \qquad \forall \, k \in K_{di}, i \in I, d \in D, t \in T \tag{22}$$

Figure 6 shows an example where a given truck k_1 from depot D_1 is delivering product to customers c_1, c_2, and c_3 (i.e., customer set s_1). Two alternative routes are shown: $r_1 = (D_1,P_1,s_1)$ and $r_2 = (D_1,P_2,s_1)$. Thus, the minimum distance needed to make the delivery is $dis_{D_1,s_1}^{\min} = dis_{D_1,P_1,s_1}$, where P_1 is the closest plant. Moreover, the additional distance if plant P_2 is selected is given by $\delta_{D_1,P_2,s_1} = dis_{D_1,P_2,s_1} - dis_{D_1,s_1}^{\min}$. If $y_{k_1,s_1,t} = 1$ and $Y_{k_1,P_1,t} = 1$ then $DIS_{k_1,t} = dis_{D_1,s_1}^{\min} = dis_{D_1,P_1,s_1}$. Otherwise, if $y_{k_1,s_1,t} = 1$ and $Y_{k_1,P_2,t} = 1$ then $DIS_{k_1,t} = dis_{D_1,s_1}^{\min} + \beta_{k_1,t}$ and because of Equations (18)-(20) we have $\beta_{k_1,t} \ge \delta_{D_1,P_2,s_1}$ from which $DIS_{k_1,t} \ge dis_{D_1,P_2,s_1}$ can be derived.

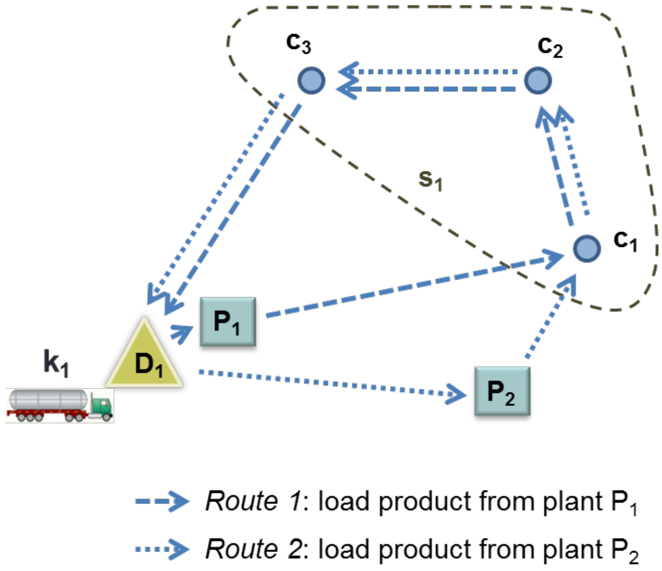

Figure 6. Routes with alternative sources for the same depot and customer set.

4.2.5 Inventory constraints at customer sites

The inventory level at customer locations must also be tracked by the model. For each customer $c \in C_i$, the level of product i inventory at the end of time t (L_{ct}) must lie between the minimum desired level (safety stock) and the maximum storage capacity of the tank as established by Eqn (23). Notice that the safety stock can be given as a parameter with variations over the planning horizon based on the consumption profile of that particular customer.

$$Q_{ct}^{\min} \le L_{ct} \le Q_c^{\max} \qquad \forall c \in C_i, i \in I, t \in T \tag{23}$$

Constraint (24) represents the material balance constraint for the inventory of product i at each customer location. In particular, the amount of product i in the customer storage tank in time period t is equal to the inventory of that product at the previous time period, plus the product delivered to the customer in time period t, D_{ct}, less the amount of product consumed by the customer, R_{ct}, in the same time t. For the first time period t_0, the inventory at the previous time period $t - 1$ will be given by the initial inventory of product i at customer c (L_c^{ini}).

$$L_{ct} = L_{c(t-1)} + D_{ct} - R_{ct} \qquad \forall c \in C_i, i \in I, t \in T \tag{24}$$

As an alternative to Equations (23) and (24), Appendix B presents the constraints required when the volumes and time windows for each delivery are specified beforehand (planned deliveries).

4.2.6 Deliveries from alternative sources

Industrial gas customers usually have strict requirements on product availability bound by specific contractual obligations. In general, the golden rule for any industrial gases provider is that a customer must never run out of product. Thus, if the available inventory at the owned plants is not enough to fulfill some required obligations, then the product must be provided by purchasing it from an alternative source in order to replenish any customer inventory levels that are subject to redline conditions in a timely manner.

In this section we indicate the changes in the mathematical model required to handle the possibility of purchasing product from an alternative source. To this end, the set of plants P is split into two disjoint subsets P^{own} and P^{alt}, standing for the owned plants and the alternative sources (i.e. typically plants owned by other companies). By doing this, the set P must be replaced by P^{own} in the Eqns (1)-(9) that model the production side of the supply chain (see Section §4.1). However, constraints (11), (12), (14), (15), and (18)-(21) defined in Sections §4.2.1 to §4.2.4 remain unchanged, since now the set $P = P^{own} \cup P^{alt}$ also includes the alternative sources. For each additional source $p \in P^{alt}$, variables Y_{kpt} and E_{kpt} are also included. Given these modifications, the total amount of product i purchased at an alternative source $p \in P^{alt}$ at time period t ($D^{truck}_{pi,t}$) is still defined by Eqn (14). Besides, the maximum amount of product i that can be purchased at time t is now given by the parameter $Q^{purchase}_{pi,t}$, as indicated by Eqn (25).

$$D^{truck}_{pi,t} \le Q^{purchase}_{pi,t} \qquad \forall\, i \in I_p, p \in P^{alt}, t \in T \tag{25}$$

4.2.7 Sourcing and product grade constraints

Different grades of industrial gas products can be easily handled by the proposed method. For example, when liquid oxygen (LOX) is considered, a distinction may be made between industrial LOX and medical LOX, since they have different product purities. While it is possible to handle the different product grades as different products, this approach may turn out to be over-restrictive. For example, a customer requesting a lower grade product could also receive a higher grade, as long as the required purity specifications are met. Given that each plant p produces a grade j = grade(p, i) for product i, let us consider the binary relation of product grades R(j, j') such that the demand of a customer requiring j can be fulfilled by delivering j'. Thus, R should be a reflexive and transitive relation. Based on this relation a set J_c including all product grades that can be delivered to customer c can be obtained. Consequently, the set of plants from which product i can be sourced to a given customer set s is:

$$P_s = \bigcap_{c \in s} \{ p \in P : grade(p,i) \in J_c \} \qquad \forall\, s \in S_i, i \in I$$

A customer set s should not be considered in the model if $P_s = \varnothing$. Eqn (26) defines the set S_{pi}, which includes all customer sets where product i of plant p can be delivered. The definition $S_{di} = \bigcup_{p \in P_{di}} S_{pi}$ should be employed when Eqn (26) is used.

$$S_{pi} = \{ s \in S_i : p \in P_s \} \qquad \forall\, p \in P, i \in I \tag{26}$$

With the above definitions, constraint (27) must be added to the mathematical formulation to handle multiple product grades for a given product i, representing the possibility to deliver products of higher purity if available to customers who require a lower grade of the same product.

$$Y_{kpt} \le \sum_{s \in S_{pi}} y_{kst} \qquad \forall p \in P_{di}, k \in K_{di}, i \in I, d \in D, t \in T \tag{27}$$

In general, constraint (27) can be applied to restrict the selection of the customer sets s that can be sourced from plant p, when using a truck from depot d. To this end the set S_{pi} must be replaced by a set $S_{pi,d}$, which also takes into account the depot. For example, this situation appears when a route given by d, p, and s exceeds a given maximum distance.

4.2.8 Tightening constraints

Valid cuts that do not eliminate integer solutions from the feasible space are added to the mathematical model in order to improve its computational performance. The proposed cuts are intended to tighten the LP relaxation by improving the calculation of the distribution cost.

Let $\mu_c(t_1,t_2)$ with $t_1 \le t_2$ be the summation of the product consumed by customer c in the interval from time period t_1 to time period t_2, as stated by Eqn (28).

$$\mu_c(t_1,t_2) = \sum_{t=t_1}^{t_2} R_{ct} \qquad \forall c \in C, \{t_1,t_2\} \subset T : t_1 \le t_2 \tag{28}$$

Tightening constraint (29) imposes that at least one delivery must be made to each customer c within a given interval [t_1, t_2]. The LHS of (29) is the number of trucks visiting all customer sets s that include c within the proposed interval.

$$\sum_{t=t_1}^{t_2}\left(\sum_{d\in D}\sum_{k\in K_{di}}\sum_{s\in S_{di}:(c\in s)} y_{kst}\right)\geq 1 \qquad \forall c\in C_i, i\in I, \{t_1,t_2\}\subset T \\ :(t_1\leq t_2)\wedge \mu_c(t_1,t_2-1)\leq Q_c^{\max}-Q_c^{\min}<\mu_c(t_1,t_2) \qquad (29)$$

The selection of the intervals for which Eqn (29) is defined is explained next. The maximum inventory available at customer c between two consecutive replenishments is given by the expression $Q_c^{\max}-Q_c^{\min}$. If the product consumed between t_1 and t_2, i.e. $\mu_c(t_1,t_2)$, is higher than this difference a delivery must be made to customer c within [t_1, t_2]. This condition is therefore necessary for Eqn (29). Besides, to avoid redundant additional constraints the condition $\mu_c(t_1,t_2-1)\leq Q_c^{\max}-Q_c^{\min}$ is also needed. For instance, let Eqn (29) be defined for a given interval [t_1, t_2]. Then, for any $t_3 > t_2$ the condition $\mu_c(t_1,t_3-1)\leq Q_c^{\max}-Q_c^{\min}$ does not hold because $\mu_c(t_1,t_3-1)\geq \mu_c(t_1,t_2)>Q_c^{\max}-Q_c^{\min}$. In this way the constraint (29) is included only for the shortest time interval starting at each time period t_1. When t_1 is the first period of the time horizon, $Q_c^{\max}$ can be replaced by the initial inventory of customer c (L_c^{ini}) without loss of generality.

4.3 Objective Function

The proposed mathematical model seeks to minimize the overall cost of production and distribution for the entire time horizon. The objective function is given by Equation (30). Equation (31) defines the production cost for each time period t, which is given by the start-up and variable production costs of each plant. Besides, Equation (32) sets the distribution cost at time t as the cost of all deliveries made by every truck plus the cost of the product purchased from the alternative sources at the given time period. We should note that we are not including inventory cost as it is normally a minor cost compared with the production and distribution costs. However, it is clear that inventory costs can be trivially included in (30).

$$Minimize \sum_{t \in T} \left(PCost_t + DCost_t \right) \quad (30)$$

$$PCost_t = \sum_{p \in P^{own}} \left(F_{pt}^{start} \cdot b_{pt}^{start} + PW_{pt} \cdot \Delta_t \cdot u_{pt} \right) \quad (31)$$

$$DCost_t = \sum_{d \in D} \sum_{i \in I} \left(\sum_{k \in K_{di}} c_k \cdot DIS_{kt} \right) + \sum_{p \in P^{alt}} \left(C_{p,i,t}^{purch} \cdot D_{p,i,t}^{truck} \right) \quad (32)$$

4.4 Modeling different levels of production-distribution coordination

The simultaneous production-distribution coordination model is given by Equations (1)-(27) and objective function (30). This fully coordinated model is referred as model (M1). Sequential models are derived from (M1) by decomposing the production and distribution optimization into two separate programs that will be connected through a sequence of decisions involving both. We introduce first the production optimization model (M2) generating the production side schedule that minimizes the total cost of production. This model includes the constraints (1)-(9), with objective function (33).

$$Minimize \sum_{t \in T} \left(PCost_t \right) \quad (33)$$

Two options have been considered to set the production targets: either trucks withdrawals $D_{pi,t}^{truck}$ are forecasted directly based on historical frequencies (M2.a) or planned deliveries are set for each customer based on its consumption forecast, storage capacity, and historical delivery data (M2.b). Equation (34) fixes the variable $D_{pi,t}^{truck}$ for the model (M2.a). In this case, the truck withdrawal volume in each time period t is given by the parameter $U_{pi,t}^{withdraw}$. Alternatively, when the production side model (M2.b) is used, Equations (35) and (36) are employed to determine how much product is delivered to each customer in order to fulfill their forecasted demands. The parameter U_{c,t_1,t_2}^{deliv} indicates the volume of product i to be delivered to customer $c \in C_i$ during the time interval $[t_1, t_2]$, where $t_1 \leq t_2$. Besides, the variable σ_{pct} is introduced representing the product delivered from plant p to customer c at time t.

$$D_{pi,t}^{truck} = U_{pi,t}^{withdraw} \qquad \forall\, i \in I_p, p \in P, t \in T \quad (34)$$

$$D_{pi,t}^{truck} = \sum_{c \in C_{pi,t}} \sigma_{pct} \qquad \forall\, i \in I_p, p \in P, t \in T \tag{35}$$

$$\sum_{t=t_1}^{t_2} \left(\sum_{p \in P_{c,t}} \sigma_{pct} \right) = U_{c,t_1,t_2}^{deliv} \qquad \forall\, c \in C_i, i \in I, \{t_1, t_2\} \subset T : U_{c,t_1,t_2}^{deliv} > 0 \tag{36}$$

Finally, we introduce the distribution side optimization program (M3) generating the distribution schedule that minimizes the total distribution cost. Constraints (7)-(9) and (11)-(27), with objective (37) are used. We assume that the variables that handle production mode selection (B_{pmt}) and production rate ($W_{pmi,t}$) are fixed taking into account a solution of a previously solved model (M2).

$$Minimize \quad \sum_{t \in T} \left(DCost_t \right) \tag{37}$$

To show the potential impact of a better coordination of production and distribution decisions, we compare the simultaneous model (M1) with the sequence (M2) → (M3), the later being to determine the production decisions first and then observing the consequences on the availability of product before solving the distribution model.

5. Results and Discussion

The models (M1), (M2) and (M3) were implemented in GAMS 24.1.3 and solved using the commercial solver CPLEX 12.5.1. Computational results were obtained on an Intel Core i7-960 (3.20 GHz, 4 cores) machine with 16 GB of RAM. All instances were solved using the parallel processing capacities of the machine and a relative gap tolerance of 0.01, otherwise default solver setting were used. Two examples including simultaneous production decisions at multiple plants and distribution decisions at multiple depots are presented. Besides, the application of the proposed model to industrial size problem instances is discussed.

5.1 Example 1

A first small test case is presented featuring two plants and two main products (LIN i.e. liquid nitrogen and LOX i.e. liquid oxygen). A unique grade is considered for each product. The plants can be operated in two production modes (High LIN and High LOX) with specific capacity limits. For each plant and product, Table 2 includes the maximum rate for each production mode together with the inventory levels, maximum storage capacity, and redline (minimum level). The minimum production rates are established by a turndown ratio of 60% for plant P_1 and 70% for plant P_2. All product quantities are given in thousand standard cubic feet (Mcf). Figure 7 shows the feasible production rates for each plant and production mode. The unit specific power is 20 kWh/Mcf for every plant, product, and production mode. Besides, we assume that both plants are initially running, and the associated start-up costs are $7,000 for plant P_1 and $4,000 for plant P_2.

Table 2. Plant production and storage data for Example 1

Plant		P_1		P_2		Unit
Product		LIN	LOX	LIN	LOX	
w^{max}	Mode Hi LIN	108	95	100	105	Mcf/h
	Mode Hi LOX	190	37	185	48	
Inventory	Initial	3,500	4,800	4,700	4,000	Mcf
	Maximum	9,000	6,300	8,100	7,000	
	Redline	3,000	2,100	2,500	1,750	

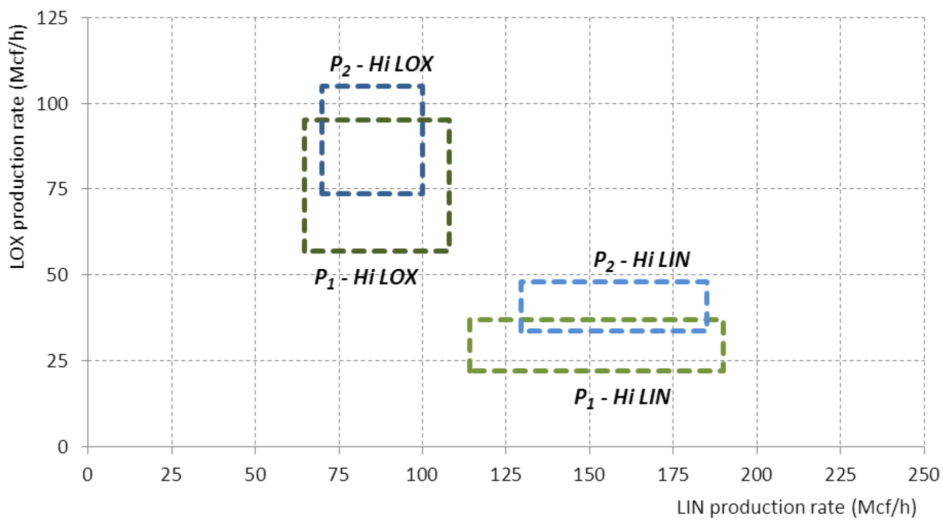

Figure 7. Production rate limits for each operating mode and plant for Example 1.

There is a depot located beside each plant. Depot D_1 is located at plant P_1 and has 5 trucks, 3 with a trailer for LIN and 2 with a trailer for LOX. Also, depot D_2 is located at plant P_2 and has 4 trucks available, 2 for LIN and 2 for LOX. The transportation cost of trucks is 2.85 $/mile, and each trailer has a capacity of 630 Mcf. The supply-chain includes 9 customers (5 LIN customers and 4 LOX customers) to be served by truck delivery. Figure 8 shows a map including all plant/depot and customer locations, which are also indicated in Table 3. Straight line paths are used to calculate route distances. Table 4 includes the liquid product initial inventory level, storage capacity, and redline for each customer, together with the average consumption per day. The default source for LIN customers c_1, c_2, and c_3 and LOX customers c_6 and c_7 is plant P_1. The remaining customers are associated with plant P_2. Thus, as it can be observed in Figure 8, the default source for each customer is the plant in closest proximity.

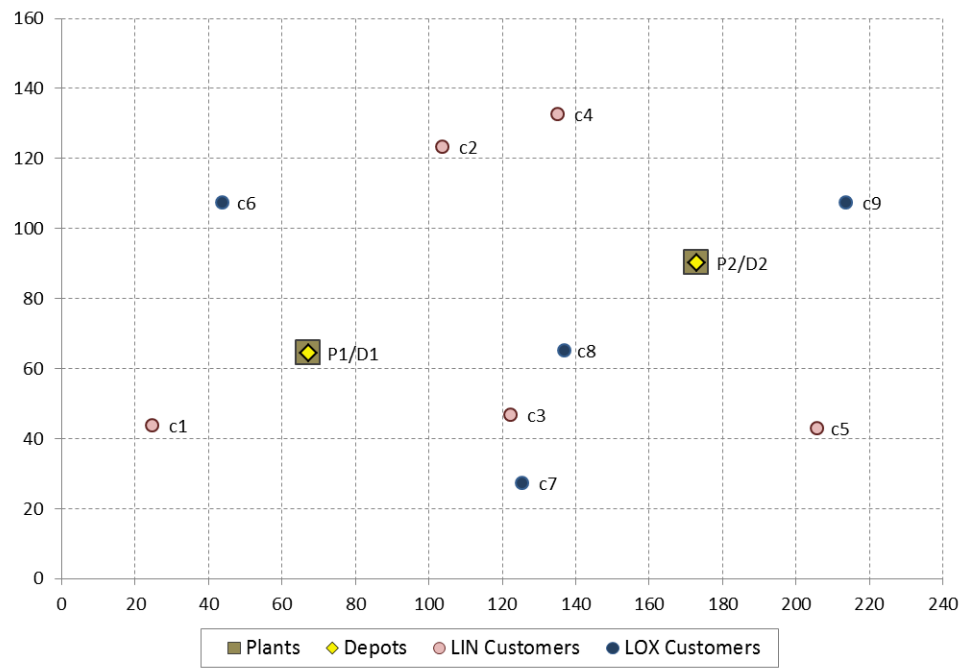

Figure 8. Map for Example 1.

Table 3. Location of plants, depots, and customers for Example 1 (miles).

	Plant / Depots		LIN customers					LOX customers			
	P_1/D_1	P_2/D_2	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9
X coord.	67.2	173.1	24.5	103.6	122.2	135.0	205.7	43.7	125.4	136.8	213.5
Y coord.	64.5	90.2	43.9	123.3	46.7	132.6	43.1	107.5	27.3	65.1	107.4

Table 4. Customer inventory and consumption data for Example 1.

	LIN customers					LOX customers				Unit
	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	
Initial inventory	1750	2620	320	280	2640	920	2380	320	1760	Mcf
Storage capacity	2940	3900	510	350	4380	1560	3800	430	2250	
Redline	940	1750	280	180	1520	590	1670	190	960	
Average consumption	900	1480	280	200	1260	440	1360	180	1000	Mcf/day

The time horizon of one week is discretized into 14 time periods, each one with half day duration and corresponds to peak and off-peak electricity prices on a day. The forecast of the electricity cost at each plant over the whole time horizon is presented in Table 5. For both plants and customers, the lower bounds on inventory levels at the end of the time horizon are set as the available inventories when the time horizon begins.

Table 5. Forecasted electricity prices (cent/KWh) at each plant for Example 1

Time period		t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}	t_{11}	t_{12}	t_{13}	t_{14}
Electricity prices (cent/KWh)	P_1	4.76	4.06	4.37	4.17	4.45	4.06	4.22	3.94	4.25	3.98	4.17	3.94	4.51	4.07
	P_2	3.12	2.98	3.11	2.96	3.13	2.72	2.98	2.67	2.91	2.74	3.04	2.80	3.14	2.84

Given all the problem data presented above, Example 1 has been tested with different levels of production-distribution coordination and alternatives for plant sourcing. We consider the levels of coordination presented in Table 1. In particular, when analyzing the sequential approach of solving first the production and then the distribution model, we test both production targets presented in Section 4.4, either (a) truck withdrawal forecasts or (b) planned delivery forecasts. Table 6 presents the truck withdrawal targets for model (M2.a), and Table 7 presents the delivery targets per customer for model (M2.b). Data in Tables 6 and 7 is used to solve the production models (M2.a) and (M2.b), respectively, under the sequential coordination strategy. However, notice that customer consumptions given in Table 4 are still the targets for the distribution side model (M3). For every level of coordination between production and distribution, both fixed and dynamic sourcing alternatives are also tested. When fixed sourcing is considered each customer is served only by its default source, and no coordination between plants is possible. In this case a mathematical model is solved for each plant. Dynamic sourcing allows some customers to be served by different plants using vehicles from different depots during the time horizon. In this example, the shared customers that can receive product from both P_1/D_1 or P_2/D_2 are LIN customers c_2, c_3, and c_4, and LOX customers c_7 and c_8. All possible routes that visit up to two customers are included in the respective models. Overall, six alternative levels of coordination were tested with Example 1.

Table 6. Forecasted truck withdrawals for model (M2.a) – Example 1.

		t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}	t_{11}	t_{12}	t_{13}	t_{14}
P_1	LIN	3	2	2	2	3	2	2	2	3	2	2	2	3	2
	LOX	2	1	2	1	2	1	2	1	2	1	2	1	2	1
P_2	LIN	2	1	1	1	2	1	1	1	2	1	1	1	2	1
	LOX	1	1	1	1	1	1	1	1	1	1	1	1	1	1

Table 7. Planned deliveries per customer and time period for model (M2.b) – Example 1.

		t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}	t_{11}	t_{12}	t_{13}	t_{14}
LIN	c_1	630	630	-	-	630	630	-	630	630	-	630	630	630	630
	c_2	630	630	1260	630	630	630	630	1260	630	630	630	630	910	630
	c_3	330	-	280	-	280	-	280	-	280	-	280	-	230	-
	c_4	-	270	-	200	-	200	-	200	-	150	-	200	-	180
	c_5	630	630	630	630	630	630	630	630	630	630	630	630	630	630
LOX	c_6	-	630	-	630	-	-	630	-	630	-	-	560	-	-
	c_7	630	630	630	630	1260	630	630	630	630	630	630	700	630	630
	c_8	-	290	-	-	270	-	-	270	-	-	270	-	-	160
	c_9	-	630	630	630	630	-	630	630	630	-	630	630	630	700

Table 8 shows the optimal solution values obtained by applying the proposed models for each alternative level of production-distribution coordination. If no coordination between plants or between production and distribution is considered, a total cost of $70,039 is obtained when the production is based on a forecast of truck withdrawals. However, by improving the coordination the total cost of production and distribution decreases. For instance, when better production targets based on planned deliveries are used and multiple plants/depots are considered, the total cost drops to $65,252, which is almost 7% less than the previous solution. Moreover, the best solution for the fully-coordinated model (M1) with dynamic sourcing has a total cost of $63,089, featuring potential savings of almost 10%. In a similar way, the remaining savings that can be obtained by a better coordination are also shown in Table 8.

Figure 9 compares the production, distribution, and total cost for each level of coordination. In most cases the savings of the fully coordinated model comes from a lower production cost, obtained by redistributing the production load between plants. From Table 5 it can be observed that plant P_2 has lower electricity prices, which makes it convenient to allocate more production there. Thus, shared customers c_2, c_3, and c_7 are candidates to shift sourcing from P_1 to P_2. For each alternative level of coordination with dynamic sourcing, the total volume sourced from each plant and to each customer is presented in Table 9. These amounts can be easily calculated based on the distribution side variables E_{kpt}, e_{kst}, d_{sct}, and D_{ct}. Besides, a comparison of the overall volume sourced per plant and product is shown in Figure 10.

Table 8. Total costs ($) and potential savings (%) due to better coordination for various levels of Production-Distribution Coordination (Example 1).

Coordination Strategy	Sequential		Simultaneous
	Production based on truck withdrawal	Production based on planned deliveries	
Single plant/depot (Fixed Sourcing)	70,039.73 (reference)	67,807.54 (3.19%)	67,145.51 (4.13%)
Multi-plant/depot (Dynamic Sourcing)	69,239.94 (1.14%)	65,252.69 (6.83%)	63,089.46 (9.92%)

Table 9. Total volume sourced (Mcf) from each plant to each customer for the sequential and simultaneous coordination levels with dynamic sourcing – Example 1.

			c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9
Sequential approach	Truck Forecast	P_1	6,300	8,970	1,960	270		3,080	9,180	120	
		P_2		1,390		1,130	8,820		340	1,140	7,000
	Delivery Forecast	P_1	6,300					3,080	4,244		
		P_2		10,360	1,960	1,400	8,820		5,276	1,260	7,000
Simultaneous Model		P_1	6,300	2,940	1,960			3,080	1,000		
		P_2		7,420		1,400	8,820		8,520	1,260	7,000

A brief explanation of the results of Table 9 and Figures 9 and 10 is as follows. The sequential production model based on truck withdrawals (M2.a) uses as its production target an estimation of the number of full-load trucks required at each plant. Because this estimation is higher than the actual demand, some extra production is made in addition to the amount required by the default customers. Thus, the simultaneous distribution model (M3) is able to source some product to c_2 and c_7 using plant P_2. However, the volume sourced from P_2 to the shared customers is still restricted by the production targets. In turn, the best solution of the sequential model (M2.b) with production based on planned deliveries is different. In this case most of the volume required by customers c_2, c_3, and c_7 is sourced from plant P_2. While it reduces the overall production cost, it turns out that trucks from depot D_1 are required to deliver the product from plant P_2, which increases the distribution cost. Finally, the fully coordinated model takes into account both production and distribution resources to find a balanced solution that shifts most of the demand of c_2 and c_7 to plant P_2, without significantly penalizing the distribution cost.

The model size and computational statistics obtained by the application of the simultaneous coordination strategy (M1) with dynamic sourcing is presented in Table 10. The remaining production-distribution models applied to Example 1 have shown similar computational performance, with CPU times varying between 10 and 150 s.

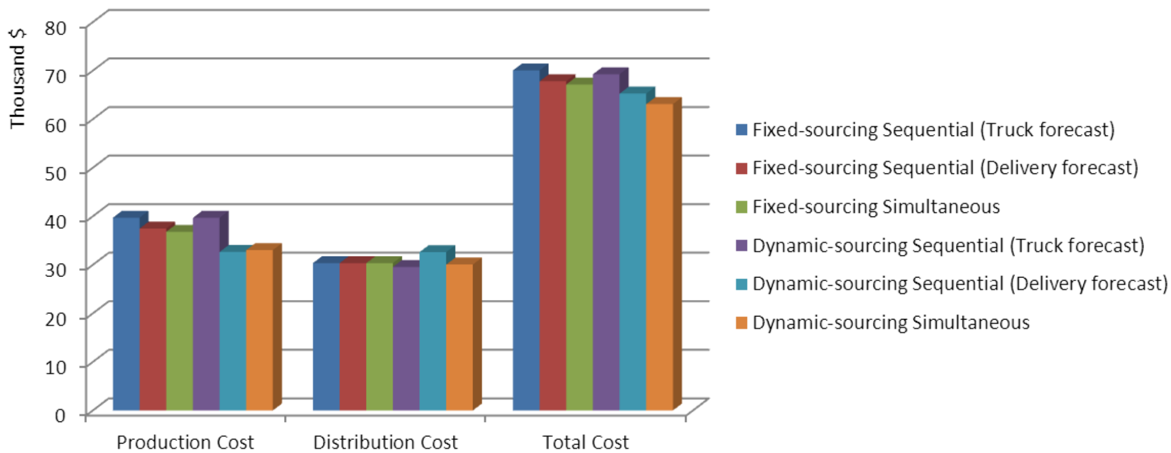

Figure 9. Cost comparison for the alternative levels of Production-Distribution Coordination.

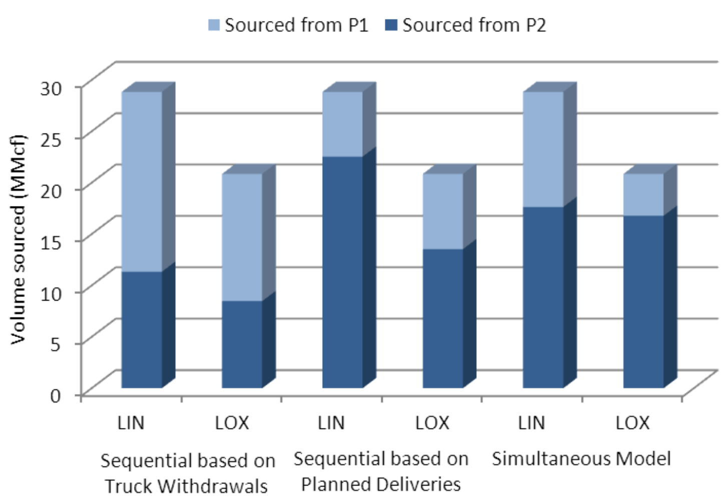

Figure 10. Product sourced per plant for each multi-plant coordination strategy (Example 1).

Table 10. Model size and performance for the multi-plant simultaneous production-distribution model (Example 1).

	Multi-plant Simultaneous Model
Binary variables	1,344
Continuous vars.	2,395
Constraints	2,916
MIP solution	63,089.45
CPU time	11.09 s
Relative gap	1%
Nodes	3,618

5.2 Example 2

A medium size example adapted from a real industrial size test case is presented next. Example 2 includes three plants producing two main products (LIN and LOX). Similar to Example 1, there is a unique grade for each product, and each plant can operate in two different production modes (High LIN and High LOX). Production rate limits for each facility and production mode are shown in Figure 11.

Besides, the supply chain includes 3 depots and one alternative source. Depots D_1 and D_3 are located at plants P_1 and P_3, respectively. Both have 5 trucks, 3 with a trailer for LIN and 2 with a trailer for LOX. Depot D_2 is a standalone depot located nearby plant P_2. It has 4 trucks available, 2 for LIN and 2 for LOX. The alternative source Alt_1, which produces both products, is located at the north-east of depot D_1 and the west of depot D_2, at a similar distance from both. Only trucks from these depots are allowed to load product at plant Alt_1. Thus, the distribution capacity is given by 14 trucks, 8 for LIN and 6 for LOX. Both 28 LIN customers and 22 LOX customers with varying consumption profiles require inventory replenishment during a time horizon of one week. Figure 12 shows the plant, depot, alternative source, and customer locations for the entire supply-chain. Overall, it includes 3 plants, 3 depots, 1 alternative source, and 50 customers. All problem data for Example 2 are provided as Supplementary Information. Similar to Example 1, for every plant and customer we assume that the inventory levels at the end of the time horizon must be at least the same than when the time horizon begins. The overall forecasted product to be replenished is 50,896 Mcf for LIN and 28,059 Mcf for LOX.

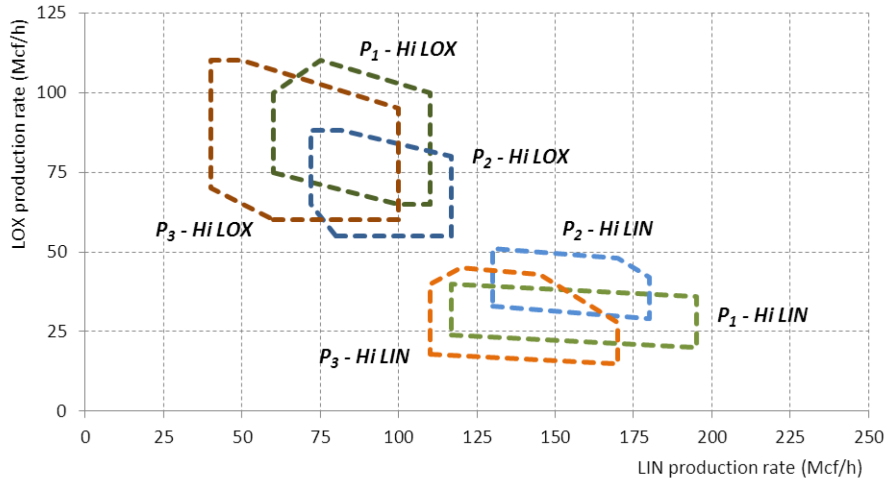

Figure 11. Production rate limits for the alternative modes and plants of Example 2.

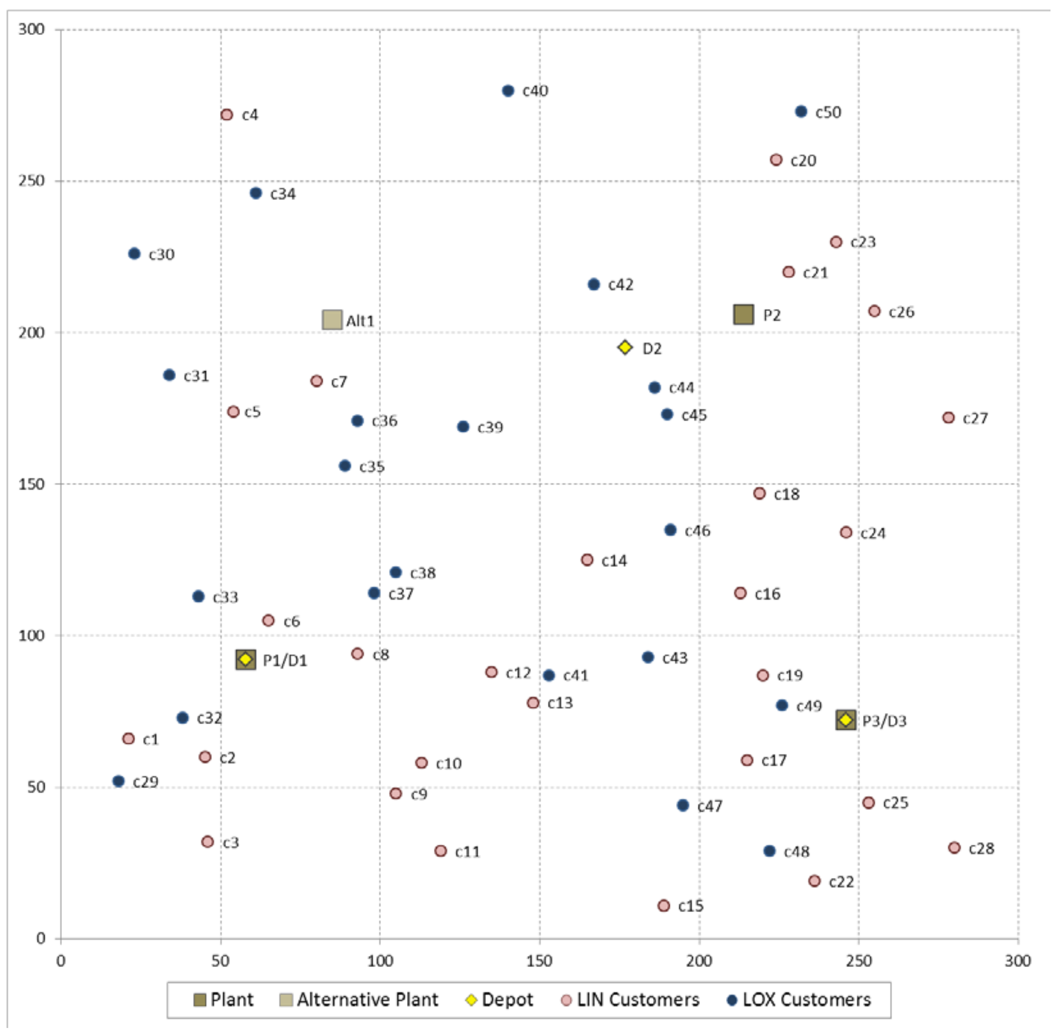

Figure 12. Supply-chain map for Example 2.

Route distances are calculated by using the straight line distance for any pair of locations. The route generation procedure described in Appendix A is used to propose a sufficiently large route set *S*. The parameters of the algorithm and size of the set of routes obtained for each product are shown in Table 11. Overall, 245 customer sets and 505 alternative routes are proposed. For conciseness we do not report results for other route sets, although it is clear that changing the parameters in Table 11 can impact the routes available for the model, and thus the quality of the distribution schedule found.

The model statistics and computational results considering a CPU time limit of 1 h. are shown in Table 12. The model features good computational performance taking into account the model size and the number of possible routes being tested. On one hand, Table 12 shows that the relaxed solution is close to the best MIP solution, which clearly indicates that the proposed MILP model has a tight relaxation. On the other hand, due to the model size and complexity, the convergence rate of the bounds is quite slow and the best possible solution is nearly midway the relaxed and the MIP solution even after 1 h. of CPU time. However, taking into account the authors' experience, a solution with a relative gap of ~2.5% is excellent for the problem being solved. For instance, while limited by the set of routes proposed, the best solution cannot improve more than $2,800. At the best solution found, plant P_1 produces and sources a total of 26,962 Mcf of liquid product (LIN + LOX), while plants P_2 and P_3 produce and source 26,664 Mcf and 25,329 Mcf, respectively.

5.2.1 Impact of electricity price variations

An alternative scenario (A) is introduced to further show the impact of using a coordinated model that simultaneously takes into account production and distribution decisions. In this case, the electricity prices (for all time periods, *t*) are increased by 1 cent for plant P_2 and decreased by 0.5 cents for plants P_1 and P_3. Example 2 is solved again with the computational results also shown in Table 12. The model features the same size reported previously. The best solution found decreases from $109,841 to $107,756 with the modified electricity prices. While the difference between both solutions is small, the impact that the change of electricity cost has on the selection of production and distribution activities throughout the entire supply-chain is significant. Figure 13 shows how the total cost of each production and distribution facility changes between both solutions, and Figure 14 presents a comparison of the product being sourced from each plant. As it can be seen, the production in Plant P_2 decreases, while the production in P_1 and P_3 increases due to the changes in electricity prices. The distribution costs of the three depots are similarly changed.

Table 11. Route generation parameters and statistics for Example 2.

		LIN	LOX
Parameters	*cmax*	3	3
	dmax	500	500
	smax	60	45
	vmin	2	2
	vmax	5	5
# of customer sets		140	105
# of routes		286	219

Table 12. Computational results for Example 2.

	Example 2	Example 2 (A) change of electricity prices	Example 2 (B) shut-down of plant P_2 starting at time t_3
Binary variables	13,832	13,832	13,808
Continuous vars.	21,533	21,533	21,533
Constraints	19,993	19,993	19,993
Relaxed LP sol.	104,070	101,032	115,392
MIP solution	109,841	107,756	123,135
Best possible sol.	107,061	104,451	118,505
Rel. gap	2.5%	3.1%	3.7%
CPU time	3,600 s	3,600 s	3,600 s
Nodes	135,991	183,854	69,979

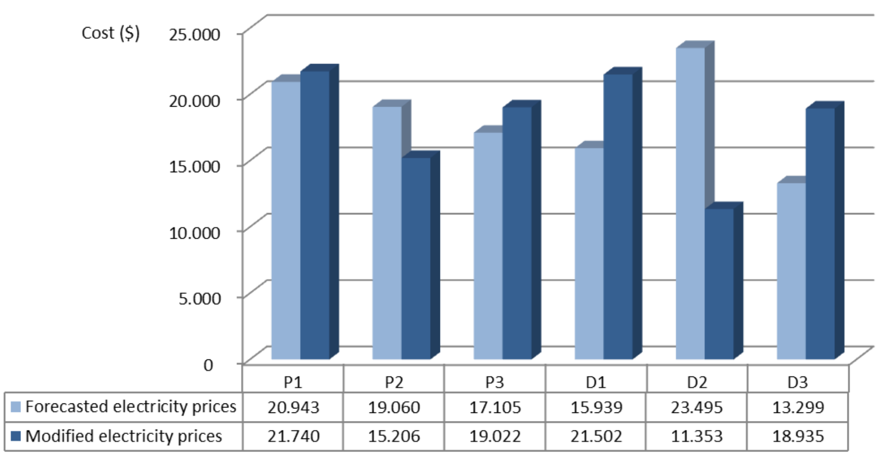

Figure 13. Comparison of total cost at each production and distribution site for the best solution of Example 2 considering alternative electricity prices.

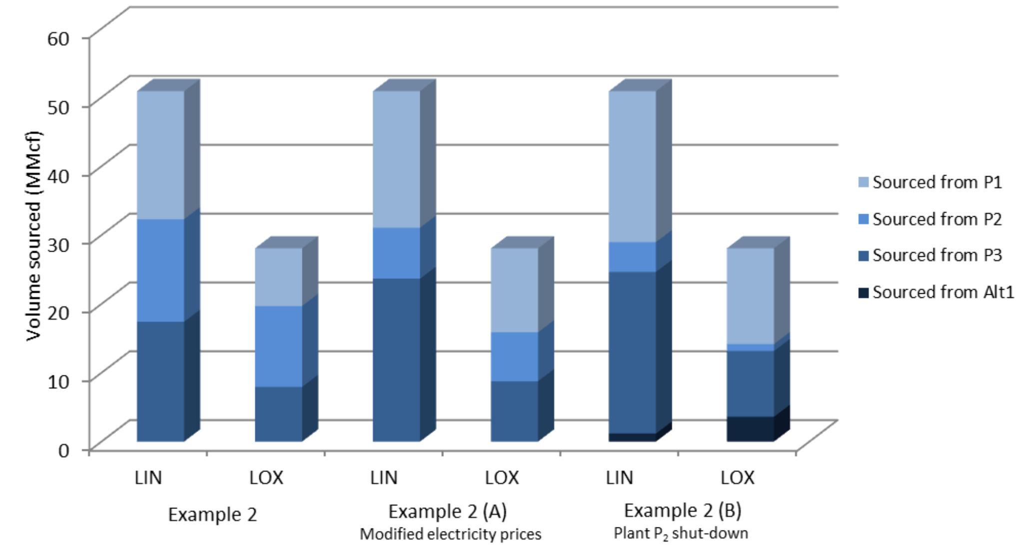

Figure 14. Product sourced per plant for each scenario of Example 2.

In order to quantify the impact of the aforementioned change of electricity prices, we take the best solution obtained with the original forecast prices and calculate the production cost of each plant P_1 to P_3, but using the modified electricity prices instead. Based on the total product sourced from each plant and because usp_{pmi} is constant, it is possible to derive from the objective function (30) that the additional cost for plant P_2 is $5,332.8, while the cost reduction for plants P_1 and P_3 is $2,696.2 and $2,532.9, respectively. Using these results to obtain the production cost of each plant (see Figure 13), both the total production cost and the simultaneous production and distribution costs are 2% higher than the solution with the modified prices. Conversely, the same can be observed with the production cost of P_1 to P_3 of the best solution for the modified electricity price scenario by using the original forecasted prices instead.

5.2.2 Production capacity disruptions

A second scenario (B) is also considered, this time assuming a shut-down is required for maintenance at plant P_2. The maintenance starts at time t_3 (start of second day) and lasts until the end of the week. To model the shut-down, the RHS of Eqn (1) is set to zero for plant P_2 at all time periods t_3-t_{14}. The same route set is used and the computational results are also shown in Table 12. The best solution features a total cost of $123,135, which is 12% higher than the best solution of the original example. The total volume sourced from each plant is also included in Figure 14. This scenario requires product to be purchased from alternative source Alt_1 in order to ensure that customer demands are satisfied. However, as shown in Figure 14, almost all deliveries come from plants P_1 and P_3. By considering a plant shut-down, Example 2 (B) illustrates a possible situation in industrial gases supply-chains in which the proposed computational tool can help to optimally re-organize production and distribution decisions.

5.3 Application of the proposed method to industrial size test cases

The proposed simultaneous production and distribution model has been applied to several real test cases involving the current supply-chain of Air Liquide, a multinational industrial gases company with operations in 70 countries. The examples include 4 to 15 plants, hundreds of customers, and more than 1000 alternative proposed routes. Because of the problem size, in some examples additional methods such as clustering and assumptions such as planned deliveries were incorporated in the model to reduce the complexity of the routing alternatives (see Appendix B). The mathematical model has been applied to several industrial-size test cases, including both historical and future scenarios. When dealing with historical test cases, some model variables were fixed based on historical data (plant withdrawals, for example). Both historical and a fully-coordinated mathematical models were solved and the results compared. Potential savings around 9% of the total historical cost were identified due to better production-distribution coordination.

In order to illustrate the complexity of the test cases considered, for medium size examples similar to Example 2, the model features good computational performances by finding solutions with a relative gap of ~2% in one hour of CPU time. When large examples are considered (100+ customers) the computational performance decreases, although the model is still able to find good quality solutions in reasonable CPU times. As an example, Table 13 includes the problem size, model size, and computational statistics of a large example related to a segregated market region. The example features a planned delivery forecast scenario (i.e. customer inventory constraints are not included) with 168 customers and 282 planned deliveries. After applying the route generation procedure, with a total of 1235 alternative routes the model requires 11,053 binary variables (routes are not available for every time *t*). Although the model size is large, a realistic and good quality solution with a 3.6% gap was obtained after 5 h. of CPU time. The best solution obtained is composed of 48.5% production cost and 51.5% distribution cost. Out of the 235 trips needed for product distribution, 171 trips (~73%) feature a truck filling ratio (truck load / truck capacity) higher than 95%. In addition, only four trailers visit an alternative source to purchase additional product, which amounts to approximately 1% of the total cost.

To improve the model accuracy at the distribution side, traveling distances were obtained using geographic information system (GIS) software. An efficient implementation of the route generation method described in Appendix A allowed exploring several thousands of potential candidate routes, depending on the selected parameters. The algorithm enumerates possible routes by traversing a search tree where each node represents a route with a given customer set. New nodes are created by adding an extra customer to each parent node. Time windows, filling ratios and traveling distances are considered when appropriate to select or reject possible routes. While there is always a correlation between the number of routes proposed and the difficulty to converge to an optimal solution, testing alternative sets of routes clearly demonstrates the relevance of an appropriate route selection to decrease the distribution cost.

Table 13. Statistics for an industrial size production-distribution coordination test case featuring planned deliveries.

Problem size	Time periods	14
	Plants	4
	Products	2
	Prod. modes	1 or 2
	Alt. sources	4
	Depots	4
	Trucks	32
	Customers	168
	Planned deliveries	282
	Customer sets	440
	Routes	1235
Model size	Binary vars.	11,053
	Cont. vars	22,086
	Constraints	16,243
CPU performance	CPU time	5 h
	Rel. gap	3.6%
	Nodes	144,178

6. Conclusions and Future Work

This paper has presented an MILP formulation for the simultaneous coordination of production and distribution decisions on industrial gases supply-chains. On the production side, the model accounts for multiple plants running various production modes while producing one or more products. Because air separation is an energy intensive process, the main component of the production cost is the cost of electricity, and thus the operation of each plant follows electricity market conditions. On the distribution side, a combined vehicle routing and inventory management problem, known as an inventory routing problem (IRP), is considered. The vendor is responsible for inventory replenishments so that customers do not run out of product. Since the entire supply-chain is included, the IRP considered here includes multiple products, and multiple sources for each product. A forecast of customer consumption is given to solve the problem. Trucks departing from depots (located or not at plants) are used to deliver product from a given source to one or more customers. A route is given by the specification of a depot, a plant, and a customer set to which the product is delivered. Because hundreds of customers are considered, the number of possible routes grows exponentially. To handle this complexity, the model selects the routes to be used from a set of proposed routes. Alternative routes for this set are chosen by a pre-processing route generation algorithm, which inspects thousands of feasible routes taking into account alternative parameters and a sorting criterion. Overall, the fully-coordinated model includes production decisions at multiple plants, and distribution decisions at multiple depots.

To asses the impact of a better coordination, different levels of production-distribution coordination were proposed. While the fully-coordinated model combines dynamic sourcing (ability to serve the same client from multiple plants) with simultaneous production and distribution, alternatives taken into account include: (a) either one or multiple plants per customer (fixed sourcing vs. dynamic sourcing), and (b) either a sequential (production before distribution) or a simultaneous (production and distribution together) approach. As was shown in Example 1, the capability of the model to perform simultaneous optimization yields significant cost savings, in both the fixed and dynamic sourcing cases.

The proposed model has been successfully illustrated with a small and a medium size test case, showing both the capabilities of the model, as well as its computational efficiency that is due to a tight MILP formulation. The latter allows to readily explore different scenarios such as changes in pricing of electricity or disruption in the plant operations, as was illustrated in Example 2. Finally, the application to industrial case studies was discussed in which despite longer computational times, savings of the order of 9% were identified. As for future directions, two areas that deserve attention are the use of decomposition techniques for reducing the computational times in large industrial problems, and addressing the uncertainty of model parameters.

Acknowledgments

Financial support from the Center for Advanced Process Decision-making (CAPD) at Carnegie Mellon University and American Air Liquide is greatly appreciated.

Appendix A - General framework for generating a list of feasible routes

The mathematical model presented in this paper requires a set of alternative routes given as input. This set represents the possible routes that can be selected by the model to obtain a feasible solution. While it is possible to enumerate all of them, the number of routes grows exponentially with the number of depots, plants, and customers. Therefore, it is convenient to reduce the alternatives by filtering out those routes that are more unlikely to be part of the optimal distribution schedule. By limiting the set of routes proposed, the size of the model and the computational effort required to find its best solution both decrease. However, this approach can potentially cut off some of the routes needed to obtain an optimal solution, and thus the optimality of the proposed model is limited by the quality of the set of routes proposed.

A general framework to generate a set of routes based on the data of depots, plants and customers is described here. As mentioned in Section 4.2, each route is defined by a tuple (*d*, *p*, *s*), where *d* is a depot, *p* is a plant (source) and *s* is a set of customers. We assume here that plant, depot, and customer locations are given and it is possible to calculate the traveling distance between any pair of them. Thus, given a tuple (*d*, *p*, *s*) it is also possible to calculate the distance dis_{dps} for the shortest path to deliver product from plant *p* to the customers of *s* using a truck from depot *d*.

The procedure ROUTEGENERATION is presented in Table A1. The main parameters of the proposed method, which can vary for each plant *p* and product *i*, are:

- *cmax*: maximum number of customers visited in a trip,
- *dmax*: maximum distance for the shortest path of the route,
- *smax*: maximum number of routes,
- *vmin*: required number of routes for each customer *c* and time $t \in T_c$, and
- *vmax*: limit on the number of routes for each customer *c* and time $t \in T_c$.

The proposed ROUTEGENERATION procedure iterates over all combinations of plants and products adding the routes obtained to a list of routes *R*. At each iteration (i.e. for a given plant *p* and product *i*), all possible routes subject to a limit *cmax* (a given positive integer) on the number of visited customers are inspected. Customer sets *s* are generated as combinations (subsets) of *n* = 1, 2, …, *cmax* elements taken from C_i. After inspecting the possible routes, GENERATEFEASIBLEROUTES returns a set with all the tuples *r* = (*d*, *p*, *s*) that verify the following conditions: (a) trucks from depot *d* can source from plant *p* ($p \in P_{di}$), (b) all customers of *s* can receive product from plant *p* ($p \in P_s$, where $P_s \neq \emptyset$ as required in Section 4.2.7), (c) the customer set *s* verifies $|s| \leq cmax$, and (d) the TSP distance of route *r* does not exceed the limit ($dis_{dps} \leq dmax$). Additional conditions can be imposed so that the number of feasible routes does not become too large. Once all feasible routes are obtained, the resulting set *FR* is sorted based on a criterion selected beforehand. To implement the SORT procedure, both the route distance and the logistics ratio (i.e. cost per volume sourced) were the alternatives evaluated to quantify the convenience of selecting a given route. The logistics ratio, generating the most economically convenient routes, was used in the test cases. It is calculated using the maximum volume that can be delivered to the customer set *s* in a given time *t*. Given the sorted list of routes *FR*, two selection stages are executed to choose the routes required by the model. SELECTMIN ensures the selection of at least *vmin* routes for each customer *c* and time period $t \in T_c$ when a delivery can be made to this customer. Only if the set of feasible routes *FR* does not include enough alternatives, *vmin* different routes are not found. SELECTMAX completes the selection of routes seeking at least *vmax* routes for each customer *c* and time period $t \in T_c$. However, it finishes earlier whenever the number of selected routes reaches the maximum quantity *smax*. The procedure TESTROUTE is as an auxiliary procedure used for both selection methods. The list of routes *R* is returned by the algorithm, from which customer sets S_{di} are derived.

Table A1: Route Generation procedure.

```
procedure ROUTEGENERATION
input: {integers} cmax, dmax, smax, vmin, vmax
output: {route-list} R
begin
        R ← [] {empty list}
        for each p∈P, i∈I_p
                FR ← GENERATEFEASIBLEROUTES(p, i, cmax, dmax)
                SORT(FR)
                SELECTMIN(FR, R, vmin)
                SELECTMAX(FR, R, smax, vmax)
        end for
        return R
end
```

```
procedure SELECTMIN
input: {route-list} FR, R, {integer} vmin
output: {route-list} FR, R
begin
        i ← 1
        while i ≤ size(FR)
                r ← FR(i)
                if TESTROUTE(r, vmin) then
                        Select route r
                        Add r to list R
                end if
                i ← i + 1
        end while
end
```

```
procedure SELECTMAX
input: {route-list} FR, R, {integer} smax, vmax
output: {route-list} FR, R
begin
        i ← 1
        while (i ≤ size(FR)) and (# of selected routes in FR < smax)
                r ← FR(i)
                if TESTROUTE(r, vmax) then
                        Select route r
                        Add r to list R
                end if
                i ← i + 1
        end while
end
```

procedure TESTROUTE
input: {route} *r*, {integer} *limit*
output: {boolean: *whether to select a route or not*}
begin
    if *r is not selected* **and**
      ∃ *customer c and time t such that*:
        (*c can receive a delivery at time t using route r*
        and # of selected routes for c at time t < limit)
    then return true
    else return false
end

Appendix B - Large scale problems

In this section two alternatives to reduce the complexity of the distribution side problem by incorporating additional assumptions are described: clustering and planned deliveries. These methods can be used, either solely or combined, to facilitate the solution of industrial size problems otherwise limited by the computational effort needed to solve a large MILP model.

Clustering methods may be used to reduce the model size when a given problem instance includes hundreds of customers, which leads to a large increase of the number of alternative routes. Let $q \in Q$ be a group or cluster of customers and C_q the subset of customers belonging to cluster *q*. We assume that the set *Q* is obtained a priori through the application of some clustering algorithm (Jain et al., 1999) and that every customer belongs to a unique cluster (at least for each time *t*.) The location of a cluster *q* is calculated as a weighted average of the locations of the customers belonging to *q*. In turn, the weight κ_c given to a customer $c \in C_q$ is based on an estimation of the minimum number of deliveries necessary to replenish the consumption of *c* over the entire time horizon. Eqn (B-1) defines the location ($\bar{x}_q$) of each cluster and the minimum number of deliveries (κ_c), where $Q_c^{avg,\min}$ is the average redline, R_c^{avg} is the average consumption for customer $c \in C_i$, and $U_i^{\max} = \max_{k \in K_i}\{U_k^{truck}\}$ is the capacity of the largest truck available to deliver product *i*.

$$\bar{x}_q = \frac{\sum_{c \in C_q} \bar{x}_c \cdot \kappa_c}{\sum_{c \in C_q} \kappa_c} \quad \forall q \in Q \, ; \quad \kappa_c = \left\lceil \frac{\sum_{t \in T} R_{ct}}{\min\{Q_c^{\max} - Q_c^{avg,\min} + R_c^{avg}, U_i^{\max}\}} \right\rceil \quad \forall c \in C_i, i \in I \qquad \text{(B-1)}$$

Routes distances from a given depot *d* and plant *p* are obtained using the location of the cluster given by Eqn (B-1). An internal distance can be added to each route visiting cluster *q*, to account for the distance traveled between customers inside the cluster. Two types of routes are considered to deal with clusters on the distribution side: (i) intra-cluster routes, that only deliver product to all or a subset of the customers C_q and (ii) inter-cluster routes, where two or more clusters are visited on a given round-trip delivery (i.e. a truck visits one or more customers of cluster q_1 and then one or more customers of cluster q_2, etc.)

To handle these alternatives, we extend the routing scheme presented by adding simple conditions when customer sets are defined. Alternative (i) means that there is only one customer set *s* for each cluster *q*. Moreover, there is a one-to-one correspondence between customer sets *s* and clusters *q*. Since it is less straightforward, alternative (ii) is discussed in more detail. In this case, for every cluster *q* and customer set *s*, either $C_q \subset s$ or $C_q \cap s = \varnothing$. In other words, customer sets are defined based on cluster data, so that each *s* includes all customers $c \in C_q$ or none of them (an alternative point of view is that each set *s* now includes clusters instead of customers).

Given an appropriate definition of the sets *s*, the variable d_{sqt} is introduced representing the volume delivered to some or all the customers of *q* through *s*. Thus, Eqns (16) and (17) are replaced by the following constraints:

$$\sum_{d \in D:(s \in S_{di})} \sum_{k \in K_{di}} e_{kst} = \sum_{q \in Q_i:(C_q \subset s)} d_{sqt} \qquad \forall\, s \in S_i, i \in I, t \in T \tag{B-2}$$

$$\sum_{c \in C_q} D_{c,t} = \sum_{s \in S_i:(C_q \subset s)} d_{sqt} \qquad \forall\, q \in Q_i, i \in I, t \in T \tag{B-3}$$

Notice that Eqns (B-2) and (B-3) must be considered together with customer inventory constraints (23) and (24). No additional changes are introduced in the model, only the cluster locations given by Eqn (B-1) are used to calculate route distances. Overall, by aggregating customers into clusters the number of delivery sets *s* is significantly reduced, which in turn reduces the number of binary variables y_{kst}. The tradeoff between the accuracy of route distance calculations and the CPU time required to solve the problem must be evaluated to select between using a detailed customer-based routing approach or an approximate cluster-based method.

The second alternative to handle large test cases is a reduction of scope of the distribution side problem by assuming that the amount of product to be delivered to each customer throughout the time horizon is given. In this case, customer inventory constraints are not needed, and the problem data only specifies the forecast of planned deliveries instead of the customer consumption profiles. Thus, the distribution side full inventory routing problem reduces to a smaller vehicle routing problem with time windows (VRP-

TW). The complexity of the problem decreases, mainly because the number of routes available at any given time is restricted by the possible deliveries (open time windows) at that time.

We assume that, for each delivery of product *i* to customer *c*, the volume to be delivered and the specific time window during which the delivery takes place are given. The parameter U^{deliv}_{c,t_1,t_2} introduced in Section 4.4 is used, where the length of the time intervals (t_1, t_2) is usually one day. Let T_c be the set of time periods *t* when a delivery can be made to customer *c*, then the accumulated volume of product *i* that must be delivered to customer *c* up to time period *t* is calculated as:

$$U^{sum}_{c,t_0} = 0 \qquad \forall\, c \in C \tag{B-4}$$

$$U^{sum}_{c,t} = U^{sum}_{c,t-1} + \sum_{t' \le t} U^{deliv}_{c,t',t} \qquad \forall\, c \in C, t \in T : t > t_0 \tag{B-5}$$

In order to guarantee that the right amount of product is delivered to customer *c* by the end of each time window (t_1, t_2), constraint (B-6) is used. This constraint is defined when $t \in T^*_c$, where $T^*_c = \left\{ t \mid t \in T_c \wedge \exists t' : U^{deliv}_{c,t',t} > 0 \right\}$ includes the upper bound limits of all the time windows of customer *c*, and it works properly even if two deliveries have overlapping time windows.

$$\sum_{\substack{t' \in T_c \\ t' \le t}} D_{c,t'} = U^{sum}_{c,t} \qquad \forall\, c \in C, t \in T^*_c \tag{B-6}$$

To handle planned deliveries, customer inventory constraints (23) and (24) are replaced by Equations (B-4) to (B-6). To use these equations it is important to ensure that, for all model constraints, variables y_{kst}, e_{kst}, d_{sct}, and d_{sqt} are only defined at time periods such that $t \in T_s$, where $T_s = \bigcap_{c \in s} T_c$.

As a particular case, if every time window is restricted to a unique period (i.e., $t_1 = t_2$), then Equation (B-6) reduces to Equation (B-7). In this case, the amount of product delivered to customer *c* at time period *t* ($D_{c,t}$) becomes a problem parameter.

$$D_{c,t} = U^{deliv}_{c,t,t} \qquad \forall\, c \in C, t \in T^*_c \tag{B-7}$$

Finally, planned deliveries are particularly useful when combined with clustering methods. For instance, notice that when definition (B-7) is applied the LHS of Eqn (B-3) can be calculated a-priori.

Nomenclature

Subscripts

c	customer
d	depot
i	product
j	product grade
k	truck
m	production mode
p	plant
q	cluster
s	customer set (subset of customers visited in a given route)

Sets

C	customers
C_i	customers for product i
C_q	customers belonging to cluster q
D	depots
I	products
I_p	products of plant p
I_{pm}	products produced by plant p while running in mode m
J_c	product grades that can be delivered to customer c
J_i	product grades of product i
K	trucks
K_{di}	trucks for product i available at depot d
M	production modes
M_{pt}	production modes available at plant p in time period t
P	all plants
P^{alt}	alternative sources
P^{own}	plants owned by the company
P_{di}	plants associated to depot d and product i
$P_{di,s}$	plants from which a truck from depot d can source product i to customer set s
Q	clusters of customers
Q_i	clusters of customers for product i
S_i	customer sets for product i
S_{di}	customer sets available for product i and depot d
S_{pi}	alternative customer sets to source product i from plant p
T	time periods
T_c	time periods when a delivery to customer c is possible
T_s	time periods when a delivery to all customers in s is possible

Parameters

$\alpha_{pm,i\lambda}$	coefficient of the production rate of product i for the limiting hyperplane λ

δ_{dps} difference between actual distance dis_{dps} and minimum distance $dis_{ds}^{\min}$

$\delta_{dp,i}^{\max}$ maximum δ_{dps} for all possible sets $s \in S_{di}$

Δ_t duration of time period t

$\mu(t_1, t_2)$ total product consumed by customer c in the interval $[t_1, t_2]$

η_p turndown ratio for plant p

$\pi_{pm,\lambda}$ upper bound for the hyperplane λ limiting the feasible rates of production mode m, plant p

κ_c estimation of the minimum number of deliveries for customer c

b_p^{init} whether plant p is running (1) or shut down (0) at time t

c_k traveling cost per distance unit for truck k

$C_{pi,t}^{purchase}$ cost of product i if purchased at alternative source p in time t

dis_{dps} shortest traveling distance of route (d, p, s) obtained by application of a TSP method (customers of set s are visited using the shortest path starting at plant p and finishing at depot d)

$dis_{ds}^{\min}$ minimum distance required for a truck of depot d to deliver product from any valid source to customer set s

$F_{p,t}^{start}$ start-up cost of plant p at time t

H time horizon

L_c^{ini} initial inventory of customer c

L_{pi}^{ini} initial inventory of product i at plant p

$Q_{ct}^{\min}$ redline (safety stock level) for customer c at time t

$Q_{pi,t}^{\min}$ safety stock in time period t for product i at plant p

$Q_c^{\max}$ storage capacity of customer c

$Q_{pi}^{\max}$ storage capacity of product i at plant p

$Q_{pi,t}^{purchase}$ maximum volume of product i available at alternative source p in time t

R_{ct} product consumption forecast of customer c at time t

$R_{pi,t}^{site}$ forecast of gaseous customer pipeline demand for product i at plant p in time period t

U_{c,t_1,t_2}^{deliv} volume of product required by customer c between time t_1 and time t_2 (planned delivery)

$U_{c,t}^{sum}$ accumulated volume required by customer c at time t

U_k^{truck} trailer capacity for vehicle k

$U_{pi,t}^{withdrawal}$ fixed truck withdrawal volume of product i from plant p at time period t

u_{pt} electricity price forecast of plant p during time period t

usp_{pmi} unit specific power

$w_{pmi}^{\max}$ maximum production rate of product i at plant p running production mode m

$w_{pmi}^{\min}$ minimum production rate of product i at plant p running production mode m

Binary Variables

b_{pt}^{start} denotes that plant *p* starts operation at time period *t*

B_{pmt} denotes that plant *p* operates in mode *m* during time period *t*

Y_{kpt} denotes that truck *k* loads product at plant *p* in time period *t*

y_{kst} denotes that truck *k* visits the customers in set *s* during time period *t*

Continuous Variables

β_{kt} additional distance traveled by truck *k* to load product from a given plant at time *t*

$D_{c,t}$ total volume delivered to customer *c* in time period *t*

d_{sct} volume delivered to customer *c* distributed among customers of set *s* in time *t*

d_{sqt} volume delivered to cluster *q* distributed among members of set *s* in time *t*

$D_{pi,t}^{site}$ volume of product *i* to be gasified and sent by pipeline at plant *p* in time period *t*

$D_{pi,t}^{truck}$ volume of product *i* withdrawn for truck delivery from plant *p* at time period *t*

$DCost_t$ total distribution cost at time *t*

DIS_{kt} distance traveled by truck *k* at time *t*

E_{kpt} volume of product withdrawn from plant *p* and loaded into truck *k* at time *t*

e_{kst} volume of product delivered by truck *k* to the customers *s* in time *t*

L_{ct} inventory of customer *c* at time *t*

L_{pit} inventory of product *i* available at plant *p* at the end of time period *t*

$PCost_t$ total production cost at time *t*

$PW_{p,t}$ power consumption of plant *p* at time period *t*

$W_{pmi,t}$ production rate of product *i* at plant *p*, when *p* is running in mode *m* at time *t* (zero otherwise).

References

Anily, S., & Federgruen, A. (1993). Two-echelon distribution systems with vehicle routing costs and central inventories. *Operations Research, 41*(1), 37-47.

Bell, W. J., Dalberto, L. M., Fisher, M. L., Greenfield, A. J., Jaikumar, R., Kedia, P., Mack, R. G., & Prutzman, P. J. (1983). Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer. *Interfaces*, *13*(6), 4-23.

Benoist, T., Gardi, F., JeanJean, A., & Estellon, B. (2011). Randomized local search for real-life inventory routing. *Transportation Science*, *45*(3), 381-398.

Bertazzi, L., Paletta, G., & Speranza, M. G. (2002). Deterministic order-up-to level policies in an inventory routing problem. *Transportation Science, 36*(1), 119-132.

Bertazzi, L., Savelsbergh, M, & Speranza, M. G. (2008). Inventory routing. In: Golden, B., Raghavan, S., & Wasil, E. (Eds). *The vehicle routing problem, latest advances and new challenges. Operations Research/Computer Science Interfaces* (vol. 43, pp. 49-72). Springer US.

Campbell, A. M., Clarke, L., Kleywegt, A. J., & Savelsbergh, M. W. P. (1998). The inventory routing problem. In Crainic, T. G. & Laporte, G. (Eds.), *Fleet management and logistics* (pp. 95-112). Boston: Kluwer Academic Publishers.

Campbell, A. M., Clarke, L. W., & Savelsbergh, M. W. P. (2002). Inventory routing in practice. In P. Toth and D. Vigo (Eds.), *The Vehicle Routing Problem* (pp. 309-330). Philadelphia: Society for Industrial and Applied Mathematics.

Campbell, A. M. & Savelsbergh, M. W. P. (2004). A decomposition approach for the inventory routing problem. *Transportation Science*, *38*(4), 488-502.

Castle, W. F. (2002). Air separation and liquefaction: recent developments and prospects for the beginning of the new millennium. *International Journal of Refrigeration*, *25*, 158-172.

Çetinkaya, S., & Lee, C. (2000). Stock replenishment and shipment scheduling for vendor-managed inventory systems. *Management Science*, *46*(2), 217-232.

Chan, L. M. A, Federgruen, A., and Simchi-Levi, D. (1998). Probabilistic analysis and practical algorithms for inventory-routing models. *Operations Research*, *46*(1), 96-106.

Chandra, P., Fisher, M. L. (1994). Coordination of production and distribution planning. *European Journal of Operational Research*, *72*, 503-517.

Coelho, L., Cordeau, J.-F., & Laporte, G. (2013). Thirty years of inventory routing. *To appear in Transportation Science*.

Daryanian, B., Bohn, R. E., & Tabors, R. D. Optimal demand-side response to electricity spot prices for storage-type customers. *IEEE Transactions on Power Systems*, *4*(3), 897-903.

Dror, M. & Ball, M. (1987). Inventory/routing: reduction from an annual to a short period problem. *Naval Research Logistics Quarterly*, *34*(6), 891-905.

Dror, M., Ball, M., & Golden, B. (1985). Computational comparison of algorithms for the inventory routing problem. *Annals of Operations Research*, *4*(1), 3-23.

Fumero, F. & Vercellis, C. (1999). Synchronized development of production, inventory, and distribution schedules. *Transportation Science*, *33*(3), 330-340.

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

Ierapetritou, M. G., Wu, D., Vin J., Sweeney, P., & Chigirinskiy, M. (2002). Cost minimization in an energy-intensive plant using mathematical programming approaches. *Industrial and Engineering Chemistry Research*, *41*(21), 5262-5277.

Jaillet, P., Bard, J. F., Huang, L., & Dror, M. (2002). Delivery cost approximations for inventory routing problems in a rolling horizon framework. *Transportation Science*, *36*(3), 292-300.

Jain, A. K., Murty, M. N., & Flynn, P. J. (1999). Data clustering: a review. *ACM Computing Surveys, 31*(3), 264-323.

Karwan, M. H., & Keblis, M. F. (2007). Operations planning with real time pricing of a primary input. *Computers and Operations Research*, *34*(3), 848-867.

Kleywegt, A. J., Nori, V. S., & Savelsbergh, M. W. P. (2002). The stochastic inventory routing problem with direct deliveries. *Transportation Science*, *36*(1), 94-118.

Mitra, S., Grossmann, I. E., Pinto, J. M., & Arora, N. (2012a). Optimal production planning under time-sensitive electricity prices for continuous power-intensive processes. *Computers and Chemical Engineering*, *38*, 171-184.

Mitra, S., Grossmann, I. E., Pinto, J. M., & Arora, N. (2012b). Integration of strategic and operational decision-making for continuous power-intensive processes. *Computer Aided Chemical Engineering*, *30*, 232-236.

Park, Y. B. (2005). An integrated approach for production and distribution planning in supply-chain management. *International Journal of Production Research*, *43*(6), 1205-1224.

Ross, J. G. & Lane, I. E. (1998). Industrial power demand response analysis for one-part real time pricing. *IEEE Transactions on Power Systems*, *4*, 897-903.

Smith, A. R. & Klosek, J. (2001). A review of air separation technologies and their integration with energy conversion processes. *Fuel Processing Technologies*, *70*(2), 115-134.

Thomas, D. J. & Griffin, P. M. (1996). Coordinated supply chain management. *European Journal of Operational Research*, *94*, 1-15.

You, F., Pinto, J. M., Capón, E., Grossmann, I. E., Arora, N., & Megan, L. (2011). Optimal distribution-inventory planning of industrial gases. I. Fast computational strategies for large-scale problems. *Industrial and Engineering Chemistry Research*, *50*(5), 2910-2927.

Supplementary Information – Problem data for Example 2

Table S1. Inventory at each production facility (Mcf).

	P_1		P_2		P_3	
	LIN	LOX	LIN	LOX	LIN	LOX
Initial	10,000	8,000	8,500	6,500	6,500	7,000
Maximum	18,000	12,000	12,000	9,000	14,000	10,000
Redline	5,000	3,500	3,000	4,000	4,000	3,000

Table S2. Production rate limits for each plant and production mode, given as vertices of the product space (Mcf/h).

Production mode		P_1		P_2		P_3	
		LIN	LOX	LIN	LOX	LIN	LOX
Hi LOX	1	60	75	80	55	60	60
	2	100	65	117	55	100	60
	3	110	65	117	80	100	95
	4	110	100	82	88	50	110
	5	75	110	72	88	40	110
	6	60	100	72	65	40	70
Hi LIN	1	117	24	130	33	110	18
	2	195	20	180	29	170	15
	3	195	36	180	42	170	28
	4	117	40	170	48	145	43
	5			130	51	120	45
	6					110	40

Table S3. Additional parameters for plants.

	USP[a] (kWh/Mcf)	Start-up cost ($)	Initial state
P_1	20	7,000	operating
P_2	20	4,000	shut-down
P_3	20	6,000	operating

[a] The same unit specific power coefficient applies to all production modes.

Table S4. Location of plants and depots (miles).

	Owned plants/depots				Alternative sources
	P_1/D_1	P_2	D_2	P_3/D_3	Alt_1
X coord.	58	214	177	246	85
Y coord.	92	206	195	72	204

Table S5. Vehicles available at each depot.

	Number of trucks/trailers		Trailer capacity	Cost
	LIN	LOX	(Mcf)	($/mile)
D_1	3	2	630	2.75
D_2	2	2	630	2.85
D_3	3	2	630	2.65

Table S6. Customer location and storage data.

		Location (miles)		Inventory (Mcf)		
		X coord.	Y coord.	Initial	Maximum	Redline
LIN Customers	c_1	21	66	1,200	1,674	504
	c_2	45	60	837	1,715	413
	c_3	46	32	381	837	168
	c_4	52	272	561	873	261
	c_5	54	174	264	420	105
	c_6	65	105	2,900	4,250	961
	c_7	80	184	1,788	2,793	837
	c_8	93	94	249	453	114
	c_9	105	48	498	699	141
	c_{10}	113	58	1,611	2,571	642
	c_{11}	119	29	459	873	174
	c_{12}	135	88	1,803	2,514	753
	c_{13}	148	78	234	453	90
	c_{14}	165	125	1,615	3,837	710
	c_{15}	189	11	1,338	2,514	753
	c_{16}	213	114	255	420	105
	c_{17}	215	59	546	699	210
	c_{18}	219	147	1,518	3,354	672
	c_{19}	220	87	1,980	3,072	615
	c_{20}	224	257	627	837	168
	c_{21}	228	220	1,245	2,514	630
	c_{22}	236	19	1,158	1,674	436
	c_{23}	243	230	429	873	174
	c_{24}	246	134	1,335	2,514	630
	c_{25}	253	45	1,884	3,072	615
	c_{26}	255	207	1,725	2,514	753
	c_{27}	278	172	1,407	2,514	504
	c_{28}	280	30	276	420	84
LOX Customers	c_{29}	18	52	384	546	108
	c_{30}	23	226	274	612	196
	c_{31}	34	186	318	519	105
	c_{32}	38	73	597	1,140	342
	c_{33}	43	113	1,122	2,070	519
	c_{34}	61	246	1,323	2,244	450
	c_{35}	89	156	756	1,035	228
	c_{36}	93	171	420	558	111
	c_{37}	98	114	360	519	186
	c_{38}	105	121	819	1,035	321
	c_{39}	126	169	1,276	2,149	790
	c_{40}	140	280	918	1,176	420
	c_{41}	153	87	733	1,258	261
	c_{42}	167	216	648	1,822	442
	c_{43}	184	93	270	519	105
	c_{44}	186	182	546	759	153
	c_{45}	190	173	282	573	114
	c_{46}	191	135	1,734	2,760	690
	c_{47}	195	44	795	1,035	312
	c_{48}	222	29	966	1,935	563
	c_{49}	226	77	744	1,038	312
	c_{50}	232	273	474	900	222

Table S7. Consumption forecast for each customer and time period (Mcf).

		t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}	t_{11}	t_{12}	t_{13}	t_{14}
LIN Customers	c_1	96	156	132	108	108	102	114	156	90	126	168	132	168	150
	c_2	102	150	105	114	150	129	123	123	102	105	144	165	111	150
	c_3	75	54	45	63	48	78	54	54	66	72	69	60	60	48
	c_4	57	45	57	51	72	69	78	57	72	78	72	57	48	48
	c_5	33	33	21	39	24	24	33	36	24	21	33	33	30	30
	c_6	227	227	197	407	263	368	242	194	359	182	215	164	233	293
	c_7	162	165	186	183	186	273	279	192	192	246	195	267	156	243
	c_8	24	36	24	24	33	33	36	27	33	30	27	45	33	27
	c_9	48	48	45	69	42	45	39	48	39	48	45	42	39	57
	c_{10}	195	135	228	249	204	207	174	204	216	195	150	213	195	144
	c_{11}	87	54	72	78	84	51	57	51	54	72	75	48	51	63
	c_{12}	219	141	186	225	189	192	126	195	183	249	150	147	129	147
	c_{13}	27	39	39	27	30	39	42	36	45	24	27	24	36	27
	c_{14}	180	168	312	264	168	204	324	228	264	216	276	192	240	228
	c_{15}	126	159	228	228	234	153	183	132	165	180	165	144	228	177
	c_{16}	33	24	36	27	39	39	21	30	36	33	21	39	42	33
	c_{17}	45	51	36	63	54	60	60	57	39	69	60	69	66	57
	c_{18}	324	246	180	228	240	234	186	294	300	174	246	216	318	168
	c_{19}	276	237	213	240	297	207	156	300	240	204	306	270	171	306
	c_{20}	72	72	75	66	45	63	45	72	51	51	75	84	72	72
	c_{21}	144	183	162	237	135	177	168	186	207	207	168	138	240	153
	c_{22}	138	162	162	156	156	114	126	156	132	90	132	132	138	96
	c_{23}	57	57	66	60	51	72	66	69	75	72	45	81	75	51
	c_{24}	201	210	150	171	204	219	177	144	126	126	246	249	246	234
	c_{25}	174	162	240	183	306	279	201	168	219	207	171	243	222	216
	c_{26}	240	153	153	195	204	168	144	243	156	135	165	132	228	156
	c_{27}	177	165	183	216	201	141	147	165	198	159	129	213	210	168
	c_{28}	36	24	36	36	39	33	27	33	30	24	39	27	24	33
LOX Customers	c_{29}	36	45	54	39	33	42	36	54	30	42	39	33	42	42
	c_{30}	60	60	42	48	48	60	36	48	54	48	30	54	42	42
	c_{31}	36	42	42	39	42	51	33	45	51	51	33	33	33	30
	c_{32}	99	72	99	93	75	99	57	105	57	81	78	96	81	87
	c_{33}	105	180	198	162	153	177	156	195	183	117	132	174	150	204
	c_{34}	141	120	117	129	162	132	177	210	117	153	150	168	141	213
	c_{35}	78	51	78	96	54	90	66	63	75	72	81	57	99	54
	c_{36}	51	51	45	39	39	51	54	45	36	42	39	39	30	51
	c_{37}	39	45	48	36	51	36	36	45	45	42	30	42	42	45
	c_{38}	63	60	99	75	57	99	69	102	75	78	63	63	84	93
	c_{39}	198	180	222	222	156	198	288	198	168	216	204	174	234	144
	c_{40}	90	114	114	66	78	108	78	72	114	78	108	66	84	72
	c_{41}	102	90	108	96	84	96	102	90	90	84	90	108	84	108
	c_{42}	156	204	108	168	156	108	120	180	180	180	216	132	156	132
	c_{43}	48	45	48	27	51	42	27	42	30	33	45	27	36	42
	c_{44}	57	51	72	45	72	54	54	42	42	48	51	72	69	69
	c_{45}	45	57	42	39	51	51	30	54	30	45	30	39	48	45
	c_{46}	180	153	192	264	174	153	159	192	144	141	159	147	252	192
	c_{47}	72	102	63	69	84	63	57	72	78	60	78	99	54	78
	c_{48}	180	204	180	192	132	180	192	150	174	144	138	132	132	150
	c_{49}	102	72	54	96	90	102	72	60	54	96	54	96	78	54
	c_{50}	54	78	60	72	72	90	66	72	72	48	90	72	48	72

Table S8. Electricity prices (cent/kWh) at each production facility.

	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}	t_{11}	t_{12}	t_{13}	t_{14}
P_1	4.15	3.72	4.28	3.86	3.98	3.66	3.91	3.62	4.15	3.74	4.30	3.84	4.04	3.86
P_2	3.08	2.85	2.98	2.53	3.18	2.86	2.90	2.68	3.15	2.79	3.03	2.50	3.04	2.65
P_3	3.75	3.50	3.64	3.04	3.84	3.25	3.53	3.45	3.52	3.01	3.54	3.13	3.89	3.13

Table S9. Alternative source parameters.

Plant Alt_1	LIN	LOX
Selling price ($/Mcf)	1.6	1.8
Availability (Mcf)	4,000	3,000