

A hybrid heuristic for a stochastic production-inventory-routing problem

Agostinho Agra ^{a,2}, Cristina Requejo ^{a,2}, Filipe Rodrigues ^{a,2}

^a *Department of Mathematics, University of Aveiro, 3810-193 Aveiro, Portugal.*

Abstract

We consider a stochastic single item production-inventory-routing problem with a single producer and multiple clients. At the clients, demand is allowed to be backlogged incurring a penalty cost. Demands are considered uncertain.

A recourse model is presented where the production and routing decisions are taken before the scenario is known, and the quantities to deliver to the clients and the inventory levels are adjustable to the scenario. Valid inequalities are introduced and a hybrid heuristic that combines ideas from the sample average approximation method and from relax-and-fix approaches is proposed.

Preliminary tests based on randomly generated instances are reported showing that the hybrid heuristic performs better than the classical sample approximation algorithm for hard instances.

Keywords: Inventory routing, stochastic programming, sample approximation algorithm, hybrid heuristic.

¹ Published in *Electronic Notes in Discrete Mathematics* **64**, 345–354, 2018. [doi: 10.1016/j.endm.2018.02.009]

² The research of the authors has been partially supported by Portuguese funds through the *Center for Research and Development in Mathematics and Applications (CIDMA)* and FCT, the Portuguese Foundation for Science and Technology, within project UID/MAT/04106/2013. The research of the third author has been funded by FCT under Grant PD/BD/114185/2016.

1 Introduction

We consider a single item stochastic production-inventory-routing problem (SPIRP) with a single supplier/producer and multiple retailers/clients. A vendor managed inventory approach is followed where the supplier monitors the inventory at the retailers and decides on the replenishment policy for each retailer. Inventory aspects are considered both at the supplier and at the retailers. Demand is allowed to be backlogged at the retailers and in that case a penalty cost is incurred. Backlog is not allowed at the supplier. Demands are considered uncertain, following a uniform distribution. A constant production capacity at the supplier is assumed. The decision maker has to decide the production and the distribution plans for a finite time horizon. The production plan consists in defining the production periods and the amount to produce in each one of those periods. The distribution plan defines the retailers that should be visited in each time period, the quantities to deliver to each visited retailer, and the corresponding route in each time period. For the distribution plan a single vehicle is considered. The goal is to minimize the production and routing cost plus the expected inventory and penalty for backlogged demand costs. We assume the production plan and the choice of which clients to visit in each time period (and consequently the routing) are first stage decisions, that is, such decisions must be taken before the scenario is revealed. The quantities to deliver to each client at each time period, and the inventory levels, can be adjusted to the scenario (known as second stage decisions). Such assumptions may hold for short planning horizons.

Complex problems combining production, inventory and routing decisions have been receiving a great attention in recent years. For a recent survey see [?]. Stochastic approaches for related inventory routing problems have been also considered, e.g. [?,?,?,?]. Among these [?,?,?] present heuristic approaches. A classical approach for handling stochastic problems is the sample average approximation (SAA) method, see [?]. In this method the expected inventory and penalty for backlogged demand costs are approximated by a sample average estimate obtained from a random sample. The resulting sample average approximating problem (SAAP) is solved for different samples in order to obtain a set of candidate solutions. Then these candidate solutions are tested on a larger sample and the best solution for that larger sample is chosen.

However, most of the deterministic production-inventory-routing problems reported in the literature are not solved to optimality. Therefore, solving each SAAP to optimality, within reasonable running time, may be impractical for

many instances. In [?] it is noticed that when the SAAP is not solved to optimality for each sample, then the use of a commercial solver based on a branch and cut algorithm with a running time limit acts as a heuristic which may lead to lack of stability, since one can obtain solutions whose objective function can significantly vary from sample to sample. Here we propose a hybrid heuristic (HH) that combines different approaches, namely the SAA method and relax-and-fix heuristics. First, following the SAA method, several small size samples are considered. For each sample, a relaxation is used to obtain a partial feasible solution. Then, the obtained partial solutions are explored in order to identify variables that are frequently fixed to zero or to one. Those variables are fixed and the restricted model is solved for a larger sample keeping all the remaining variables free. To improve the efficiency of both methods (SAA and HH), valid inequalities are added to the initial formulation.

To the best of our knowledge, the use of the information from several solutions obtained with the SAA method combined with relax-and-fix heuristics is new. The proposed heuristic has several advantages in relation to the classical SAA method. It does not require solving each sample set to optimality, it does not require the use of large sample sets, and it allows one to adjust the final solution to a larger sample set since many variables are kept free. Further, the number of variables that are fixed can be easily controlled.

The mathematical model is presented in Section 2 as well as the valid inequalities included. The SAA method and the proposed HH are introduced in Section 3. Computational tests that compare the SAA method and the HH and show the advantages of the HH are presented in Section 4. Final conclusions are given in Section 5.

2 Mathematical formulation and improvements

In this section we introduce the mathematical model for the SPIRP. We assume the production plan (production periods and quantities to produce) and the routing (which clients to visit in each period) are first stage decisions, that is, decisions taken before the demands are known, while the quantities to deliver, the amount backlogged at each time period as well as the inventory levels are adjusted to the scenario. The goal of the stochastic approach is to find the solution that minimizes the production and routing cost plus the expected value of the inventory cost plus the penalty value for backlogged demand. Following the SAA method, the true expected cost value is replaced by the mean value of a large random sample $\Omega = \{\xi^1, \dots, \xi^s\}$ of scenarios, ob-

tained by the Monte Carlo method. This larger set of s scenarios is regarded as a benchmark scenario set representing the true distribution [?].

Consider a graph $G = (N, A)$ where $N = \{0, 1, \dots, n\}$ represents the set of nodes and A represents the set of possible links, set A is a subset of set $N \times N$. Node 0 denotes the producer and $N_c = \{1, \dots, n\}$ is the set of clients. $T = \{1, \dots, nt\}$ is the set of periods.

Consider the following parameters: $d_{it}(\xi)$ is the demand of client $i \in N_c$ in period $t \in T$ in scenario $\xi \in \Omega$, I_0/I_i is the initial stock at producer/client i , \bar{S}_0/\bar{S}_i is the inventory limit at producer/client i , \bar{P} is the production limit at each time period, \underline{Q}_{it} and \bar{Q}_{it} are the lower and upper limits to the delivery quantity in period $t \in T$ at client i . For the objective function parameters define: S is the set up cost for producing in a period, P is the production cost of a unit of product, V is the vehicle usage cost, C_{ij} is the travel cost from node i to node j , $(i, j) \in A$, H_i is the holding cost of the product at node $i \in N$, B_i is the backlog cost of the product at node $i \in N$.

Consider the following non-adjustable variables: binary variables y_t indicate whether there is production in period $t \in T$, variables p_t give the production level in period $t \in T$, binary variables z_{it} indicate whether there is a visit to node $i \in N_c$ in period $t \in T$, the routing variables x_{ijt} indicate whether a vehicle travels from node i to node j , $(i, j) \in A$, in period $t \in T$, v_t is a binary variable that is one if the vehicle is used in period $t \in T$ and zero otherwise; for $(i, j) \in A$, $t \in T$, f_{ijt} is the artificial flow of a single commodity variable used to prevent cycles in the routing problem. Notice this is not the amount transported from node i to node j in period $t \in T$ since that quantity depends on the scenario. Such adjustable variables could be used but would imply the use of an unnecessarily large model.

For each scenario $\xi \in \Omega$, we define the adjustable variables $q_{it}(\xi)$ indicating the quantity delivered at node $i \in N_c$ in period $t \in T$; $s_{it}(\xi)$ indicating the stock level of a node $i \in N$ at the end of period $t \in T$; and $r_{it}(\xi)$ is the quantity backlogged at node $i \in N_c$ in period $t \in T$.

The SPIRP is modeled by the following formulation.

$$\min \sum_{t \in T} (S y_t + P p_t + V v_t + \sum_{(i,j) \in A} C_{ij} x_{ijt}) + \frac{1}{|\Omega|} \sum_{\xi \in \Omega} \sum_{t \in T} \sum_{i \in N} (H_i s_{it}(\xi) + B_i r_{it}(\xi)) \quad (1)$$

$$s.t. \quad s_{0,t-1}(\xi) + p_t = \sum_{i \in N_c} q_{it}(\xi) + s_{0t}(\xi) \quad \forall t \in T, \xi \in \Omega \quad (2)$$

$$s_{i,t-1}(\xi) + r_{it}(\xi) + q_{it}(\xi) = d_{it}(\xi) + s_{it}(\xi) + r_{i,t-1}(\xi) \quad (3) \\ \forall i \in N_c, t \in T, \xi \in \Omega$$

$$s_{i0}(\xi) = I_i \quad \forall i \in N, \xi \in \Omega \quad (4)$$

$$s_{it}(\xi) \leq \bar{S}_i \quad \forall i \in N, t \in T, \xi \in \Omega \quad (5)$$

$$p_t \leq \bar{P} y_t \quad \forall t \in T \quad (6)$$

$$\underline{Q}_{it} z_{it} \leq q_{it}(\xi) \leq \bar{Q}_{it} z_{it} \quad \forall i \in N_c, t \in T, \xi \in \Omega \quad (7)$$

$$\sum_{j \in N} x_{ijt} = z_{it} \quad \forall i \in N_c, t \in T \quad (8)$$

$$\sum_{j \in N} x_{jit} - \sum_{j \in N} x_{ijt} = 0 \quad \forall i \in N, t \in T \quad (9)$$

$$\sum_{j \in N} x_{0jt} = v_t \quad \forall t \in T \quad (10)$$

$$\sum_{i \in N} f_{ijt} - \sum_{i \in N_c} f_{jit} = z_{jt} \quad \forall j \in N_c, t \in T \quad (11)$$

$$f_{ijt} \leq n x_{ijt} \quad \forall i, j \in N, t \in T \quad (12)$$

$$y_t, v_t \in \{0, 1\} \quad t \in T \quad (13)$$

$$z_{it} \in \{0, 1\} \quad \forall i \in N, t \in T \quad (14)$$

$$x_{ijt} \in \{0, 1\} \quad \forall i, j \in N, t \in T \quad (15)$$

$$p_t \geq 0 \quad t \in T \quad (16)$$

$$f_{ijt} \geq 0 \quad \forall i, j \in N, t \in T \quad (17)$$

$$r_{it}(\xi), s_{it}(\xi), q_{it}(\xi) \geq 0 \quad \forall i \in N, t \in T, \xi \in \Omega \quad (18)$$

For brevity we omit the explanation of the model. This model can be tightened through the inclusion of valid inequalities. From the underlying network flow structure of the SIRP one can derive several families of valid inequalities. Next we introduce three such families.

$$r_{it}(\xi) \geq \left(\sum_{\ell=1}^t d_{i\ell}(\xi) - I_i \right)^+ \left(1 - \sum_{\ell=1}^t z_{i\ell} \right), \forall i \in N_c, t \in T, \xi \in \Omega \quad (19)$$

$$\sum_{i \in N} s_{i,k-1}(\xi) + \sum_{i \in N_c} r_{it}(\xi) \geq \sum_{i \in N_c} \sum_{\ell=k}^t d_{i\ell}(\xi) \left(1 - \sum_{\ell=k}^t y_\ell \right), \forall k, t \in T, k \leq t, \xi \in \Omega \quad (20)$$

$$\sum_{i \in N_c} r_{it}(\xi) \geq r \left(\left\lceil \frac{d}{\bar{P}} \right\rceil - \sum_{\ell=1}^t y_\ell \right), \forall t \in T, \xi \in \Omega \quad (21)$$

where $d = \left(\sum_{i \in N_c} \sum_{\ell=1}^t d_{i\ell}(\xi) - \sum_{i \in N} I_i \right)^+$, $r = d - \left(\left\lceil \frac{d}{\bar{P}} \right\rceil - 1 \right) \bar{P}$, and $(x)^+ = \max\{x, 0\}$. Inequalities (19) are adapted from the uncapacitated lotsizing

problem with backlogging [?], and are derived for each client and force the net demand (demand that cannot be satisfied by the initial stock) at client i until period t to be satisfied with backlog in case there is no delivery until period t , that is, if $\sum_{\ell=1}^t z_{i\ell} = 0$. Inequalities (20) ensure that the total demand from period k to period t must be satisfied either from stock (at the producer or at the clients) or from backlog at the clients if there is no production in that period. For the particular case of $k = 1$, inequalities (20) can be

written in the stronger format $\sum_{i \in N_c} r_{it}(\xi) \geq d \left(1 - \sum_{\ell=1}^t y_\ell \right)$. Inequalities (21)

are mixed integer rounding inequalities (see [?]) derived from the constraint $\bar{P} \sum_{\ell=1}^t y_\ell + \sum_{i \in N_c} r_{it}(\xi) \geq d$ which is a relaxation of the feasible set that imposes the total set up capacity plus the backlogged demand to cover the net demand of all clients until period t .

These inequalities can be further strengthened and/or extended. We present only the ones used in the computational tests, reported in Section 4, and omit further discussion. The formulation (1)-(18), (19), (21) is denoted by SPIRF.

3 Solution approaches

In this section we discuss two approaches to solve the SPIRF that will be tested in Section 4. A first approach is a classical SAA method described in [?]. In the SAA method m separate sample sets $\Omega_k, k \in M = \{1, \dots, m\}$, are considered, each one containing $\ell \ll s$ scenarios. The model is solved for each one of the scenarios set, Ω_k , (replacing Ω by Ω_k in model SPIRF) with a time limit of α seconds, giving m candidate solutions. For each candidate solution, the first stage solution is fixed, and the value of the objective function for a very large sample with s scenarios is computed by solving a pure linear programming problem on the recourse variables. The solution with the minimum average cost is chosen.

Next we describe the hybrid heuristic HH in four stages.

First stage: obtain m solutions. Generate m samples of small dimension ℓ_1 . For each sample a partial solution is obtained by solving the SPIRF assuming the routing variables x_{ijt} are continuous (relaxing constraints (15)) with a time limit of α_1 seconds. Variables z_{it} are kept binary.

Second stage: fix a set of z_{it} variables. First assign a different weight w^k to each one of the m partial solutions obtained in the first stage. Those weights are computed according to the objective function value, denoted by c^k , computed on a larger sample of dimension $\ell_2 \geq \ell_1$. Those weights are normal-

ized as follows: let $\bar{c} = \max_{k \in M} \{c^k\}$ and $\bar{d} = \min_{k \in M} \{\bar{c} - c^k \mid c^k \neq \bar{c}\}$. Then for each solution k define its weight $w^k = \frac{\bar{c} + \bar{d} - c^k}{\sum_{i \in M} (\bar{c} + \bar{d} - c^i)}$. So $\sum_{k=1}^m w^k = 1$. Let z_{it}^k be the value of variable z_{it} in solution k . Then, for each pair (i, t) , $i \in N, t \in T$, compute the weighted average of the values of z_{it} obtained in the first stage: $W(i, t) = \sum_{k=1}^m w^k z_{it}^k$. The values $W(i, t)$ are between 0 and 1. Define two threshold values γ_1 and γ_2 and set $z_{it} = 0$ if $W(i, t) \leq \gamma_1$, and $z_{it} = 1$ if $W(i, t) \geq \gamma_2$. In order to define the percentage of variables to be fixed *a priori*, values γ_1 and γ_2 can be obtained from the empirical distribution of the $W(i, t)$ values.

Third stage: solve the simplified model for a larger sample. With the variables z_{it} fixed in the previous stage and without the integrality constraints (15), solve the resulting model SPIRF for a new and larger sample with $\ell_3 \geq \ell_2$ scenarios. Then, with all the variables z_{it} fixed to their optimal value, the routing variables are determined by solving a TSP problem for each time period.

Fourth stage: compute objective function value. Fix the first stage decision variables. Then, compute the objective function value for the very large sample, by computing the value of the recourse variables for each one of the s scenarios. This can be done by solving a pure linear programming model again.

4 Computational Results

In this section we report preliminary computational experience performed to assess the quality of the proposed HH, and compare it with the SAA method. All tests were run using a computer with an Intel Core i7-4750HQ 2.00GHz processor and 8GB of RAM, and were conducted using the Xpress-Optimizer 28.01.04 solver with the default options.

We present computational results for 7 sets of instances to the SPIRF, A_{20}, \dots, A_{80} . For each set A_n with n clients, 5 instances are generated from different samples of demand values. The coordinates of the n clients are randomly generated in a 100 by 100 square grid, and the producer is located in the center of the grid. For each value of n , a complete graph is obtained and symmetric traveling costs are associated to the set of arcs. The traveling costs C_{ij} are the Euclidean distance between nodes i and j in the grid. The number of periods nt considered is 5. For each client and each period, the nominal demand value d_{it} is randomly generate between 40 and 80 units, and we consider uncertain demands varying in $[0.7d_{it}, 1.3d_{it}]$. The initial stock at producer I_0 is zero, and the initial stock I_i at client i is randomly generated

between 0 and three times the average demand of client i . The maximum inventory level \bar{S}_i is 500 for all $i \in N$. The production capacity \bar{P} is 50% of the average demand. For all $i \in N$, the holding cost H_i is 2 and the backlog cost B_i is 3. The production set up cost, the unit production cost and the vehicle usage cost are given by $S = 100$, $P = 1$ and $V = 50$, respectively.

Preliminary tests (not reported here) have shown that inequalities (19)-(21) reduce the integrality gap by more than forty percent, on average. In general, when a given inequality (19) - (21), derived for a scenario ξ , cuts off the linear relaxation solution, then the same happens for most of the inequalities derived for the remaining scenarios in Ω . Hence, the number of cuts can be quite large when the number of scenarios is large. For the SPIRF we opted to omit constraints (20) as the number of inequalities added was too large for some instances.

The obtained computational results are displayed in Table 1. For each problem instance $s = 1000$ scenarios were generated. The number of candidate solutions used is $m = 10$.

When using the SAA method, each candidate solution is obtained by using a sample of $\ell = 5$ and 25 scenarios and a time limit of $\alpha = 300$ seconds. However, when no integer solution is found within $\alpha = 300$ seconds the model is solved until the first integer solution is obtained. For each strategy, the obtained results, average computational time in seconds (in columns “ \bar{T} ”) and average solution cost (in columns “ \bar{C} ”), are displayed in columns SAA_1 and SAA_2 in Table 1. Columns SAA_1 and SAA_2 display the results for $\ell = 25$ and $\ell = 5$, respectively.

For the HH, in the first stage, each candidate solution is obtained by using a sample of $\ell_1 = 5$ scenarios and a time limit of $\alpha_1 = 300$ seconds. In the second stage a sample of $\ell_2 = 15$ scenarios is used. In the third stage, when solving the simplified model, a sample of $\ell_3 = 50$ scenarios is used. In the second stage, several strategies for fixing variables z_{it} were compared corresponding to different values of parameters γ_1 and γ_2 . The results for two of these strategies are reported in Table 1. Columns HH_1 and HH_2 display the results for $\gamma_1 = 0.05$ and $\gamma_1 = 0.15$, respectively. The value of γ_2 is set to $1 - \gamma_1$.

Now we compare strategies SAA_1 and SAA_2 . For instances with a small number of clients ($n = 20, 30$), the problems used to obtain the candidate solutions can be solved to optimality. Therefore, using the larger sample sets with 25 scenarios (thus more representative samples), better solution values are obtained. For instances with a higher number of clients the problems to obtain the candidate solutions could not be solved to optimality. In these

Table 1

Computational results for the SAA method and the HH. For each instance set, the lowest average cost solution is displayed in bold and, in parenthesis, we show the number of instances for which the corresponding solution method produces the best solution.

n	SAA_1		SAA_2		HH_1		HH_2	
	\bar{T}	\bar{C}	\bar{T}	\bar{C}	\bar{T}	\bar{C}	\bar{T}	\bar{C}
A_{20}	406	9145 (2)	431	9175 (0)	49	9153 (2)	49	9166 (2)
A_{30}	804	13749 (5)	576	13772 (0)	108	13783 (0)	108	13783 (0)
A_{40}	3716	19326 (0)	4143	19261 (0)	513	19210 (3)	510	19206 (4)
A_{50}	4194	21268 (0)	4220	20678 (0)	672	20514 (4)	704	20510 (1)
A_{60}	7182	25139 (0)	5968	24037 (0)	950	23468 (2)	964	23434 (3)
A_{70}	13499	30627 (0)	9576	31782 (0)	1310	27680 (1)	1314	27881 (4)
A_{80}	17396	36162 (0)	16879	35736 (0)	2027	32190 (2)	2032	32224 (3)

cases, better solutions are obtained when a smaller number of scenarios is used (strategy SAA_2).

Results not reported here for the HH with different threshold values, γ_1 and γ_2 , and solving, in the third stage, the model SPIRF with and without the integrality constraints (15), show that the best results occur when the integrality constraints (15) are not used.

Reported results for the SAA method and the proposed HH allow us to conclude that for instances with a small number of clients the SAA method solves problems until optimality and lower cost solutions are obtained. However, for harder instances, with a higher number of clients, the proposed HH performs better than the classical SAA method, since it provides lower cost solutions with smaller computational times. Although the variance values for each set of instances are not reported in Table 1, they reveal that lower computational time variance values are always obtained by the HH and the same happens, in general, for the solution cost variance values, where the variance of the HH are in average lower than half of the variance costs of the SAA. Furthermore, the HH allows to obtain a large number of different low costs solutions in short time by varying the value of parameters γ_1 and γ_2 .

5 Conclusion

We consider a stochastic production-inventory-routing problem for which a recourse model is introduced and valid inequalities are presented. A heuristic that combines the sample average approximation method (SAA) with relax-and-fix approaches is proposed. This heuristic uses the statistical information

obtained from several solutions derived for small scenario samples in order to determine which variables to fix. The heuristic is fast and quite flexible, and can be easily extended to other complex stochastic problems.