International Workshop on Lot Sizing - IWLS 2021

August 30th - September 1st 2021, online

Foreword

Dear members of the lot sizing community,

Let's start with the bad news. Due to obvious reasons, we will hold this 2021 edition of the International Workshop on Lot Sizing on-line. This means that we will not have any boat cruises, guided museum outings, city tours or visits to a whiskey distillery, beer brewery, or port cellar. We will not be able to sit all together and watch the presentations in a little castle, near the beach or even in a regular university aula. That is, of course, unfortunate.

On the upside, we have a great program with 15 excellent presentations, all on lot-sizing. As the previous 2020 on-line edition has shown, this still allows to have interesting discussions and fruitful exchanges (you just need to unmute your microphone). And that is of course what it is all about. Moreover, you won't have to make any long journey to participate, and you can enjoy this workshop in the comfort of your own home or office. We would like to thank the organisations listed on the next pages who are supporting this workshop, and finally, we hope you have a great workshop.

Raf and Matthieu.

Workshop organisation

'Local' organizing committee

- Raf Jans, HEC Montréal, Canada (raf.jans@hec.ca)
- Matthieu Gruson, ESG UQÀM, Canada (gruson.matthieu@uqam.ca)

Members of the board of the EURO working group on lot sizing

- Nabil Absi (France)
- Kerem Akartunali (UK)
- Christian Almeder (Germany)
- Stéphane Dauzère-Pérès (France)
- Silvio de Araujo (Brasil)
- Raf Jans (Canada)
- Safia Kedad Sidhoum (France)
- Wilco van den Heuvel (Netherlands)

List of sponsors

EURO Working Group on Lot Sizing

EURO
The Association of European Operational Research Societies

GERAD
Groupe d'Études et de Recherche en Aide à la Décision

Chair in Supply Chain Operations Planning

CIRRELT
Centre Interuniversitaire de Recherche sur les Réseaux d'Entreprise, la Logistique et le Transport

HEC Montréal
École des hautes études commerciales

ESG UQÀM
École des sciences de gestion

Program

Day 1: Monday, August 30, 2021

8h - 9h30 Session 1: Integrated lot sizing problems

Chair: Silvio de Araujo

9h30 - 10h Break

10h - 11h Session 2: Exact approaches for lot sizing problems

Chair: Christian Almeder

Day 2: Tuesday, August 31, 2021

8h - 9h30 Session 1: Stochastic lot sizing I

Chair: Raf jans

9h30 - 10h Break

Heuristic for the 3-Level Integrated Lot-Sizing and Cutting Stock Problem and Extensions

Gislaine Mara Melega
UFSCar - Federal University of São Carlos, Production Engineering Department, 13565-905, SP, Brazil
gislainemelega@gmail.com

Silvio Alexandre de Araujo
UNESP - São Paulo State University, Mathematics Department, 15054-000, São José do Rio Preto, SP, Brazil
silvio.araujo@unesp.br

Raf Jans
HEC Montreal and GERAD/CIRRELT, Montreal, H3T 2A7, QC, Canada
raf.jans@hec.ca

Reinaldo Morabito
UFSCar - Federal University of São Carlos, Production Engineering Department, 13565-905, SP, Brazil
morabito@ufscar.br

Abstract

In this study, we deal simultaneously with lot-sizing and cutting stock decisions in an integrated way, so as to capture the interdependency between these decisions in order to obtain a better global solution. For this, a generalized 3-level integrated problem (*G3ILSCS*) proposed in the literature is extended to take into account other relevant decisions of the supply chain. The extensions consist of the supplier selection of the raw materials used in the cutting process and the distribution of the final products from the production plant to a warehouse. To solve the integrated problem, a hybrid heuristic is proposed aiming to overcome the difficulties present in the integrated problem, mainly comprising the high number of variables, the multi-level structure, and the integrality requirements. The hybrid method combines two decomposition approaches, the column generation and the relax-and-fix procedures, in each iteration of the algorithm. The models and solution approaches are analyzed in an extensive computational study aiming to evaluate the impact of incorporating other decisions of the supply chain into the integrated problem.

Introduction

The idea of integrating processes in a production plant is to take into account, simultaneously, the decisions related to the problems involved so as to capture the interdependency between the decisions in order to obtain a better global solution. The manufacturing setting addressed in this study has its production processes linked to the cutting of raw materials (objects) and the production planning of end products (final products). In these industries, objects of large sizes are kept in stock to be cut later into smaller pieces of different sizes, using cutting patterns, in order to meet internal demand. These pieces then go to downstream levels of the production plant in order to produce and assemble the final products. The production planning of final products takes into account the tradeoff between setup and inventory holding costs to meet the clients' demand, considering capacity limitations. Therefore, it is necessary to plan the acquisition, production, and cutting of these objects, as well as the production of final products, in order to minimize the negative effects of these processes, which can be seen as the waste of material, delays in downstream levels, high costs, among others. These two problems are known in the literature as the cutting stock problem and the lot-sizing problem, respectively ([6, 3, 4, 1]).

A generalized 3-level integrated problem, *G3ILSCS*, proposed in [5], is extended and computationally analyzed considering other relevant decisions/levels of the supply chain. The first extension consists of an alternative means to the acquisition of objects, besides producing these objects in the company, and comprehends the supplier selection of objects used in the cutting process. In this alternative level, suppliers can also provide the inputs (objects) to the production plant. The other extension addresses an additional level after the production of final products, related to the distribution of the final products from the production plant to a warehouse. The problem with all these features comprise an 4-level integrated lot-sizing and cutting stock problem with supplier selection and distribution (*G4ILSCS*).

Mathematical Models and Solution Methods

The *G3ILSCS* model consists of a production environment composed of three levels and multi-periods in a deterministic setting. Level 1 corresponds to the production planning of objects, that have to be produced considering a capacitated environment in order to fulfill the downstream level (level 2). Level 2 is associated with the cutting process, in which the produced objects are cut into pieces according to cutting patterns [2] by a cutting machine with limited resources. The cut pieces can be used as components to assemble the final products or directly as final products. It is at level 3 that the production of the final products occurs and the independent demand for final products has to be met in each time period. The link between the different

time periods is provided by inventory at each level. There is a bill-of-material relationship, for which the dependent demand of final products triggers a dependent demand for pieces, and indirectly, for objects. Therefore, the decisions of the 3-level integrated problem determines simultaneously a production planning that defines for every time period: the production quantities for products and objects and the cutting patterns, with corresponding frequencies, considering limited resources, while searching for a global optimal minimal solution to the 3-level integrated problem.
In the *G4ILSCS*, an alternative means to the acquisition of objects needed in the cutting process is considered, i.e., objects can be produced at level 1 of the production plant, as well as be purchased from external suppliers at level 1A. The supplier selection at level 1A takes into account a set of suppliers offering different types of objects, which can be purchased considering fixed and variable costs, proportional to the quantity purchased. The price of objects does not vary according to the number of objects ordered, i.e., there is no discount rate. In this integrated approach, the supplier selection decisions define the number of objects to be purchased considering the different supplier costs, which allied to an optimized production planning of the cutting process, and the production of the final products and objects, aims to reduce the total costs in the integrated problem.
The other extension modeled in the *G4ILSCS* manages the distribution of final products. Level 4 is responsible for the distribution costs incurred from shipments between the production plant and the warehouse. The distribution decisions are related to the load/arrangements of final products into vehicles, i.e., they are associated with the number of vehicles needed to transport the final products, hence, such decisions are directly linked to the production lot-sizing decisions of final products. In this integrated problem, the demand of final products is shifted from the production plant to the warehouse and in both sites, there is the possibility of inventory. Inventory at the production plant arises from cargo consolidation on the vehicles, whereas at the warehouse, inventory is addressed to keep the final products in stock in order to meet the clients' orders. Therefore, the distribution decisions define the number of vehicles utilized in the transport of final products and the transportation costs incurred in each time period of the planning horizon, i.e., this problem combines the lot-sizing and cutting stock decisions with vehicle loading decisions.
Considering that the problems addressed in this study are classified as NP-hard, we proposed a hybrid heuristic solution method to solve the integrated problem, aiming to overcome the difficulties, comprising the multi-level structure, the hight number of the variables, and integrality requirements. The goal of the hybrid heuristic is to provide a good trade-off between solution quality and computational effort while solving the integrated problem. The column generation procedure is used as a first step to generate an initial matrix of columns (cutting patterns for cutting the objects into pieces and cargo configurations for loading the final products into de vehicles)

at levels 2 and 4 of the integrated problem. After that, the column generation is applied in each step of the relax-and-fix procedure in the hybrid heuristic, aiming to find more attractive columns, for cutting patterns and cago configurations, while the hybrid heuristic searches for a feasible solution to the integrated problem.
The models and solution methods are analyzed in an extensive computational study. Therefore, the main objective of this study is to evaluate the impact of incorporating levels of the supply chain into the integrated problem, as well as assess the performance of the hybrid approach in different environments when solving the 3-level integrated lot-sizing and cutting stock problem with supplier selection and distribution decisions.

Acknowledgments

This research was funded by the Conselho Nacional de Desenvolvimento Científico e Tecnológico (*CNPq*), Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (*CAPES*) and the Fundação de Amparo a Pesquisa do Estado de São Paulo (FAPESP process n 2016/01860-1, 2018/19893-9, 2019/03302-4).

References

[1] N. Brahimi, N. Absi, S. Dauzère-Pérès, and A. Nordli. Single-item dynamic lot-sizing problems: An updated survey. *European Journal of Operational Research*, 263(3):838 – 863, 2017.

[2] P. C. Gilmore and R. E. Gomory. A linear programming approach to the cutting-stock problem. *Operations Research*, 9(6):849 – 859, 1961.

[3] A. M. Gomes, J. F. Gonçalves, R. Alvarez-Valdés, and J. M. V. de Carvalho. Preface to the special issue on cutting and packing. *International Transactions in Operational Research*, 23:3 – 4, 2016.

[4] R. Jans and Z. Degraeve. Modeling industrial lot sizing problems: a review. *International Journal of Production Research*, 46(6):1619–1643, 2008.

[5] G. M. Melega, S. A. de Araujo, and R. Jans. Classification and literature review of integrated lot-sizing and cutting stock problems. *European Journal of Operational Research*, 271(1):1 – 19, 2018.

[6] G. Wäscher, H. Haußner, and H. Schumann. An improved typology of cutting and packing problems. *European Journal of Operational Research*, 183(3):1109 – 1130, 2007.

Novel formulations for demand fulfillment in an integrated procurement and lot-sizing problem

Caio Paziani Tomazella
University of São Paulo
caio.tomazella@usp.br

Maristela Oliveira Santos
University of São Paulo
mari@icmc.usp.br

Douglas Alem
University of Edinburgh Business School
douglas.alem@ed.ac.uk

Raf Jans
HEC Montréal
raf.jans@hec.ca

Abstract

This study addresses a problem that integrates decisions on supplier selection, multi-stage assembly production, and demand fulfillment with multiple customers and backlogging. This problem is motivated by a case of a manufacturing company that produces refrigeration equipment. To avoid using backlogging penalties, traditionally used in cost-minimization formulations, which involve intangible costs such as loss of customer satisfaction, we propose models that optimize service levels in the form of minimum backlog and maximum fill rates. We use budgetary constraints to keep operational costs under a limit. The results show that the proposed models bring improvements to the service levels, and they can also be used to enforce service equity among customers. We also propose an adaptation of our model that considers customer and product importance when optimizing demand fulfillment.

Introduction

In this article, we approach the Integrated Procurement and Lot-Sizing Problem (IPLSP) focusing on demand fulfillment decisions. This approach differs from the traditional models found in the literature since our proposed models optimize service level objectives while keeping the costs, which are usually minimized, under a limit. The motivation of this work is a case presented by a refrigeration equipment manufacturer, which assembles products with extensive and complex bill-of-materials, consisting of components that are produced in house and others that are purchased from third-party suppliers. At the same time, there is a need to fulfill the demand from multiple customers while working within set and restricted budgets.
Usually, production and procurement decisions in a manufacturing plant are made separately. However, the literature shows that integrating both problems leads to more cost efficient solutions, especially by lowering the costs involved in the purchase of raw materials [3]. Aside from this integration, we treat the demand individually by customer, instead of an aggregated value. This turns demand fulfillment decisions into a variation of the one-warehouse multi-retailer problem [6] in which each customer has its associated demand and backlog, although we do not manage their inventory and cannot make shipments in advance.
Due to an inherent difficulty in calculating inventory shortage penalties [7], especially when different customers with different profiles are involved, we avoid using backlogging costs and propose models that minimize the total backlog, and maximize fill rates. Fill rates are indicators that measure the percentage of demand fulfilled in time and can be calculated in several ways, either globally or for individual products [5]. Since our model involves different customers, we also present functions that maximize the worst individual cases, which aim at enforcing an equity between service levels between customers.

Problem Definition and Modeling

In order to solve the IPLSP, we propose a Mixed Integer Programming (MIP) formulation for it. Since the production process involves the assembly of several components, we use a multistage lot-sizing model, as presented by [1]. The purchase of components is modeled as a supplier selection problem with delivery lead-times and total quantity discounts, which are given due to the fact that material is usually bought in large quantities [4]. With the demand separated by customers, we use the formulation from [2], with variables indicating how much is sent to each customer per period and their associated backlogs. Also, orders can be partially fulfilled, allowing the model to prioritize certain demands when optimizing service levels in periods of limited resource availability.

In a practical case, production and procurement decisions are made in different time scales (e.g., weeks and months) therefore we divide our planning horizon in macro-periods for the supplier selection variables and in micro-periods for the production and demand fulfillment variables.
The first model, MIP_{TC} is a traditional approach that minimizes the total costs from the IPLSP, including the backlog costs. The second, MIP_B, minimizes the total backlog, while keeping all costs under a budget. In our modeling, a value is made available at the start of each macro-period, and the amount that is not spent can be carried to the following period without any loss.
The other models aim at maximizing fill rates, which are identified in the literature by β. Mathematically, a fill rate is calculated using $\beta = 1 - \frac{backorder}{demand}$, and its variations are obtained depending on which demand is considered. The Global Fill Rate (β) uses the entire demand, the Customer Fill Rate (β^c) calculates one for each customer and takes the minimum value among them, while the Customer-Product Fill Rate (β^{cj}) does the same for each customer and product combination.
Given this definition, we propose models MIP_G, MIP_C and MIP_{CP} that maximize, β, β^c and β^{cj}, respectively. Note that MIP_C and MIP_{CP} are Maximin type models. Also, additional constraints are used to obtain the backorder values from the backlogging variables.

Experimentation Results

Several computational experiments were carried in order to evaluate the performance and solutions of the proposed models. Since our datasets, which included data from the manufacturer, do not contain budget values, we used the costs obtained by solving MIP_{TC} (minus the backlogging costs) as a benchmark. For all models, *a posteriori* calculations were done in order to obtain the fill rate values in each solution, so they could be compared.
In our first analysis, we found out that our service level based models were able to improve their respective indicators when compared to MIP_{TC}, even though the solutions had the same budget to work with. For example, in one of our datasets, MIP_C was able to find an improved value of β^c than MIP_{TC} in all instances. While these results were more evident on the fill rate models (especially the Maximin ones), these improvements were less frequent within the MIP_B solutions.
Regarding the trade-offs between fill rates, we showed that some solutions of MIP_G had low values of β^c and β^{cp}, meaning that, when optimizing the global demand fulfillment, some demands are left unfulfilled for longer periods of time. On the other hand, MIP_C and MIP_{CP} obtain β^c and β^{cp} values, respectively, comparable to β without deteriorating the latter too much.
Lastly, we did some experimentation that showed that these conclusions also apply

to cases with more strict budgets and lost sales instead of backlogging, as well as proposing additional models that maximize the weighted sums of Customer and Customer-Product Fill Rates. These models are relevant in cases when customers or products have different strategical importance to the manufacturer and it needs to be taken into consideration when making the demand fulfillment decisions.

References

[1] K. Akartunali and A. J. Miller. A heuristic approach for big bucket multi-level production planning problems. European Journal of Operational Research, 193(2): 396–411 (2009)

[2] S. Chand, V. N. Hsu, S. Sethi, and V. Deshpande. A dynamic lot sizing problem with multiple customers: Customer-specific shipping and backlogging costs, IIE Transactions, 39 (11): 1059–1069, (2007).

[3] A. L. Cunha, M. O. Santos, R. Morabito, and A. Barbosa-Póvoa. An integrated approach for production lot sizing and raw material purchasing, European Journal of Operational Research, 269 (3): 923–938, 2018.

[4] D. R. Goossens, A. J. T. Maas, F. C. R. Spieksma, and J. J. van de Klundert. Exact algorithms for procurement problems under a total quantity discount structure, European Journal of Operational Research, 178 (2): 603–626, 2007.

[5] M. Gruson, J. F. Cordeau, and R. Jans. The impact of service level constraints in deterministic lot sizing with backlogging, Omega, 79: 91–103, 201.

[6] O. Solyali and H. Süral. The one-warehouse multi-retailer problem: Reformulation, classification, and computational results, Annals of Operations Research, 196 (1): 517–541, 2012.

[7] H. Tempelmeier and S. Herpers. ABCβ-a heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint, International Journal of Production Research, 48 (17): 5181–5193, 2010.

A Multi-Plant Lot Sizing Problem Applied to an Aluminum Beverage Packaging Company

Carlos Eduardo de Morais Ferreira
Federal University of São Paulo (UNIFESP) / Aeronautics Institute of Technology (ITA)
cem.ferreira@unifesp.br

Antonio Augusto Chaves
Federal University of São Paulo (UNIFESP)
antonio.chaves@unifesp.br

Mariá Cristina Vasconcelos Nascimento
Federal University of São Paulo (UNIFESP)
mcv.nascimento@unifesp.br

Abstract

Many multinational companies require multiple plants to meet the worldwide demand for their products. Globalization allows such companies to locate their facilities all over the globe, in search of local infrastructure and resources that minimize production costs. In this context, the use of multiple plants in different locations allows them to meet demands but raises the problem of determining which facilities should attend the given demand. This case-oriented research presents a lot sizing model for an aluminum beverage packaging company with several plants, a diversified portfolio, internal holding, and the possibility to rent additional warehouses. A particular characteristic of this case study is the high setup costs and times of the machines. For this reason, the mathematical formulation considers the setup carryover, setup crossover, besides transfers between plants. Preliminary results suggest that a commercial solver is able to solve real-based instances in an acceptable time.

Introduction

The lot sizing problem (LSP) is the object of several studies found in the literature [1], due to its wide applicability in different industrial segments. In general, the LSP consists of finding the production plan (item quantity in each period, when and where to produce the items, etc.) to meet the item demands minimizing related costs (maximizing profit). Moreover, several constraints that vary according to the application must be considered, such as inventory balance and setup of the machines. This study models an LSP considering the characteristics of a multinational company that produces metallic packaging.
In line with this, this paper presents a mathematical model for the multi-plant LSP (MPLSP) that considers setup carry-over. We also incorporate new features inherent in this case study: restriction in production lines, storage control in the plant, and rented warehouses. The objective function aims at minimizing the total costs of the aluminum beverage packing industry.

Problem Description

The case study of a metal packaging company is composed of multiple plants and the demands of the customers by plant are deterministic. The distribution of final items between plants is allowed. Therefore, the items are transferred from the plants to other plants or storage locations to meet the demands of the different plants.
The multiple plants have distinct production lines, which may have restrictions on the manufacturing of certain products from the company portfolio. The plants have limited internal storage environments but, if necessary, the plant can rent extra storage spaces, called external storage.
The production process involves production, setup, inventory (either to store inside the plant or in external warehouses) and transportation costs, for transferring items between plants. The costs are unitary and may vary according to the item, location, and period. The backlogging is also allowed, i.e, the company may not be able to meet the customer deadline so as the delivery is delayed, by considering an extra expense (backlogging cost).
The time available for each plant is limited and comprises the machine setup, the time used for manufacturing items, and the idle time. The setup times for the production of certain items are considerably high, up to 10 days, and there is a dependency between the setups not approached in this paper. There are setup time constraints for both the use of the machinery and for its preparation.
The mathematical model presented by Belo-Filho et al [2] was the closest formulation to the case study. To address long setup times, both setup carry-over and setup crossover are allowed. Figure 1 shows some possible setup situations that are

considered in this study.

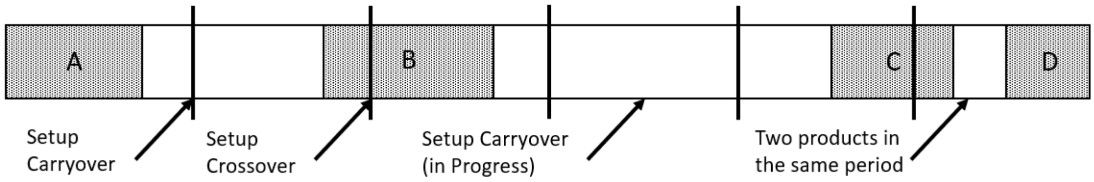

Figure 1: Various setup situations that may occur during the production process.

In summary, the MPLSP with multiple items and periods target of this study considers: internal storage, backlogging, setup crossover, and setup carryover. The model of [2] was adapted to include inter-plant transfers, external storage, production lines distinct in the plants, and the production restrictions between lines and products.

Computational Results

Using the data from the company as basis, were generated 15 instances with 12 periods on the planning horizon, each representing a month, 5 to 15 plants, 10 to 30 production lines and 2 to 10 items. The lines were randomly distributed among the factories, which could contain up to 5 production lines each. The parameters ranged between the observed intervals within the real data. Due to the tightness of the production capacities, each plant started with the setup of the first item in the planning horizon ready.
The results of preliminary tests are presented in Table 1. They were obtained using the solver Gurobi 9.1.1 with a time limit of 3600 seconds. It shows that Gurobi could solve small-sized instances very efficiently, but the instance size increases, the optimal solution is not guaranteed. However, we must point out that even for the largest instances, Gurobi achieved low gaps.

Acknowledge

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

Table 1: Results obtained using Gurobi 9.1.1

Instance Name	Number of Plants	Number of Lines	Number of Items	FO	Relative Gap	Time Elapsed
peq2	5	10	2	30114236.98	0	1.078
peq4	5	10	4	273237615.5	0	8.906
peq6	5	10	6	369895483.4	0	19.563
peq8	5	10	8	714872862.6	0	130.281
peq10	5	10	10	1369856593	0	228.172
med2	10	20	2	66428711.64	0	2.437
med4	10	20	4	184624112.8	0	285.89
med6	10	20	6	328197256.6	0	1929.734
med8	10	20	8	754572360.4	0	2522.422
med10	10	20	10	1005944950	0.032	3600.828
grd2	15	30	2	79803905.47	0	20.594
grd4	15	30	4	267030652.1	0.010	3602.468
grd6	15	30	6	647868820.6	0.028	3600.828
grd8	15	30	8	917459422.6	0.110	3601.125
grd10	15	30	10	1791259134	0.186	3600.125

References

[1] Behrooz Karimi, SMT Fatemi Ghomi, and J. M. Wilson. The capacitated lot sizing problem: a review of models and algorithms. Omega, v. 31, n. 5, p. 365-378, (2003).

[2] Márcio AF Belo-Filho, Franklina MB Toledo, and Bernardo Almada-Lobo. Models for capacitated lot-sizing problem with backlogging, setup carryover and crossover. Journal of the Operational Research Society, v. 65, n. 11, p. 1735-1747, (2014).

A Benders decomposition approach to the lot-sizing and market selection problem

Wilco van den Heuvel
Erasmus University Rotterdam
wvandenheuvel@ese.eur.nl

Semra Ağralı
MEF University
agralis@mef.edu.tr

Zeki Caner Taşkın
Boğaziçi University
caner.taskin@boun.edu.tr

Abstract

In the NP-hard lot-sizing and market selection problem one needs to select the markets to serve and construct a production plan to satisfy the demands of the markets selected. We consider single and multi-objective versions of the problem with trade-offs between revenues and costs, as well as market share and costs. We show how to apply a Benders decomposition approach in an effective way by utilizing an efficient procedure to get the dual variables of the Benders subproblems. Moreover, we establish a relationship between the Benders cuts and the core of a corresponding cooperative lot-sizing game, where in the latter game the lot-sizing costs need to be shared among players in a fair way. Finally, we propose how to reuse the Benders cuts to get an effective algorithm for the multi-objective version of the problem.

Introduction

The economic lot-sizing problem considers a production planning problem where the goal is to satisfy a set of deterministic dynamic demands over a planning horizon at minimum production cost. In this traditional problem, the set of demands to be satisfied are predetermined and there is no option for partial satisfaction by choosing some percentage of the demand. However, in today's competitive business environment, global companies make decisions not only on the supply side but also on the demand side [1]. Specifically, these companies select the set of markets whose demands they would like to satisfy instead of satisfying a predetermined deterministic demand. Hence, the problem has two decision stages: first, to determine the markets whose demand will be satisfied, and then, plan the lot sizing decisions to satisfy the whole demand. Since the selection of markets affect the lot-sizing decisions, these two types of decisions need to be made simultaneously.

[2] consider a production planning problem where a set of markets with known demands and corresponding revenues exists. The problem is to select the markets whose demands will be fully satisfied through a single product that has a lot-sizing production cost structure. Since the lot sizing decisions are affected by the total demand to be satisfied, it is important to select the markets wisely to satisfy their demands. They develop an integer programming formulation for the economic lot-sizing problem with market selection, which they refer to as market selection problem. Then, they prove that the problem is NP-Hard, introduce polynomially solvable special cases and a heuristic for the general case.

We consider a similar problem as [2] but within a multi-objective problem setting. Specifically, we consider a problem where the objective may not only consist of profit, but also entails revenue or market share. We first develop an exact solution method for the single objective version based on a decomposition algorithm, which can be used to effectively solve the multi-objective optimization model.

Problem Description and Mathematical Model

We study an integrated market selection and lot sizing problem over a planning horizon that consists of T planning periods. There is a set of markets, $m = 1, \dots, M$, each of which has a deterministic demand, d_t^m, for each period $t = 1, \dots, T$. If we decide to satisfy the demand of the market throughout the planning horizon, then we obtain a revenue R_m. There is a fixed setup cost, K_t, incurred if we decide to produce in that period and a variable production cost, p_t, per every unit we produce. We pay an inventory holding cost of h_t for every unit that we put in the inventory at the end of period t. Finally, C_{it} represents the unit cost of producing in period $i = 1, \dots, t$ and holding cost from period i to period $t = 1, \dots, T$, i.e., $C_{it} = p_i + \sum_{j=i}^{t} h_j$.

The first set of binary decision variables are the market selection variables, z_m, which take the value of 1 iff we decide to satisfy the demand of that market. Then, we have lot-sizing decision variables: y_t is a binary variable that takes the value of 1 iff we decide to produce in period $t = 1, \dots, T$, and x_{it}^m gives the amount of units produced in period $i = 1, \dots, t$ to satisfy the demand of market $m = 1, \dots, M$ in period $t = 1, \dots, T$. We can formulate our single-objective, maximization of the net profit, mathematical model as follows:

$$\text{(MIP)} \quad \max \quad \sum_{m=1}^{M} R_m z_m - \sum_{t=1}^{T} \left(K_t y_t + \sum_{i=1}^{t} \sum_{m=1}^{M} C_{it} x_{it}^m \right) \tag{1}$$

$$\text{s.t.} \quad \sum_{i=1}^{t} x_{it}^{m} = d_t^m z_m, \quad t = 1, \dots, T; m = 1, \dots, M \tag{2}$$

$$x_{it}^{m} \leq d_t^m y_i, \quad i \leq t = 1, \dots, T; m = 1, \dots, M \tag{3}$$

$$x_{it}^{m} \geq 0, \quad i \leq t = 1, \dots, T; m = 1, \dots, M \tag{4}$$

$$y_i, z_m \in \{0, 1\}, \quad i = 1, \dots, T; m = 1, \dots, M. \tag{5}$$

The objective function (1) maximizes the net profit that is obtained by subtracting the lot-sizing cost from the revenue obtained from the markets that are selected. Constraints (2) are demand satisfaction constraints for the markets that we selected throughout the planning horizon. Constraints (3) make sure that we can produce only if we pay the setup cost and the amount of production for period t does not exceed the demand at that period. Constraints (4) and (5) are the nonnegativity and binary restrictions for the decision variables, respectively.
In practice, one does not necessarily prefers to maximize profit, but market share may also be an important objective. With z_1 (resp. z_2) be the first (resp. second) objective, sensible versions of the multi-objective market selection problem are:
Revenue vs. cost: $(z_1, z_2) = \left(\sum_{m=1}^{M} R_m z_m, \sum_{t=1}^{T} \left(K_t y_t + \sum_{i=1}^{t} \sum_{m=1}^{M} C_{it} X_{it}^m \right) \right)$
Market share vs. cost: $(z_1, z_2) = \left(\sum_{m=1}^{M} \sum_{t=1}^{T} d_t^m z_m, \sum_{t=1}^{T} \left(K_t y_t + \sum_{i=1}^{t} \sum_{m=1}^{M} C_{it} X_{it}^m \right) \right)$
Furthermore, it will turn out that finding the maximum return on investment (ROI), i.e., maximizing the objective

$$\frac{\sum_{m=1}^{M} R_m z_m}{\sum_{t=1}^{T} \left(K_t y_t + \sum_{i=1}^{t} \sum_{m=1}^{M} C_{it} X_{it}^m \right)}$$

can be done by using the results from the "Revenue vs. cost" multi-objective problem.

Decomposition Algorithm

We propose a decomposition approach to solve the problem. The main idea is to decompose the problem into a master problem and a subproblem, where the information of optimal solutions of the subproblems will be used to generate cuts for the master problem. For our specific problem we propose a decomposition where the master problem mainly involves the markets to be selected, while the subproblem takes care of the lot-sizing cost. With this in mind, the master problem can be written as

$$\text{(MP)} \quad \max \sum_{m=1}^{M} R_m z_m - \nu(z) \tag{6}$$

$$\text{s.t.} \quad z_m \in \{0,1\}, \qquad m = 1, \ldots, M, \tag{7}$$

where $\nu(z)$ is the lot-sizing cost when markets given by $z = (z_1, \ldots, z_M)$ are selected. Suppose for the moment that it is possible to formulate the cost $\nu(z)$ in a linear way for a given market selection z, i.e.,

$$\nu(z) = \sum_{m=1}^{M} \beta_m^z z_m, \tag{8}$$

and that the linear expression is valid for any other market selection z', i.e., $\sum_{m=1}^{M} \beta_m^z z'_m \geq \nu(z')$. Then the master problem can be rewritten as

$$\text{(MP2)} \quad \max \quad \sum_{m=1}^{M} R_m z_m - \theta \tag{9}$$

$$\text{s.t.} \quad \theta \geq \sum_{m=1}^{M} \beta_m^z z_m \qquad \text{for each market selection } z \in \{0,1\}^M \tag{10}$$

$$z_m \in \{0,1\}, \qquad m = 1, \ldots, M. \tag{11}$$

The issue with this formulation is that (i) we need a procedure to find the β_m^z coefficients in (10), and (ii) there is an exponential number of constraints (10). To deal with issue (ii), instead of pre-computing all constraints (10) upfront, we generate them in a cutting plane fashion. This leads to the following algorithm.

Step 0 Initialization: start with set of cuts (possibly empty)

Step 1 Solve (MP2) with current set of cuts, resulting into an optimal solution $\tilde{z}$

Step 2 If $\theta < \nu(\tilde{z})$, generate cut of type (8) and add it to set of cuts, Go to Step 1

Step 3 Otherwise, output the optimal solution $\tilde{z}$.

We find the Benders cuts (issue (ii)) by using a relationship with the core of a corresponding cooperative lot-sizing game, where in the latter game the lot-sizing costs need to be shared among players in a fair way. The decomposition approach turns out to be competitive compared to using a commercial solver. Finally, we propose how to reuse the Benders cuts to get an effective algorithm for the multi-objective version of the problem.

References

[1] Geunes J, Levi R, Romeijn HE, Shmoys DB, Approximation algorithms for supply chain planning andlogistics problems with market choice, Mathematical programming 130(1):85–106, 2011.

[2] Van den Heuvel W, Kundakcioglu OE, Geunes J, Romeijn HE, Sharkey TC, Wagelmans AP, Integrated market selection and production planning: complexity and solution approaches, Mathematicalprogramming 134(2):395–424, 2012.

New valid inequalities for a multi-echelon multi-item lot-sizing problem with returns and lost sales

Franco Quezada
Sorbonne Université, LIP6, 75005 Paris, France
Universidad de Santiago de Chile, LDSPS, Santiago, Chile
franco.quezada@usach.cl

Céline Gicquel
Université Paris Saclay, LISN, 91190 Gif-sur-Yvette, France
celine.gicquel@lri.fr

Safia Kedad-Sidhoum
Conservatoire National des Arts et Métiers, CEDRIC, 75003 Paris, France
safia.kedad_sidhoum@cnam.fr

Abstract

This work studies a multi-echelon multi-item lot-sizing problem with remanufacturing and lost sales. The problem is formulated as a mixed-integer linear program. A new family of valid inequalities taking advantage of the problem structure is introduced and used in a customized branch-and-cut algorithm. The provided numerical results show that the proposed algorithm outperforms both the generic branch-and-cut algorithm embedded in a standard-alone mathematical solver and a previously published customized branch-and-cut algorithm.

Introduction

The present work considers a remanufacturing system involving three production echelons: disassembly of used products brought back by customers, refurbishing of the recovered parts and reassembly into like-new finished products. We aim at optimizing the production planning for the corresponding three-echelon system over a multi-period horizon. Within a remanufacturing context, production planning includes making decisions on the used products returned by customers, such as how much and when used products should be disassembled, refurbished or reassembled in order to build new or like-new products. The main objective is to meet customers' demand for the remanufactured products in the most cost-effective way.

Only a few works have addressed such multi-echelon production systems through exact solution approaches. A first attempt at tackling this difficulty can be found in [3]. Quezada et al. [3] considered the problem in a stochastic setting, taking into account uncertainties on the problem input parameters. They proposed a multi-stage stochastic approach based on the use of scenario trees. The problem was formulated as a MILP and solved through a new customized branch-and-cut algorithm. This algorithm relied on valid inequalities focused on strengthening the formulation of the single-echelon uncapacitated lot-sizing sub-problems embedded in the main problem. Although this approach was successful at providing near optimal solutions for small to medium size instances, some numerical difficulties were encountered to solve the larger instances. Intuitively, this difficulty might be partly due to the fact that the valid inequalities used to strengthen the formulation considered uncapacitated single-echelon sub-problems. They did not take into account the fact that, even if the production resources are assumed uncapacitated, the amount of products that can be processed on a resource at a given time period is limited among others by the amount of available used products returned up to this time period and by the yield of the disassembly process, i.e. by the proportion of disassembled parts that are in a sufficiently good state to be refurbished and reused in a remanufactured product. Hence, using valid inequalities taking into account this aspect of the problem might contribute in further strengthening its MILP formulation and decrease the computational effort needed to solve large-size instances.

To the best of our knowledge, the formulation of valid inequalities that explicitly take into account the impact of a limited returns quantity on the production plan has not yet been studied for a multi-echelon remanufacturing system. The present work aims at partially closing this gap by proposing new valid inequalities for this problem. The numerical results show the usefulness of the proposed inequalities at solving the problem under study.

Problem description and modeling

Production system

We consider a remanufacturing system comprising three main production echelons: disassembly, refurbishing and reassembly. We seek to plan the production activities in this system over a horizon comprising a discrete set $\mathcal{T} = \{1, .., T\}$ of periods. The system involves a set $\mathcal{I}$ of items. Among these ones, item $i = 0$ represents the used products returned by customers in limited quantities at each period. A used product is composed of I parts. Let α_i be the number of parts i embedded in a used

product. The returned products are first disassembled to obtain a set $\mathcal{I}_r = \{1, ..., I\}$ of recoverable parts. Due to the usage state of the used products, some of the parts obtained during disassembly have to be discarded. In order to reflect the variations in the quality of the used products, i.e. the yield of the disassembly process, we let π_i^t denote the proportion of parts which will be recoverable at each time period t for each item $i = \{1, \ldots, I\}$. The recoverable parts are then refurbished on dedicated refurbishing processes. The set of $\mathcal{I}_s = \{I+1, ..., 2I\}$ of serviceable parts obtained after refurbishing are reassembled into remanufactured products which have the same bill-of-material as the used products. These remanufactured products, indexed by $i = 2I + 1$, are used to satisfy the dynamic demand of customers.
The system comprises a set $\mathcal{P} = \{0, ..., I+1\}$ of production processes: $p = 0$ corresponds to the disassembly process, $p \in \{1, ..., I\}$ corresponds to the process refurbishing the recoverable part indexed by p into the serviceable part indexed by $p + I$ and $p = I + 1$ corresponds to the reassembly process. All these processes are assumed to be uncapacitated. All input parameters of the problem are time-dependent: r^t denotes the quantity of collected used products, d^t the customers' demand and π_i^t the proportion of recoverable parts $i \in \mathcal{I}_r$ obtained by disassembling one unit of returned product at period t. As for the costs, for each period t, we have the setup cost f_p^t for process $p \in \mathcal{P}$, the unit inventory cost h_i^t for part $i \in \mathcal{I}$, the unit lost-sales penalty cost l^t, the unit cost q_i^t for discarding item $i \in \mathcal{I}_r \cup \{0\}$ and the unit cost g^t for discarding the unrecoverable parts obtained while disassembling one unit of returned product.

Echelon stock reformulation

We now provide a formulation of the problem using the echelon stock concept [2]. The echelon stock of a product in a multi-echelon production system corresponds to the total quantity of the product held in inventory, either as such or as a component within its successors in the bill-of-material. In order to build a mathematical model for the problem, we introduce the following decision variables at time period $t \in \mathcal{T}$: we denote by E_i^t the echelon stock level of item $i \in \mathcal{I} \setminus \{0\}$ at the end of period t, X_p^t the quantity of parts processed on process $p \in \mathcal{P}$, $Y_p^t \in \{0, 1\}$ the setup variable for process $p \in \mathcal{P}$, S_i^t the inventory level of part $i \in \mathcal{I}$, Q_i^t the quantity of part $i \in \mathcal{I}_r \cup \{0\}$ discarded and L^t the lost sales of remanufactured products. This variables definition leads to the following MILP formulation:

$$\min \sum_{t \in \mathcal{T}} \Big(\sum_{p \in \mathcal{J}} f_p^t Y_p^t + h_i^t S_0^t + \sum_{i \in \mathcal{I} \setminus \{0\}} eh_i^t E_i^t + \sum_{i \in \mathcal{I}_r \cup \{0\}} q_i^t Q_i^t + g_0^t X_0^t \Big) \tag{12}$$

$$X_p^t \leq M_p^t Y_p^t \qquad \forall p \in \mathcal{J}, \forall t \in \mathcal{T} \tag{13}$$

$$S_0^t = S_0^{t-1} + r^t - X_0^t - Q_0^t \qquad \forall t \in \mathcal{T} \tag{14}$$

$$E_i^t = E_i^{t-1} + \pi_i^t \alpha_i X_0^t - \alpha_i(d^t - L^t) - Q_i^t \qquad \forall i \in \mathcal{I}_r, \forall t \in \mathcal{T} \qquad (15)$$
$$E_i^t = E_i^{t-1} + X_{i-I}^t - \alpha_i(d^t - L^t) \qquad \forall i \in I_s, \forall t \in \mathcal{T} \qquad (16)$$
$$E_{2I+1}^t = E_{2I+1}^{t-1} + X_{I+1}^t - d^t + L^t \qquad \forall t \in \mathcal{T} \qquad (17)$$
$$S_0^0 = 0 \qquad (18)$$
$$E_i^0 = 0 \qquad \forall i \in \mathcal{I} \setminus \{0\} \qquad (19)$$
$$E_i^t - E_{I+i}^t \geq 0 \qquad \forall i \in \mathcal{I}_r, \forall n \in \mathcal{T} \qquad (20)$$
$$E_i^t - \alpha_i E_{2I+1}^t \geq 0 \qquad \forall i \in \mathcal{I}_s, \forall n \in \mathcal{T} \qquad (21)$$
$$E_i^t \geq 0 \qquad \forall i \in \mathcal{I}, \forall t \in \mathcal{T} \qquad (22)$$
$$S_0^t, L^t \geq 0 \qquad \forall t \in \mathcal{T} \qquad (23)$$
$$X_p^t \geq 0, Y_p^t \in \{0,1\} \qquad \forall p \in \mathcal{J}, \forall t \in \mathcal{T} \qquad (24)$$

The objective function (12) aims at minimizing the total cost over the whole planning horizon. Constraints (13) link the production quantity variables to the setup variables. Constraints (14)-(17) are the inventory balance constraints. Constraints (14) use the classical inventory variables, whereas Constraints (15)-(17) make use of the echelon inventory variables. Constraints (18)-(19) translate the fact that the initial inventory of each item is assumed to be equal to 0. Constraints (20)-(21) ensure consistency between the echelon inventory at the different echelons of the bill-of-material and guarantee that the physical inventory of each product, S_i^t, remains non-negative for all $i \in \mathcal{I}$. Finally, Constraints (22)-(24) define the domain of the decision variables.

Single-echelon (ℓ, k, U) inequalities

We now seek to strengthen the single-echelon (k, U) inequalities investigated in [1] and [3] by considering the limited quantity of returned products available at each time period in the system. The (ℓ, k, U) inequalities are defined as follows:

Proposition 1. *Let $0 \leq \ell \leq k \leq T$ be two periods of the planning horizon. Let $U \subseteq \{k+1, ..., T\}$ be a subset of periods and $t^* = \max\{\tau : \tau \in U\}$ be the last time period belonging to U.*
The following inequalities are valid for Problem (13)-(24):

$$S_0^\ell \hat{\pi}_i^{\ell,t^*} + \alpha_i^{-1} E_i^k + \sum_{k<t\leq t^*} \phi_i^t Y_0^t \geq \sum_{t\in U} (d^t - L^t) \qquad \forall i \in \mathcal{I}_r \qquad (25)$$
$$S_0^\ell \hat{\pi}_i^{\ell,t^*} + \alpha_i^{-1}(E_i^\ell - E_{i+I}^\ell) + \alpha_i^{-1} E_{i+I}^k + \sum_{k<t\leq t^*} \phi_i^t Y_i^t \geq \sum_{t\in U} (d^t - L^t) \qquad \forall i \in \mathcal{I}_r \qquad (26)$$
$$S_0^\ell \hat{\pi}_i^{\ell,t^*} + (\alpha_i^{-1} E_i^\ell - E_{2I+I}^\ell) + E_{2I+1}^k + \sum_{k<t\leq t^*} \phi_i^t Y_{I+1}^t \geq \sum_{t\in U} (d^t - L^t) \qquad \forall i \in \mathcal{I}_r \qquad (27)$$

with $\phi_i^t = \min\left\{\sum_{\ell<\nu\le t} r^\nu \hat{\pi}_i^{\nu,t}, \sum_{\nu\in U: t\le\nu} d^\nu\right\}$

It is worth mentioning that the (k, U) inequalities used in [3] to strengthen the formulation (12)-(24) can be seen as a particular case of the more general family of (ℓ, k, U) inequalities (25)-(27) proposed in this work. Namely, by setting ℓ to 0 and by computing the value of ϕ^t without taking the returns into account (i.e. by setting ϕ^t to $\sum_{\nu\in U: t\le\nu} d^\nu$), each (ℓ, k, U) inequality (25)-(27) becomes a (k, U) inequality.
We will present numerical results that show the usefulness of the proposed inequalities at solving the problem under study.

References

[1] Marko Loparic, Yves Pochet, and Laurence A Wolsey. The uncapacitated lot-sizing problem with sales and safety stocks. *Mathematical Programming*, 89(3):487–504, 2001.

[2] Yves Pochet and Laurence A Wolsey. *Production planning by mixed integer programming*. Springer Science & Business Media, 2006.

[3] Franco Quezada, Céline Gicquel, Safia Kedad-Sidhoum, and Dong Quan Vu. A multi-stage stochastic integer programming approach for a multi-echelon lot-sizing problem with returns and lost sales. *Computers & Operations Research*, 116:104865, 2020.

Tackling Uncertainty in Lot Sizing – Comparison of Different Disciplines

Viktor Bindewald
Karlsruhe Institute of Technology (KIT), Institute for Operations Research
viktor.bindewald@kit.edu

Fabian Dunke
Karlsruhe Institute of Technology (KIT), Institute for Operations Research
fabian.dunke@kit.edu

Stefan Nickel
Karlsruhe Institute of Technology (KIT), Institute for Operations Research
stefan.nickel@kit.edu

Abstract

We study a new lot sizing variant where we assume that neither the planning horizon length nor the demands are known exactly. This setting is common in practice, since demand information is only accessible for a narrow time window and is revealed gradually over time. Moreover, problem parameters are often based on forecasts and hence are inherently uncertain. This results in a rolling horizon procedure, i.e., the multi-stage problem dissolves into a series of coupled snapshot problems with uncertain parameters. Depending on the available information, different approaches from online optimization, stochastic programming and robust optimization have to be selected to solve the snapshot problems. To evaluate the impact of the selected methodology on the solution quality we use a recent methodology-agnostic framework for multi-stage decision making. We discuss some computational results regarding different solution approaches and the value of available information here.

Introduction

Operational decision making in practice often suffers from uncertainty in the input data. This applies in particular when the uncertain information concerns future

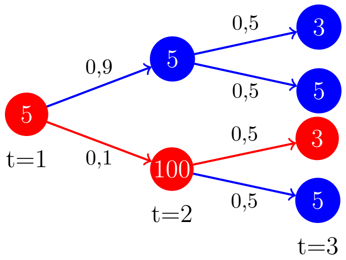

Variant	x_1	x_2	x_3	x_4	costs
OO	5	103	0	-	367
SP	10	98	0	-	372
RO	5	100	3	-	384
$OO_{\tau=1}$	5	103	0	8	411
$OO_{\tau=2}$	5	111	0	0	407

Figure 1: Illustrating differences between approaches for $T = 3$ and 4.

events and developments. In this case, decisions must be made at certain points in time, even though there is no chance of knowing how the future will unfold. Many problems exhibiting these features arise in settings which are driven by the release of customer orders, e.g., in production, supply chain management, and logistics. One prominent example is lot sizing. In its basic form, a lot sizing model seeks a production plan over T periods such that given customer demands $d_1, d_2, \ldots, d_T$ are met and the sum of production, setup and storage costs is minimized. The problem in practice, however, is that usually neither the planning horizon length T nor demand values are known precisely and decisions have to be made in a rolling horizon procedure. Existing research mainly focuses on extensions for deterministic models or on specific uncertainty types. Our approach is different insofar as we allow for an open planning horizon and different specifications of information about upcoming demands. The recently proposed methodology-agnostic framework for multi-stage decision making under uncertainty [2] has been designed exactly with the goal of deploying and quantifying the effects of different types of data uncertainty and related solution techniques. According to [1], such an inter-disciplinary comparison between different approaches is required by any comprehensive analysis of a specific problem setting involving uncertain data. Our research contributes a cross-methodological analysis of multi-stage lot sizing under different types of uncertainty and related solution approaches.

We present an example to illustrate our setting beginning with $T = 3$. The scenario tree is depicted in Figure 1. Each path from the root to a leaf represents a possible scenario where demand realization probabilities are displayed as edge labels. Production costs are 3, setup costs 20, storage costs 1, shortage costs 100 monetary units, and we start with an empty stock. The red path represents the realized scenario. However, only information about the subsequent stage is present at each decision moment in the rolling horizon procedure. The production plans obtained by solving the snapshot problems through online optimization (OO), stochastic programming (SP), and robust optimization (RO) are summarized in the table. Here, OO with its access to the deterministic lookahead performs best; SP is second, due to the distraction by the more probable outcome in $t = 2$; RO finishes last as it produces in every stage neglecting available information. To show the effect of increasing the available amount of information, consider an additional stage $t = 4$ with a realized demand of 8 in the OO setting. Then, if we increase the lookahead size from $\tau = 1$ to $\tau = 2$, the changed production plan allows to save 4 units by not producing in $t = 4$.

Multi-stage lot sizing under demand uncertainty

In the framework for multi-stage decision making under uncertainty [2], parameters are distinguished in two groups. The static data D_0 consists of the the production costs c^p, setup costs c^f, storage costs c^h, shortage costs c^s and the initial stock s_0. The time-dynamic data D_t consists of the current stock level s_{t-1} and the latest demand d_t. Furthermore some deterministic ($\mathbb{L}_t$) and uncertain ($\mathbb{U}_t$) information is available about the future demands d_{t+1}, d_{t+2} etc. The production decision x_t to be made and implemented in each stage results from solving a snapshot problem based on D_0 and D_t. The overall procedure then consists of a series of snapshot problems to be solved and decisions to be implemented. Each time period t a new demand d_t becomes known, the interaction chain is started with the updated data (see Figure 2). Which methodology is used to model and solve the snapshot problem is prescribed by $\texttt{paradigm}_t$. If online optimization is requested, then the snapshot model is a deterministic lotsizing problem with $1 + |\mathbb{L}_t|$ time periods and we use the standard MIP model. In case of stochastic programming or robust optimization the deterministic equivalent formulation of the usual SP or RO models are used for the snapshot problems. The result in the form of a tentative production plan is reported to the decision maker and made available to evaluation. The framework provides the user with freedom concerning demand representation and the related choices of snapshot models and algorithms. For example, hybrid settings are also possible, i.e., combining lookahead for the near future with uncertainty sets for the distant future.

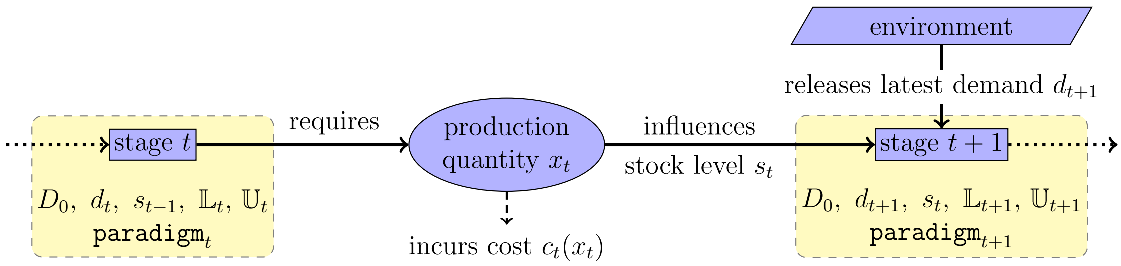

Figure 2: Illustration of the transition process from stage t to stage $t + 1$.

Computational Results

The algorithms are implemented in C++ and tested on a Ubuntu Linux machine. MIPs are solved using Cplex 20.1. For each setting, 100 instances are generated randomly with $T = 100$ and demands drawn uniformly from $[10, 100]$. Here, only the storage costs c^h vary, other cost parameters are set to $c^p = 0, c^f = 15, c^s = 5$. For RO and SP, two possible outcomes are known for each upcoming time period. We compare the settings with respect to the average costs accumulated over T.

Value of lookahead in OO and forecasts in RO and SP The value of information is highly dependent on the storage costs c^h and decreases fast with increasing $|\mathbb{L}_t|$ and $|\mathbb{U}_t|$ respectively. This is also true for the utopic setting $c^h = 0$ (depicted in gray). In summary you can see that RO performs slightly worse that SP and OO is always ahead because of its deterministic information.

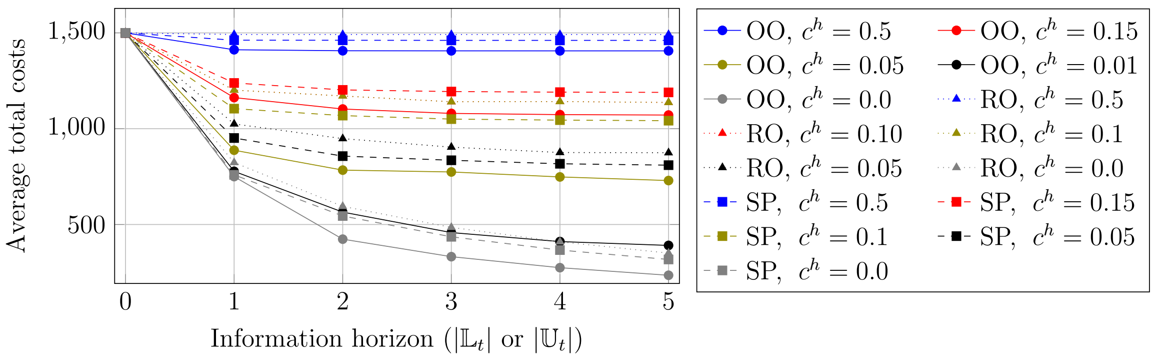

Inter-disciplinary comparison We compare hybrid approaches with $\mathbb{U}_t = 1, 2, 3$ combined with a fixed-size $\mathbb{L}_t$. We show the case $c^h = 0.1$ only and use OO as the baseline (depicted in black). As expected, SP is always behind OO and is usually followed by RO. Surprisingly, RO with $|\mathbb{U}_t| = 2$ and $|\mathbb{L}_t| = 2$ outperforms its SP counterpart. Note that in the more distant future ($|\mathbb{U}_t| + |\mathbb{L}_t| \geq 4$) however, the differences between the disciplines are almost completely eradicated.

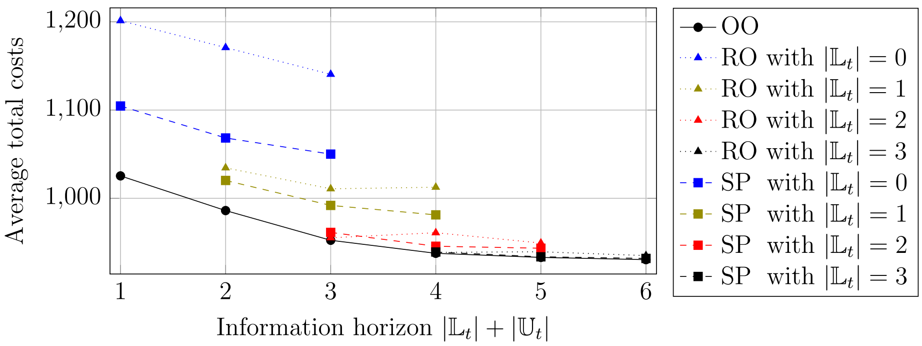

References

[1] Bakker H., Dunke F., Nickel, S., A structuring review on multi-stage optimization under uncertainty: Aligning concepts from theory and practice, Omega, 96, (2019)

[2] Bindewald V., Dunke F., Nickel S., Modeling multi-stage decision making under incomplete and uncertain information, www.optimization-online.org, (2020)

A Deep Reinforcement Learning approach for the Stochastic Inventory Problem

Henri Dehaybe, Daniele Catanzaro, Philippe Chevalier
Center for Operations Research and Econometrics (CORE), Université catholique de Louvain, Voie du Roman Pays 34, B-1348, Louvain-la-Neuve, Belgium
{*henri.dehaybe, daniele.catanzaro, philippe.chevalier*}*@uclouvain.be*

Abstract

We investigate the problem of approximating the optimal policies to backlog, lost sales and delivery lead times versions of the Single Location Stochastic Inventory Problem, by using Deep Reinforcement Learning. The neural network at the core of the heuristic applies to the general continuous inventory level problem and enables the approximation of optimal policies of unseen instances of the problem after an appropriate tuning and training. The proposed solution approach is novel in the literature and constitutes an example of the benefit that a cooperation between optimization and machine learning can bring not only to the respective research areas but also to practical problem solving.

Introduction

The more and more profound integration between *Neural Network* (NN) approximators and *Reinforcement Learning* (RL) methods in recent years boosted the development of *Deep Reinforcement Learning* (DRL), a machine learning framework able to model and solve *Markov Decision Processes* (MDPs). The Q-Networks, AlphaGo, AlphaZero [3, 6], and more recently, MuZero [5], are all examples of DRL methods able to surpass human experts in complex video games such as Chess, Go or Shogi. DRL has also enabled robots to learn how to perform complex tasks such as walking, running, jumping, or even dancing [2, 1]. The stochastic and time-delayed consequences involved in these examples can be tackled by DRL without requiring any particular assumption on the system under study. The underlying NN allows to model high-dimensional action-spaces as well as to approximate arbitrarily complex functions or to learn stochastic dynamics, when provided with an appropriate training environment characterized by the same behavior as the considered system. We aim to demonstrate that DRL is a promising way to tackle and solve stochastic dynamic Operations Management problems without having to concede on simplifying assumptions. To this end, we focus on finding close to optimal order policies

to different versions of the Single Location Stochastic Inventory Problem (SL-SIP) by adapting a DRL technique. An intrinsic advantage of our approach is to be able to find solutions under different assumptions about the system of interest such as whether the unsatisfied demand is lost or backlogged. We achieve this task by representing complex policies with high-dimensional state spaces in terms of specific Neural Networks trained with Deep Reinforcement Learning.
A comparison of the performance of the proposed approach on the non-stationary single-location inventory problem of [4] with various lead times and its lost sales version to baseline (s_t, S_t) policies shows that the trained NN can indeed interpolate close to optimal policies to previously unseen instances. We also show via an ablation study that both the discrete-continuous hybridization and the linear annealing components that we propose contribute greatly to improve the stability and the performance of the learning task.

Notation and formulations of the SL-SIP

The SL-SIP as described by [4] is a single-location single-item finite horizon non-stationary dynamic lot sizing problem under demand uncertainty. For each period $t \in 1, \ldots, T$, we have an uncertain Gaussian demand distribution $d_t \sim N(\mu_t, \sigma_t^2)$. Because the demand is naturally non-negative, we make the assumption that the coefficient of variability $CV = \sigma_t/\mu_t$ is low enough to make negative demand negligible. The process starts with an initial inventory level z_1. This model assumes non-stationary demand distributions, a fixed order cost (K), linear variable order costs (c), linear holding (h) and stock-out (b).
At the beginning of each period, one must take the decision to raise the inventory level by ordering a quantity q_t. The inventory reaches a level $y_t = z_t + q_t$, available to satisfy the demand of period t. After production, the demand realizes according to its distribution and the inventory level reaches $z_{t+1} = y_t - d_t$.
In the context of Reinforcement Learning, the Markov Decision Process to be solved is called an *environment* and is actually implemented as a simulation of the sequential decision problem we aim to solve. A RL *agent* can observe the state of the process and interact with it by exerting an action that causes the transition of the environment to a new state and the return of a reward. We implemented a Reinforcement Learning environment that simulates the SL-SIP where the action vector is $a_t = (q_t)$, the state vector $x_t = (\mu_t, \mu_{t+1}, \ldots, \mu_{t+H-1}, z_t)$ where H is a forecast horizon, and the rewards are the negative costs incurred at each period. During training, all the elements of the state are randomized at the first period so that the learning agent cannot overfit an instance by memorizing the best quantities to order at each period. While Scarf's results have been generalized to other cost functions [**?**, e.g.]]porteus71, we only consider the case of a Gaussian demand error

and linear ordering, holding and stock-out costs. Any different assumption would simply involve changing the reward signal or the transition dynamics of the MDP accordingly.

A Deep Reinforcement Learning Algorithm for the SL-SIP

We propose a tailored hybrid version of the Deep Deterministic Policy Gradient algorithm (DDPG) for solving the SL-SIP. DDPG approximates two functions with dedicated neural networks. The first (called the *critic*) is an action-value function $Q(x_t, a_t)$ that predicts the discounted present value of taking action a_t when the system is in state x_t. The second (called the *actor*) is a continuous policy function $\pi(x_t)$ that predicts the action that maximizes $Q(x_t, a_t)$. During training, the action has a probability of being randomized so as to increase the exploration of the state-action space. We propose two main adaptations to the vanilla DDPG to fit the SL-SIP task.

Firstly, we add a discrete component to DDPG to leverage the knowledge that ordering 0 is always a local optimum. We use a novel two-step hybrid continuous-discrete variant of DDPG. Specifically, a standard deterministic actor neural network is trained to approximate a continuous policy $\pi_A(x_t)$. During interactions with the environment, the hybrid agent policy is to choose

$$a_t = \pi(x_t) := \underset{a}{\operatorname{argmax}} \left[Q_\pi(x_t, a) : a \in \{(0), \pi_A(x_t)\} \right].$$

That is, the agent chooses between the action proposed by the continuous policy and the local optimum $a = (0)$, according to the most valued critic estimate.

Secondly, a high fixed order cost K alone makes the learning the continuous policy hard. To help the agent find a good policy in its first learning iterations, we add a linear annealing to gradually increase its value during training, starting from a low value until it reaches the target.

Despite all the improvements we propose, it can still happen that the algorithm converge to a *never-order* policy or diverges to excessive order quantities due to explosive gradient updates. A common practice in ML to overcome such problems is *model ensembling* (or *bagging*). It consists in training several models and have them vote on the best output, effectively forming a super-model. We bag N independently trained agents as follows. Given a state s, each agent proposes an action a_n. Then each agent estimates the action-value of each proposed action. The ensemble model returns the action with the highest median action-value. Of course, the drawback is that ensembling necessitates the training of N agents and thus multiplies the time needed to obtain a model.

Experiments

We test the ability of our agent to generalize their knowledge to demand forecasts which follow various realistic trend patterns that they never encountered during training. To do so, the initial inventory of the training environment and each demand forecast are randomly drawn from uniform distributions with large supports. The forecasts horizon, and thus the optimization horizon, is 52 periods.
We perform three experiments, one for each version of the SL-SIP problem: the standard backlog version; with different lead times; and the lost sales version. Twenty agents are trained per environment configuration (i.e. with different cost parameters) and tested individually. The performance of an agent and of the baseline policies are evaluated on each instance of the test dataset as the average return over 1000 Monte Carlo policy evaluations. We effectively test our approach on 18000 unique SL-SIP instances, all unknown to the agents.
The baseline policies against which we compare our agents are the discrete (s_t, S_t) policies computed with a precision of 0.1. The median gaps of the agents are 2.4% on the backlog environments; 1.9%, 1.7%, 2.2%, 4.9% on environments with lead times of 1, 2, 4, and 8 respectively; and 1.6% on lost sales environments (where the baseline is heuristic). Model ensembling further reduces these gaps as well as their variability with magnitudes dependent on the number of constituting models.

References

[1] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1861–1870, Stockholm, Sweden, 2018.

[2] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. In *Proceedings of the 4th International Conference on Learning Representations*, San Juan, Puerto Rico, 2016.

[3] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.

[4] H. E. Scarf. The optimality of (S, s) policies in the dynamic inventory problem. Technical report, Applied Mathematics and Statistics Laboratory, Stanford University, 1960.

[5] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.

[6] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.

Solving Multi-Echelon Inventory Management Problems with Deep Reinforcement Learning

Qihua Zhong
HEC Montréal
qihua.zhong@hec.ca

Yossiri Adulyasak
HEC Montréal
yossiri.adulyasak@hec.ca

Martin Cousineau
HEC Montréal
martin.cousineau@hec.ca

Raf Jans
HEC Montréal
raf.jans@hec.ca

Abstract

We consider a multi-echelon inventory management problem that aims to minimize the system-wide total backorder costs and inventory holding costs. When backorder costs are incurred at more than one stage, the optimal policy is unknown even in a simple serial system. We apply and compare three state-of-the-art deep reinforcement learning (DRL) algorithms including Dueling Double Deep Q-network, Advantage Actor-Critic and Twin Delayed Deep Deterministic Policy Gradient. We also propose a mechanism, Heuristic-Guided Exploration, to improve the training efficiency by incorporating known heuristics into the exploration process of the DRL algorithms.
Key words: machine learning, deep reinforcement learning, inventory management, multi-echelon

Introduction

In this paper, we consider a multi-echelon inventory management problem based on the beer game introduced in [1]. The objective of the beer game is to minimize the total holding cost and backorder cost of all four facilities in a serial supply chain. The optimal ordering policy is unknown except for the special case where the backorder cost only happens at the retailer stage.

Reinforcement learning (RL) is a computational framework that can be applied to solve Markov decision processes (MDPs) and deep reinforcement learning (DRL) combines it with deep learning using deep neural networks as powerful function approximators. Oroojlooyjadid et al. [2] applied the deep Q-networks (DQNs) algorithms [3] to the beer game to train DRL agents and showed that they have comparable or better performance in different scenarios when compared to the benchmark heuristic. However, this previous study only considers the discrete action space with a small limited number of discretized action values. This limits the size of the problems it can solve and also the granularity of decisions. Furthermore, due to the curse of dimensionality, the agent can only make decisions for one facility with other facilities following predefined policies instead of making coordinated decisions for all four facilities together. This motivates us to implement the Twin Delayed Deep Deterministic Policy Gradient (TD3), a state-of-the-art DRL algorithm that has a continuous action space and can train an agent to make centralized decisions for all four facilities at the same time. Two other state-of-the-art algorithms with a discrete action space (with more details in Section) are also implemented and compared.

Problem Description

The problem considers four subsequent facilities in a serial supply chain. The retailer observes independent stochastic demand and each facility places periodic orders to its upstream supplier. An external supplier is assumed to have unlimited capacity and always ships immediately after it receives orders. We neither have control over this external supplier, nor do we observe the state or the cost of it.

There is a positive information lead time, as well as a positive product shipment lead time. At each of the facilities, inventory holding costs are incurred for every period when there is positive inventory at the facility. Backorder costs are incurred when there are unfulfilled open orders at the end of the period. The objective is to minimize the total holding costs and backorder costs of all four facilities. Depending on whether the decisions are made locally for one facility at a time or centrally for all four facilities at the same time, we consider two versions of the problem: *decentralized single-facility* and *centralized multi-facility.*

In the *decentralized single-facility* setting, the agent only observes its own state and needs to make ordering decisions for only one facility with locally observed information, while the other three facilities follow predefined ordering policies. In the *centralized multi-facility* setting, the agent can observe the states of all four facilities and needs to make ordering decisions for all four facilities at the same time with the complete state information. The objective is the same for the two settings: the four facilities take collaborative decisions to minimize the cumulative overall supply chain cost over the K time periods. To the best of our knowledge, the *centralized multi-facility* setting has not yet been experimented with in the past using DRL algorithms.

Deep Reinforcement Learning Algorithms and Numerical Experiments

DRL algorithms train agents to learn a policy that maps the states to the corresponding probability distributions over the actions. Thanks to the capability of deep neural networks, DRL algorithms are able to accommodate high-dimensional state vectors. We let the state vector contain relevant information observable by the facility including its on-hand inventory, backorder quantities, the latest demand and its purchase order pipeline.

We test scenarios with different cost structures and demand pattern, one with a known optimal policy and the other without. This allows us to investigate two questions. First, for the problems without known optimal policies, can the DRL agents find better policies compared to the benchmark approach? Second, for the problems with known optimal policies, how closely can the DRL agents perform when compared with these optimal policies?

We apply three state-of-the-art DRL algorithms to the beer game. The first algorithm Duelling Double DQN (DDDQN) [4], an improved version of DQN, is tested and the results are compared to the results of base-stock policies with two-moment approximation target stock levels, as well as to the results of the vanilla DQN reported by Oroojlooyjadid et al. [2]. Applying the popular on-policy learning algorithm Advantage Actor-Critic (A2C) [5] allows us to investigate the learning efficiency of off-policy learning versus on-policy learning in the context of inventory management. Finally, implementing the TD3 algorithm allows us to investigate the potential advantage of having a continuous action space. While the decentralized single-facility beer game is solved by all three algorithms, the centralized multi-facility beer game is solved by TD3 only as DDDQN and A2C cannot efficiently solve problems of such a large action space. Table 1 summarizes these three algorithms and their characteristics.

In the *decentralized single-facility* setting, preliminary results show that both DDDQN

Algorithms	On- /off-policy	Action space	Policy
DDDQN	Off-policy	Discrete	Deterministic
A2C	On-policy	Discrete	Stochastic
TD3	Off-policy	Continuous	Deterministic

Table 1: DRL algorithms applied in this study and their characteristics.

and TD3 agents outperform the benchmarks in scenarios for which the optimal policies are unknown, while TD3 agents have the best performance in most scenarios. In the *centralized multi-facility* setting, to improve the learning efficiency when searching in the high-dimensional action space, we propose an exploration scheme called Heuristic-Guided Exploration (HGE) that incorporates known well-performing heuristics into the exploration process of DRL algorithms. Instead of starting from scratch, HGE allows that the agent, with a decaying probability μ^2, follows a heuristic to generate meaningful experiences at the beginning of the training as guided exploration. This is analogous to using the known heuristic as a constructive heuristic, combined with using the DRL as an improvement heuristic. Preliminary results show HGE improves learning efficiency. When the optimal policy is unknown, TD3 agents trained in *centralized multi-facility* setting outperform all other tested approaches.

References

[1] Sterman, J. D., Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment, Management Science, 35, 321–339 (1989)

[2] Oroojlooyjadid, A., MohammadReza, N., Lawrence, V. S., and Martin T., A Deep Q-Network for the Beer Game: Deep Reinforcement Learning for Inventory Optimization, Manufacturing & Service Operations Management, (2021)

[3] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... and Hassabis, D., Human-level control through deep reinforcement learning, Nature, 518, 529–533 (2015)

[4] Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., and Freitas, Dueling network architectures for deep reinforcement learning, International Conference on Machine Learning, 1995–2003 (2016)

[5] Wu, Y., Mansimov, E., Grosse, R. B., Liao, S., and Ba, J., Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. Advances in Neural Information Processing Systems, 30, 5279-5288 (2017)

[6] Fujimoto, S., Hoof, H., and Meger, D. , Addressing function approximation error in actor-critic methods, International Conference on Machine Learning, 1587–1596 (2018)

Solving profit-maximizing lot size problems with customer stockpiling and discrete prices

Erik Langelo
Faculty of Logistics, Molde University College, Norway
erik.langelo@himolde.no

Lars Magnus Hvattum
Faculty of Logistics, Molde University College, Norway
lars.m.hvattum@himolde.no

Abstract

This work concerns a discrete-time dynamic lot size model with dynamic pricing and pricing lag effects. This effect corresponds to the tendency of customers to stockpile a product when it is offered for cheap, causing demand to increase in the time period with a lower price, while it decreases in subsequent periods. Having non-zero production incurs a setup cost.

The problem consists of deciding, for each time period, whether one should produce, how much to produce, and which price to set. The objective is to maximize total profit, which is defined as total revenue minus total setup cost, inventory holding cost and production cost.

The problem is time-consuming to solve using a general solver. Making the assumption that prices can be chosen from a discrete list, a dynamic programming algorithm is presented, which solves the problem with discrete prices to optimality. This solution method can be used as a heuristic for the continuous problem.

Introduction

We consider a decision maker selling a product, whose goal is to decide prices throughout a discrete planning horizon in order to maximize his or her total profit, given as total revenue minus total production cost (and for model 2, we also include total setup cost and inventory holding cost). Prices may be chosen from a predetermined list of options. These price options may be different in each time period, and the number of options in each time period may vary. Demand is a function of the price in the same period, as well as in previous periods.

Problem variations

Model 1 - pure pricing, one lag period

We consider three versions of the above presented problem. Model 1 is the simplest model, in which one must decide on a price in each period, in order to maximize total profit, given as total revenue minus total production cost. Demand is a function of the price in the current period and the price in the previous period, and production is assumed to occur in every period.

$$\begin{aligned}
\Pi &= \text{total profit} \\
d_t &= \text{demand in period } t \\
p_t &= \text{price in period } t \\
c_t &= \text{marginal production cost in period } t \\
P_t &= \text{set of allowed prices in period } t
\end{aligned}$$

$$\max \Pi = \sum_{t=1}^{T} d_t(p_t - c_t) \tag{28}$$

subject to

$$\begin{aligned}
d_1 &= f_1(p_1) && (29) \\
d_t &= f_t(p_{t-1}, p_t),\ t = 2 \dots T && (30) \\
p_t &\in P_t,\ t = 1 \dots T && (31)
\end{aligned}$$

The objective function (28) represents the sum of revenue minus production costs for all time periods. Constraints (29)-(30) define demand in period 1 as a function of its price, and demand in other periods as a function of prices in the same period and the previous period. Constraint (31) enforces that prices are chosen from the predetermined set of allowed prices.
We model this problem as a longest path problem on a layered network, see figure 1. Nodes are divided into T disjoint sets, not including the source and sink nodes. Each node represents a pricing decision, and travelling along an arc from one node to another rewards a nonnegative profit. Travelling to the sink node gives zero profit. The objective is to find the longest path from the source node o to the sink node s. Such a network problem has already been treated in the literature. See [1] for a backwards dynamic programming algorithm that solves a similar shortest path problem to optimality in $O(|A|)$ time, where A is the set of arcs. We present a similar, forwards algorithm which also finds the optimal solution in $O(|A|)$ time.

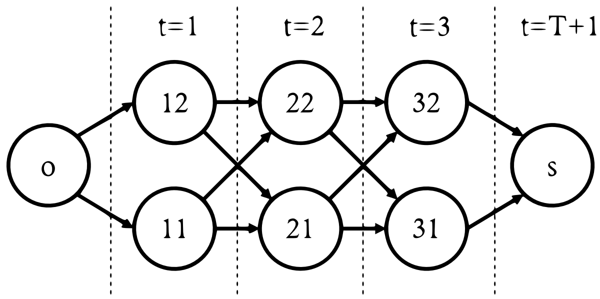

Figure 1: Graph representation of model 1.

Model 2 - pricing and setup, one lag period

In model 2, we add setup decisions to model 1. Production only occurs in periods in which setup occurs, and having setup incurs a given cost, regardless of amount produced. Product may be held in inventory between periods, also at a cost.

$$\begin{aligned}
x_t &= \text{units produced in period } t \\
h_t &= \text{cost of holding one unit of inventory from period } t \text{ to the next} \\
I_t &= \text{units held in inventory from period } t \text{ to the next} \\
s_t &= \text{setup cost in period } t \\
y_t &= \text{equals 1 if setup occurs in period } t\text{; zero otherwise} \\
i_0 &= \text{initial inventory} \\
M &= \text{a big number}
\end{aligned}$$

$$\max \Pi = \sum_{t=1}^{T} (d_t p_t - c_t x_t - s_t y_t - h_t I_t) \tag{32}$$

subject to (29)-(31) and

$$x_t + I_{t-1} - I_t = d_t, \ t = 1 \ldots T \tag{33}$$
$$x_t \leq M y_t, \ t = 1 \ldots T \tag{34}$$
$$I_0 = i_0 \tag{35}$$
$$y_t \in \{0, 1\}, \ t = 1 \ldots T \tag{36}$$
$$x_t, I_t \geq 0, \ t = 1 \ldots T \tag{37}$$

The objective function (32) is the difference between total revenue and total production, setup and inventory holding cost. Constraint (33) is the inventory balance constraint, constraint (34) enforces no production without setup, and constraint (35)

sets initial inventory to its appropriate value. Finally, constraints (36)-(37) are the binary condition on the setup variable and non-negativity conditions on production and inventory.
Under the assumption of Wagner-Whitin costs, we present a dynamic programming algorithm to solve this problem to optimality in $O(T^2K^3)$ time, where $K = \max(|P_t|)$, the highest number of prices in any period.

Model 3 - pure pricing, multiple lag periods

Model 3 is an extension of model 1 in that demand is a function of prices in any number of preceding prices, not only one. The objective functions in the two models are the same, see (28), as well as the constraint ensuring that prices are discrete, see (31). Constraints (29)-(30) are replaced by, respectively,

$$d_t = f_t(p_1, p_2, \dots, p_{t-1}, p_t), \ t = 1 \dots q \tag{38}$$
$$d_t = f_t(p_{t-q}, p_{t-q+1}, \dots, p_{t-1}, p_t), \ t = q+1 \dots T \tag{39}$$

where q is the number of preceding periods, the prices of which influence demand in the current period.
Unfortunately, the number of combinations of preceding prices for the demand function is bounded by K^q, and even calculating every possible value of a demand function is exponential in complexity. This makes it difficult to avoid such complexity in an exact solution algorithm. However, by making assumptions on the nature of the demand function, it is possible to rule out some combinations that will not appear in an optimal solution. Without assumptions on the demand function, we present an algorithm that solves the problem to optimality in $O(TK^q)$ time.

Perspectives

Although these models were conceived with the intention to describe stockpiling behaviour in customers, we make no strict assumption on the demand function. Any demand function may be used, as long as it is a function of the price in the same period and the prices in some number of preceding periods. This allows for the use of elaborate, non-linear functions, as well as other types of functions and customer behaviours not necessarily related to stockpiling. One may also note that, while prices are assumed to be discrete, these models may be used heuristically on a problem with continuous prices, in order to find a good, but not necessarily optimal solution.

References

[1] Papadimitriou, C.H. and Steiglitz, K., Combinatorial Optimization: Algorithms and Complexity, 448-450. Courier Corporation (1998).

New construction heuristics for lot-sizing problems

Daryna Dziuba
European University Viadrina Frankfurt (Oder)
dziuba@europa-uni.de

Christian Almeder
European University Viadrina Frankfurt (Oder)
almeder@europa-uni.de

Abstract

We consider the classical capacitated lot-sizing problem which is known to be NP-hard. Several construction heuristics have been proposed in the research literature, but none of them is convincing in terms of solution quality and generality – meaning that they can be applied to different variations of the problem. We propose a general greedy construction heuristic (GCH). Computational experiments on the single-level capacitated lot-sizing problem (CLSP) show that GCH outperforms Dixon-Silver and ABC heuristics. The heuristic can easily be extended to handle cases with setup times and multi-level product structures.

Introduction

Capacitated lot-sizing problems belong to complex combinatorial problems and it is usually hard to solve them to optimality or find a good feasible solution by exact methods. In this regard, a few heuristics were developed over the past 40 years. The standard CLSP with setup times is formulated as a MIP-problem as follows:

$$\min Z = \sum_{t=1}^{T} \sum_{i=1}^{N} (s_i Y_{it} + h_i I_{it}) + \sum_{t=1}^{T} c^o O_t \tag{40}$$

subject to

$$I_{it} = I_{it-1} + X_{it} - d_{it} \qquad \forall i, t \tag{41}$$

$$\sum_{i=1}^{N} (X_{it} + b_i Y_{it}) \leq C_t + O_t \qquad \forall t \tag{42}$$

$$X_{it} \leq Y_{it} \sum_{\tau=1}^{T} d_{i\tau} \qquad \forall i, t \tag{43}$$

$$X_{it}, I_{it}, O_t \geq 0,\ Y_{it} \in \{0, 1\} \tag{44}$$

where the decision variables are lot sizes X_{it} for item i in period t, setups Y_{it}, inventory I_{it}, and overtime O_t. Objective function (40) minimizes the sum of setup, holding and overtime costs. It is subject to the inventory balance constraint (41), capacity constraint (42) and setup state constraint (43).
Two well-known construction heuristics for single-level CLSP without setup time are: Dixon-Silver-Heuristic [1] and ABC-Heuristic [3]. Recently a new extension of the Dixon-Silver-Heuristic has been proposed where the local decision criterion is optimized using genetic programming (GP) [2].
A simple heuristic was presented in [5] for the single-level CLSP with setup times. All mentioned methods are based on (modified) Silver-Meal criterion and create a production plan stepwise from the first to the last period. They are rather inflexible because they can hardly be adapted to other lot-sizing problems.
Recently, [4] proposed a simple construction heuristic embedded in a metaheuristic for the practical lot-sizing problem of a pharmaceutical company. This construction heuristic is a simple, rule-based method to add new demand to an existing production plan. The order in which demand is added to the plan is optimized by a genetic algorithm. We used that idea to develop a new greedy construction heuristics.

Greedy construction heuristic

In order to solve the CLSP we first sort all non-zero demands d_{it} in an arbitrary order, i.e. we generate a so-called demand list of demand elements: $D^l = [d_{i_1 t_1}, d_{i_2 t_2}, d_{i_3 t_3}, \ldots]$.
We start with an empty production plan and add one-by-one the demand elements to the plan. Let us assume we have given a partial production plan $P^{n-1} = \{X_{it}^{n-1}, Y_{it}^{n-1}, I_{it}^{n-1}, O_t^{n-1}\}$ up to the $(n-1)$-th demand element. In order to integrate the n-th demand element $d_{i_n t_n}$ into the plan we solve the problem (1) - (5) with the additional constraints:

$$X_{jt}^{n} = X_{jt}^{n-1} \qquad I_{jt}^{n} = I_{jt}^{n-1} \qquad Y_{jt}^{n} = Y_{jt}^{n-1} \qquad j \neq i \quad \forall j, t \tag{45}$$

Constraints (6) ensure that previously taken decisions of setup and production quantities for all items except item i cannot be changed and lead to an easier optimization

problem but it is not guaranteed that the plan P^n is optimal if plan P^{n-1} was optimal. To solve the problem (1) - (6) we consider the following rules in the given order.
1. **Case $\mathbf{I_{it}^{n-1}} > \mathbf{0}$** ***(use inventory)***. If there is a positive inventory I_{it}^{n-1} for the item i in the period t, it should be used to satisfy (completely or partially) the demand $d_{i_n t_n}$. Obviously, in the case of using already available inventory no additional costs will occur that leads to the optimal solution of the subproblem. However, the inventory is required to satisfy future demand for item i. Therefore, a new demand element $d_{i_n t_n+1}$ is created and scheduled immediately in order to ensure the feasibility of the plan. New demand $d_{i_n t_n+1}$ is equal to inventory I_{it}^{n-1} which can be used to satisfy $d_{i_n t_n}$ completely or partially. If $I_{it}^{n-1} < d_{i_n t_n}$, i.e. demand was only partially satisfied, then the remaining amount is considered as a new demand element $d_{i_n t_n}$, which is scheduled according to the following rules.
2. **Case** $C_t - \sum_{i=1}^{N} X_{it}^{n-1} \geq d_{i_n t_n}$ ***(everything can be added)***. Depending on available capacity in period t, it may be possible to add the whole amount $d_{i_n t_n}$. Such a decision would either result in no additional cost (if there is already a setup in period t for item i) or in additional setup cost. Another decision may be to extend some existing lot(s) in previous periods if this leads to a feasible solution. This decision would result in additional holding cost. Therefore, we compute the cost for both decisions and the cheapest one will be realized.
3. **Case** $C_t - \sum_{i=1}^{N} X_{it}^{n-1} < d_{i_n t_n}$ ***(part or nothing can be added)***. If available capacity in period t does not allow to produce whole demand, we have several possibilities:

1. Add possible amount to t and extend existing lot(s) in previous period(s) to add remaining amount.
2. Add possible amount to t and create a new lot in previous period to add remaining amount.
3. Extend existing lot(s) in previous period(s) to add the whole amount.
4. Create a new lot in previous period to add the whole amount.

Each of the above-mentioned decisions in all cases yields a certain cost.
The total cost of adding a new production amount to a period is fourfold. First, there may be an additional setup cost incurred if there is no lot for item i in the considered period. Second, there may be an additional holding cost if we add an amount to the earlier period(s). Additionally, we introduce the ***shift of existing lots***. This routine is performed while estimating cost for any decision that involves the creation of a new lot and later while realizing that decision. It may reduce the total cost of a production plan. We compute two types of possible savings: (1) we may merge existing lots from previous periods with a later lot and reduce the

holding cost part and (2) we may merge future lots with an earlier lot and reduce the setup cost part. Both shift types may also be beneficial in terms of freeing capacity for future decisions.
The greedy character of the heuristic suggests that we always select the cheapest option while adding a demand element to a partial production plan. However, in some cases difference among additional costs does not indicate a clear preference of a certain decision. Therefore, we ***postpone*** adding such a demand element until there is a change in the plan for the considered item.

Computational experiments

The final plan depends on the order of the demand elements. We have tested various sorting rules. A period-wise sorting with items sorted according to a combination of capacity requirement, setup cost and holding cost leads to the best overall results. We were able to find feasible solutions for all tested instances of CLSP without setup time and for 95% of instances of CLSP with setup times while using the selected rules. There is a potential to find feasible solutions for those instances if we use local search to find a better sorting of the demand elements.
Results of the experiments on 1471 instances of CLSP without setup times report the average gap to MIP-solutions as 2.21%, whereas ABC-Heuristic leads to 3.71% average gap and Dixon-Silver-Heuristic to 4.77%. GCH performance is comparable to the best result of GP method (in both cases the reported average gap is 2.82% on 540 instances from the validation set).
Results of the preliminary experiments conducted on a sample of instances of CLSP with setup times report the average gap to MIP-solution as 6.23% for the instances where a feasible solution was found. [5] reports average gap of 6.46% to the best solution while finding a feasible solution for 95% of instances during preliminary experiments.

Conclusions & Outlook

The biggest advantage of GCH is its flexibility as the same design can handle both CLSP without and with setup times. The nature of the proposed heuristic allows us to use it as an offline as well as an online algorithm. The results of computational experiments showed improvement in solutions quality for CLSP without setup times and the ability to find a feasible good solution for most of the instances of CLSP with setup times. The current work focuses on finding out which problem characteristics determine the success of a sorting rule for a demand list and the optimal postponement strategy during the plan construction.

References

[1] Dixon, P.S., Silver, E.A., A heuristic solution procedure for the multi-item, single-level, limited capacity, lot-sizing problem, Journal of Operations Management, 2, 23-39 (1981)

[2] Hein, F., Almeder, C., Figueira, G., Almada-Lobo, B., Designing new heuristics for the capacitated lot sizing problem by genetic programming, Computers & Operations Research, 96, 1-14 (2018)

[3] Maes, J., VanWassenhove, L.N., A simple heuristic for the multi item single level capacitated lotsizing problem, Operations Research Letters, 4, 265-273 (1986)

[4] Oyebolu, F.B., van Lidth de Jeude, J., Siganporia, C., Farid, S.S., Allmendinger, R., Branke, J., A new lot sizing and scheduling heuristic for multi-site biopharmaceutical production, Journal of Heuristics, 23, 231-256 (2017)

[5] Trigeiro, W., A simple heuristic for lot sizing with setup times, Decision Sciences, vol. 20, no. 2, 294-303 (1989)

Managing Flexibility in Stochastic Multi-level Lot Sizing Problem with Service Level Constraints

Narges Sereshti
HEC Montréal
narges.sereshti@hec.ca

Raf Jans
HEC Montréal
raf.jans@hec.ca

Yossiri Adulyasak
HEC Montréal
yossiri.adulyasak@hec.ca

Abstract

In this research, we investigate the stochastic multi-level lot sizing problem with a service level. We consider a general setting in which it is possible to have independent demand for the components as well. The problem with uncertain demand is modeled as a two-stage stochastic program considering different demand scenarios. We first consider at all levels a static strategy in which both the setup decisions and the production quantities are determined in the first stage before the demand is realized. We also model a more adaptive strategy to be more responsive to the realized demand by considering the production quantities at some levels as recourse decisions. Through numerical experiments we investigate the value of applying such an adaptive strategy and adding more flexibility in the system under different settings.

Introduction

Being cost efficient is an important imperative in a competitive business environment. For manufacturing companies, having an efficient production plan in the context of material requirements planning (MRP) system is important to minimize different costs of production and inventory control. In MRP, time-phased production and inventory plans are crucial decisions to make a balance between customers'

demand satisfaction and cost management. While insufficient inventory will lead to shortage, unnecessary stocks will increase holding cost.
The standard lot sizing problem aims to determine the optimal timing and production quantities in order to satisfy a known demand over a finite and discrete time horizon. One of the extensions to the standard lot sizing problem is to consider the multi-level product structure which is common in MRP systems. While only the independent demand exists for each of the products in the single level lot sizing problem, there is also dependent demand due to the bill of material (BOM) structure in a multi-level lot sizing problem.
With a BOM structure, the possibility of backlog is only for the independent demand, which is usually defined for the end items. This is due to the fact that to produce the items at the lower levels, their components need to be available at the required time, and it is not possible to have backlog for the dependent demand [3]. In this research, we address a more general setting in which in addition to the end items, each of the components in the BOM may also have an independent demand, therefore it is possible to have backlog for them due to this independent portion.
This problem has practical relevance in industries with production and aftermarket services, which require spare parts. A good example is the aerospace industry where, in addition to the demand for end items, the components also have independent demand, which has to be taken into account in the planning process.
The demand is often unknown in real-world applications. In MRP systems with uncertain demand, the calculations are mostly based on the deterministic demand and the uncertainty is usually reflected in the amount of safety stock which will be added to the amount of net demand. In this research, we will use stochastic optimization models to deal with demand uncertainty. In these models the lot sizing and safety stock level decisions are jointly determined as the demand's probability distributions are considered in the model [4].
Different forms of service levels are widely used in the calculations of safety stocks to deal with demand uncertainty in stochastic lot-sizing problems. However, most of the research has been focused on the single level problem. As in the BOM it is not possible to have backlog for the dependent demand, the service levels are defined only with regard to the independent demand of the end products and the components. The service level which we consider in this research is closely related to the δ service level proposed by Helber et al. [2]. Here, instead of limiting average backlog, we limit the maximum proportion of total backlog to the total possible backlog over the whole planning horizon.
In the single level lot sizing problem, there are three main strategies to deal with multi-period lot sizing problems with stochastic demand and these have a different approach for the setup and production decisions [1]. In the static strategy, the setup and production decisions will be defined at the beginning of the planning horizon and they remain unchanged with the demand realization. In the dynamic strategy,

both setups and production decisions may be modified after the demand realization. The static-dynamic strategy is in between these two strategies in which the setups are fixed at the beginning of the planning horizon and the production decisions are made after the demand realization.
These strategies can also be applied in multi-level lot sizing problems. In the system in which we have independent demand for the components as well, we can apply different strategies at different levels in the BOM to increase the responsiveness in the system, while keeping the nervousness under control. By allowing (some) production decision to be made in the second stage, we create more nervousness [5], but we gain more flexibility and hence lower costs.
In this research, we model the stochastic multi-level lot sizing problem as a two-stage stochastic programming model which is solved using the sample average approximation (SAA) formulations. The contributions of this research can be stated as follows. First, we investigate the stochastic multi-level lot sizing problem with service level constraints. Second, we investigate the value of adding flexibility in production at different levels for different BOM structures, specially when the independent demand exists at the component level as well. Third, we apply the SAA method to empirically evaluate the solution quality with different number of scenarios.

Stochastic multi-level lot sizing problem

We propose the mathematical model for the stochastic multi-level lot sizing problem with a service level constraint. The model is an extension of the model proposed by Hung and Chien [3] in which for each type of product in the system there is a different set of inventory balance constraints. In this model, the structure of the BOM is considered by the successors of each item, which are the direct parents of the item. The capacity is also defined for each level of BOM separately.
In the first version of the problem, we assume that the strategy is static for all the items which implies that the setup and production quantity decisions are determined at the beginning of the planning horizon and cannot be changed when demands are realized [1]. Figure 1 shows the dynamics of decisions in this problem. The setup and production, and overtime variables are the first stage variables which will be defined before the demand realization. After the demand is realized for the entire planning horizon, the resulting inventory, and backlog quantities for each scenario are determined in the second stage [2]. In this problem, the model guarantees that, for each product with external demand, the worst-case ratio of backlog to maximum possible backlog among all scenarios satisfies the service level.
In the second version of the problem, we will apply a more adaptive strategy for some products which are managed by the static-dynamic strategy, while all the other products follow the static strategy. The dynamics of decision is depicted in Figure 2. In this case, the production decisions of the products under static-dynamic strategy

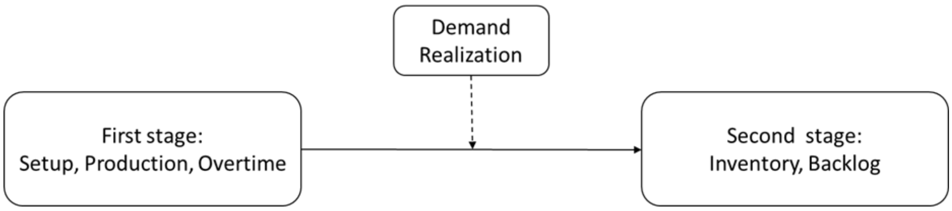

Figure 1: Sequence of events for the case without flexibility

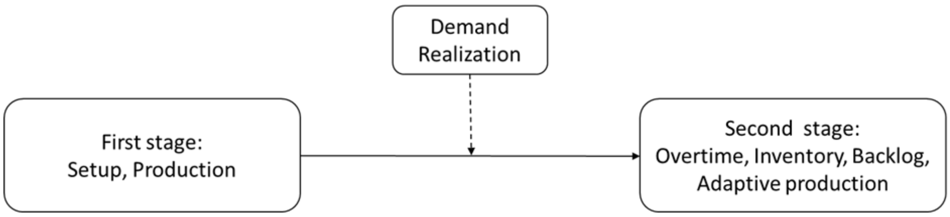

Figure 2: Sequence of events for the case with flexibility

become recourse decisions.
Both of these model are approximated using a finite number of scenarios and solved by the SAA method. Preliminary results for the serial structure show that increasing the flexibility results in cost reduction, even in the case where there is no external demand for any of the components. Sensitivity analyses have been performed to show the effect of changing different parameters on the cost reduction by adding flexibility and allowing more adaptive decisions.

References

[1] Bookbinder, James H and Tan, Jin-Yan, Strategies for the probabilistic lot-sizing problem with service-level constraints, Management Science, 34, 9, 1096–1108, (1988)

[2] Helber, Stefan and Sahling, Florian and Schimmelpfeng, Katja, Dynamic capacitated lot sizing with random demand and dynamic safety stocks, OR spectrum, 35, 1, 75–105, (2013)

[3] Hung, Yi-Feng and Chien, Kuo-Liang, A multi-class multi-level capacitated lot sizing model, Journal of the Operational Research Society, 51, 11, 1309–1318, (2000)

[4] Thevenin, Simon and Adulyasak, Yossiri and Cordeau, Jean-François, Material Requirements Planning Under Demand Uncertainty Using Stochastic Optimization, Production and Operations Management, 30, 2, 475–493 (2021)

[5] Tunc, Huseyin and Kilic, Onur A and Tarim, S Armagan and Eksioglu, Burak A simple approach for assessing the cost of system nervousness, International Journal of Production Economics, 141, 2, 619–625, (2013)

Heuristic approximation of the Pareto-front in multi-objective stochastic lot sizing

Fabian Friese
Department of Production Management, Leibniz University Hannover
fabian.friese@prod.uni-hannover.de

Stefan Helber
Department of Production Management, Leibniz University Hannover
stefan.helber@prod.uni-hannover.de

The multi-objective stochastic lot sizing problem (MO-SCLSP)

In lot sizing with stochastic demand, the minimization of operational costs is not the only conceivable objective. Minimizing the tardiness in customer demand satisfaction and ensuring production plan stability are no less important. Figure 1 shows conflicting relationships between those three objectives.

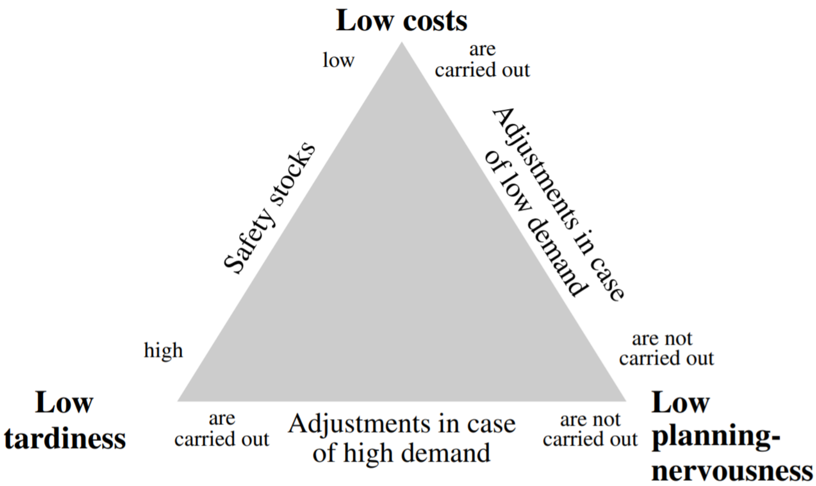

Figure 1: Triangle of tension of stochastic lot sizing

Safety stocks can be produced, which on the one hand buffer against unexpected high demand realizations. Their production and storage, on the other hand, can incur additional operational costs, e.g. setup and/or holding costs. A production plan determined based on uncertain demand forecasts can turn out to be deficient for a particular trajectory of demand realizations. In those cases, adjustments of the production plan are reasonable to adapt the production plan to the actual demand realizations and therefore reduce expected operational costs and/or tardiness. However, those adjustments induce system nervousness and therefore shall be limited. We consider those three objectives in a multi-objective formulation of the stochastic capacitated lot sizing problem (MO-SCLSP).

Solving the MO-SCLSP with an interactive approach in multiple decision stages

To allow production plan adjustments, the production plan is revisited periodically in multiple decision stages as shown in Figure 2.

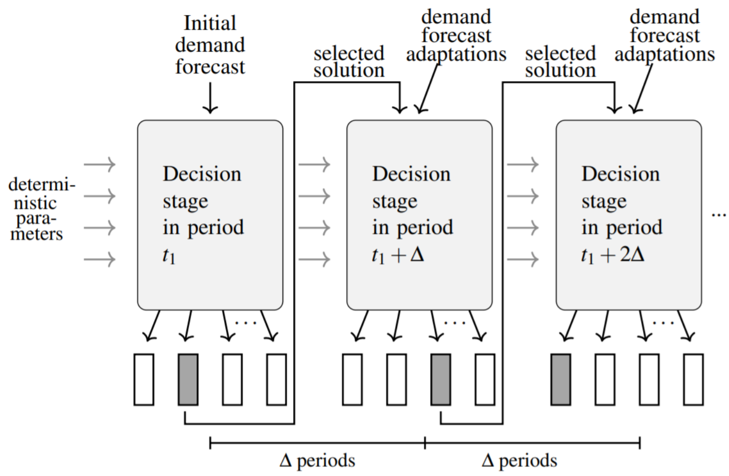

Figure 2: Solving the MO-SCLSP in multiple decision stages

The first decision is taken in period t_1 based on the initial demand forecasts. To allow the decision maker to make an informed decision, several Pareto-optimal solutions are required, from which the decision maker selects one based on their preferences. The selected production plan is executed until the next decision stage and serves as a parameter for the next optimization. After each Δ periods, the demand forecasts are updated to take into account newly observed demand information. Based on the updated demand forecasts the production plan can be adjusted, if the reduction in operational costs and tardiness justifies the induced system nervousness.

Figure 3 shows the order of events in each decision stage to determine a set of Pareto-optimal solutions. With three different scalarization techniques, the conceptionally multi-objective model is transferred into single-objective models. The

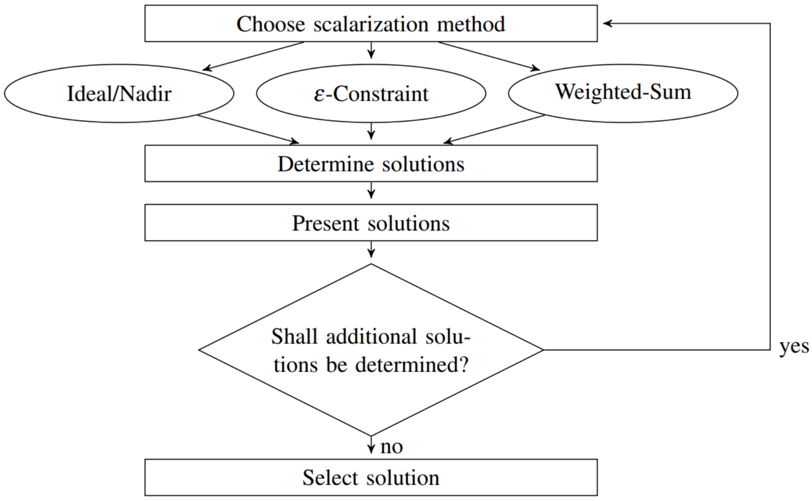

Figure 3: Order of events in each decision stage

optimal solutions of the scalarized single-objective models are also Pareto-optimal for the multi-objective model.
Ideal values are determined by optimizing just one objective function while neglecting the others. The optimal objective function values under the condition that another objective function obtains its ideal value is called nadir value. With the ε-Constraint method, each objective function but one is transferred into a constraint stipulation minimal aspiration levels of the corresponding objective function values. The remaining objective function is optimized and therefore the scalarized model is single-objective. The aspiration levels are varied systematically to derive a set of Pareto-optimal solutions. In the case of stochastic lot sizing, the minimization of tardiness and nervousness can be transferred into constraints, leaving the minimization of operational costs as the remaining objective function. Finally, with the Weighted-Sum method, the objective functions are added up to construct a single-objective scalarized model. By defining weights for each objective function, preferences among the objective functions can be expressed. The selection of the weights determines to which Pareto-optimal solution of the multi-objective problem the solution of the single-objective problem corresponds. In multi-objective lot sizing, tardiness costs and nervousness costs can be specified to construct a scalarized objective function.
After solving the single-objective problem, the subset of the Pareto-Front outlined by the Pareto-optimal solutions determined so far is presented to the decision maker. Based on this impression they can decide if and how additional Pareto-optimal solutions shall be determined. As soon as the decision maker has enough information, they can select a solution for this decision stage.

Heuristic approximation of the Pareto-front

As the deterministic CLSP is NP-hard already and the stochasticity complicates the problem even more, determining Pareto-optimal solutions is computationally expensive. However, most of the Pareto-optimal solutions are only determined to get an impression of the objective space, although the corresponding Pareto-optimal production plan will never be executed. Therefore, it is sufficient to heuristically determine bounds of the Pareto-front. We present a heuristic approach that allows determining precise upper and lower bounds. Lower bounds are obtained with a column generation approach, while upper bounds are determined with a Fix&Optimize heuristic applied to the reduced problem which can be derived from the terminal solution of the column generation approach.

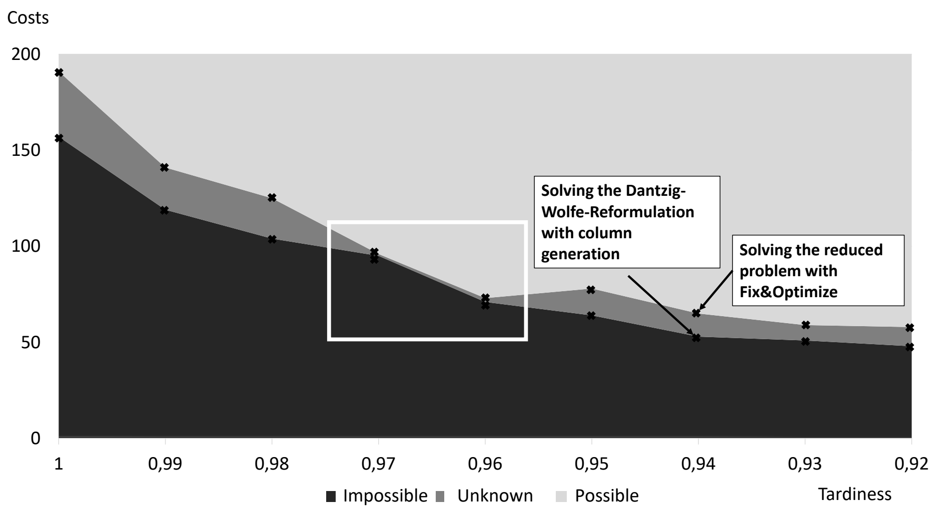

Figure 4: Approximation of the Pareto-Front

Figure 4 shows an approximate Pareto-front of a bi-objective cross-section of the Pareto-Front for a fixed level of nervousness. All combinations of the objective function values worse than the determined upper bound are known to be dominated by another solution. The combinations with objective function values lower than the lower bound are not obtainable. If the difference of the bounds is small, the decision maker can identify the interesting regions of the objective space. If requested, additional computational effort can be spent on tightening the bounds of the solutions in the interesting region and/or determining additional solutions in that area. In the presentation, we will show that the proposed heuristic approach yields strong bounds with significantly lower computational effort compared to determining Pareto-optimal solutions with an exact approach.

Distributionally robust optimization for the LSP with uncertain production yield[1]

Paula Metzker Soares
HEC Montréal, Montreal, Canada
IMT Atlantique, LS2N - UMR CNRS 6004, F-44307 Nantes, France
paula.metzker-soares@imt-atlantique.fr

Simon Thevenin
IMT Atlantique, LS2N - UMR CNRS 6004, F-44307 Nantes, France
simon.thevenin@imt-atlantique.fr

Yossiri Adulyasak
HEC Montréal, Montreal, Canada
yossiri.adulyasak@hec.ca

Alexandre Dolgui
IMT Atlantique, LS2N - UMR CNRS 6004, F-44307 Nantes, France
alexandre.dolgui@imt-atlantique.fr

Abstract

This work addresses the multi-period single-item lot-sizing problem with backorder under yield uncertainty through distributionally robust optimization. We rely on Wasserstein ambiguity set to represent the yield distribution. The resulting approach remains tractable, and it provides a production plan that remains efficient for any yield distribution described by this ambiguity set.

[1]The authors wish to thank the Region Pays de la Loire and the Canada Research Chair in Supply Chain Analytics for financial support of this research.

Introduction

In a production planning activity, lot-sizing is the process that determines production quantities to meet the demand and minimize global costs. To guarantee the production of quality goods, the production yield rate measures the expected quantity of non-defective items resulting from a lot size. However, the production system is surrounded by uncertainties that may impact the quality of these decisions. An excellent review of the lot-sizing problem with random production yield is presented in [5]. The main results of this paper are still valid, and the authors emphasize that most studies are based on strong assumptions about the yield distributions and are not very adaptable to realistic systems.

Different approaches exist to incorporate uncertainties within optimization models. While the approaches based on probability distribution become easily intractable when addressing large size problems, robust methodologies usually fail to consider realistic scenarios, being often too conservatives [4]. The distributionally robust optimization (DRO) combines the probabilistic concepts of uncertainty from the stochastic method with the tractability of the worst-case perspective from the robust optimization [2, 4]. DRO seeks to immunize the system from the worst-case probability distribution of the uncertain parameters in the ambiguity set. This ambiguity set is a family of probability distributions of the uncertain parameter characterized through partial stochastic information that can be obtained from data [2].

The distribution of the yield rate is typically inferred from historical data, but predicting true distribution of the production yield is difficult, and it is often approximated to a stationary distribution (e.g.: binomial, exponential distributions) [5]. A rather novel ambiguity set, the Wasserstein ambiguity set, was introduced by Esfahani and Kuhn [3]. Centered on the uniform distribution of S independent and identically distributed samples, this ambiguity set offers robust solutions with good performance, and it allows the decision-maker to control the conservatism of the proposed solution. In this paper, we propose a static robust lot-sizing model based on the Wasserstein ambiguity set to improve the quality of the production plan. The resulting method remains robust, computationally tractable, and it allows the decision-makers to control its conservatism by adjusting the radius of the Wasserstein metric.

The paper is organized as follows: Section describes the considered problem. Section presents the distributionally robust optimization applied to the lot-sizing problems with uncertain production yield. Finally, Section gives the main results of this work and provides some future research directions.

Problem statement

The single-item uncapacitated lot-sizing problem (LSP) with backorder and production yield rate determines the production setup Y_t and quantity X_t in a finite production horizon $t \in T$, to minimize the overall costs, and to meet the demand d_t with quality goods. Given a set of customer demands over a finite horizon, the problem determines the lot size for each period t to balance the quantity of quality goods with the backorder and the stock accumulated from previous periods to meet the demand. However, the quality goods produced can be different from the quantity ordered when the production was released. If the amount of goods available is insufficient to meet the demand, then the amount of missing goods is backordered. Any remaining amount of goods after satisfying both the demand of the current period and backordered, if any, is kept in stock. This deterministic variant of the problem does not consider uncertainty in the model. It would lead to a suboptimal or even an unrealistic plan in an uncertain context. To account for uncertain yield, we integrate stochastic information on the robust formulation through the distributionally robust optimization approach. This problem is handled within a static framework where all the decisions are defined at the beginning of the production horizon and remain fixed.

Distributionally Robust Optimization

Based on the definitions given in [3] and the notations presented in [1], we consider the Wasserstein ambiguity set $\mathcal{F}$ given in (46), where $\mathcal{P}_0$ is the family of probability distribution $\mathbb{P}$ that can describe the uncertain yield $\widetilde{\rho}$, and $d_W(\mathbb{P}, \hat{\mathbb{P}})$ is the Wasserstein metric that measures the distance between distribution $\mathbb{P}$ and the empirical distribution of the production yield $\hat{\mathbb{P}}$.

$$\mathcal{F} = \left\{ \mathbb{P} \in \mathcal{P}_0 \;\middle|\; \begin{array}{l} \widetilde{\rho} \sim \mathbb{P} \\ d_W(\mathbb{P}, \hat{\mathbb{P}}) \leq \theta \end{array} \right\}, \tag{46}$$

[3] proof that the worst case expectation for such an ambiguity can be reformulated into a tractable optimization problem. Later, [1] demonstrate that $\mathcal{F}$ can be generalized as an event-wise ambiguity set that describe each possible probability distribution in $\mathcal{P}_0$ as a scenario. Thus, $\mathcal{F}$ can be decomposed in $\|S\|$ independent and identically distributed scenarios, such that $\hat{\rho}_{\widetilde{s}}$ is the empirical distribution, here assumed to be the uniform distribution, and $u(\widetilde{\rho}, \hat{\rho}_{\widetilde{s}})$ is the epigraph of the Wasserstein metric d_W for each scenario s. Then, $\mathcal{F}$ can be represented in a lifted format associated to each scenario s in $[S]$. For that, we introduce an auxiliary random variable $\widetilde{w}$ to express the epigraph $u(\widetilde{\rho}, \hat{\rho}_{\widetilde{s}})$, resulting on the following lifted ambiguity

set $\overline{\mathcal{F}}$:

$$\overline{\mathcal{F}} = \left\{ \mathbb{P} \in \mathcal{P}_0(\mathbb{R}^{|\mathcal{T}|+1} \times [\mathcal{S}]) \;\middle|\; \begin{array}{ll} (\widetilde{z}, \widetilde{s}) = ((\widetilde{\rho}, \widetilde{w}), \widetilde{s}) \sim \mathbb{P} & \\ \mathbb{E}_{\mathbb{P}}[\widetilde{w} \mid \widetilde{s} \in [S]] \leq \theta & \\ \mathbb{P}\left[\begin{array}{l} \widetilde{\rho} \in [0,1]^{|T|} \\ u(\widetilde{\rho}, \hat{\rho}_{\widetilde{s}}) \leq \widetilde{w} \end{array} \middle| \; \widetilde{s} = s \right] = 1 & \forall s \in [S] \\ \mathbb{P}[\widetilde{s} = s] = \frac{1}{S} & \forall s \in [S] \end{array} \right\}, \tag{47}$$

Assuming $\overline{\mathcal{F}}$ given in (47) we model and implement the LSP under yield uncertainty through the DRO within a static framework with the help of the python RSOME library presented in [1]. This library makes it easier to formulate and to solve the stochastic problem, avoiding modelers to care about the reformulation step, while it allows them to directly call commercial solvers to compute a solution. Let $g(X, \tilde{\rho}) = \min \sum_{t \in T} H_t$, the inner function of the problem, be subject to $H_t \geq 0$, $H_t \geq h_t \left[\sum_{\tau=1}^{t} (\tilde{\rho}_\tau X_\tau - d_\tau)\right]$, and $H_t \geq -b_t \left[\sum_{\tau=1}^{t} (\tilde{\rho}_\tau X_\tau - d_\tau)\right]$ for all period t, where h_t, b_t and H_t are the unit inventory, the unit backordering and the inventory control cost respectively. The DRO model for the LSP with uncertain yield is given as:

$$\begin{aligned} & \min_{Y,X} \left[\sum_{t \in T} (s_t Y_t + v_t X_t) + \sup_{\mathbb{P} \in \mathcal{F}} \mathbb{E}_{\mathbb{P}}[g(X, \tilde{\rho})] \right] && \quad (48) \\ & s.t. && \\ & X_t \leq M \cdot Y_t && t \in T \\ & X_t \geq 0 \; ; \; Y_t \in \{0, 1\} && t \in T \end{aligned}$$

Main results and discussion

This paper proposes a static stochastic robust single-item multi-period uncapacitated lot-sizing problem with backordering under yield uncertainty. The presented model offers a production plan that is robust, while the production system integrates more stochastic data on the model, and it remains protected from uncertainties. The model was compared with the classical robust model under the budgeted uncertainty set. The computational experiments showed that the stochastic robust model can outperform the robust model, and it achieves satisfactory results that are free of strong assumptions, even if more stochastic information from available data is taken into account. Further studies should be carried on to test the quality of the solutions and the performance and scalability of this model.

References

[1] Zhi Chen, Melvyn Sim, and Peng Xiong. Robust stochastic optimization made easy with rsome. *Management Science*, 66(8):3329–3339, 2020.

[2] Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, 2010.

[3] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.

[4] Wolfram Wiesemann, Daniel Kuhn, and Melvyn Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.

[5] Candace Arai Yano and Hau L Lee. Lot sizing with random yields: A review. *Operations Research*, 43(2):311–334, 1995.

A Branch-and-Cut Algorithm for the Inventory Routing Problem with Product Substitution

Özlem Mahmutoğulları
ORSTAT, Faculty of Economics and Business, KU Leuven
ozlem.mahmutoullar@kuleuven.be

Hande Yaman
ORSTAT, Faculty of Economics and Business, KU Leuven
hande.yaman@kuleuven.be

Abstract

The inventory routing problem arises in vendor managed systems, in which a supplier is responsible for replenishing inventories of products at a set of retailers and manages the logistics over a finite planning horizon with a fleet of capacitated vehicles. We consider this problem in the context of the distribution of two different quality products with one-way product substitution, where the high quality product may be used to meet the demand for the low quality product. The routes of vehicles and the quantities of products sent to each retailer in each period are determined in such a way that no stockouts occur and the total cost associated with inventory holding, substitution and transportation is minimized. In this study, we derive a mixed integer linear programming formulation for the problem, strengthen this formulation with valid inequalities and develop a branch and cut algorithm as an exact solution method. We conduct experiments using benchmark and randomly generated instances to analyze the effectiveness of our solution method and investigate the relationship between product substitution decisions and system costs under different demand, supply, vehicle capacity and substitution cost settings.

Introduction

Vendor managed inventory replenishment (VMI) is a practice in which a central decision maker (the supplier) is responsible for the replenishment of inventories at a set of retailers. In VMI, the supplier monitors the inventory level of each retailer

and decides when and how many products to deliver to each retailer. The inventory routing problem (IRP) arises in this context and aims to simultaneously optimize decisions on delivery schedules, replenishment quantities and delivery routes. We incorporate the product substitution notion into the IRP for two different quality products, where a high quality product can be used to satisfy the demand for a low quality product at retailers.
We present a mixed integer programming formulation for the problem and strengthen it with valid inequalities. We apply these inequalities within a branch-and-cut algorithm. By modifying the benchmark instances in the literature, we generate instances with time dependent supplies and demands for two products. We use these instances to analyze the benefits of product substitution in decreasing costs. We also present a computational study to evaluate the effectiveness of our branch-and-cut algorithm. The results show that substitution is more advantageous when capacity limitations are tight and demand varies widely between periods.

Problem Definition and Model

The problem is defined on a complete graph $G = (V_0, E)$ where $V_0 = \{0, \dots, v\}$ is the vertex set and $E = \{\{i, j\} : i, j \in V_0, i < j\}$ is the edge set. Vertex 0 represents the supplier and the other vertices $V = V_0 \setminus \{0\}$ represent the retailers. There is a nonnegative traveling cost c_e for edge $e \in E$. The deliveries from the supplier to retailers are made by m identical vehicles, each with a capacity Q. The planning horizon consists of n periods, $T = \{1, \dots, n\}$. There are two products and product 1 has a higher quality than product 2. Each site $i \in V_0$ has $\overline{s}_{i0}^p$ units of product $p = 1, 2$ in inventory at the beginning of planning horizon and incurs a unit inventory holding cost h_{it}^p in period $t \in T$. A substitution cost f_{it} occurs for each high quality product used for substitution at retailer $i \in V$ in period $t \in T$. The supplier receives r_t^p units and retailer $i \in V$ demands d_{it}^p units of product $p = 1, 2$ in period $t \in T$. The storage capacity at retailer $i \in V$ is u_i. We assume that $\overline{s}_{i0}^1 + \overline{s}_{i0}^2 \le u_i$ and $d_{it}^1 + d_{it}^2 \le u_i$ for all $i \in V$ and $t \in T$.
To model this problem, we use the following decision variables: x_e^t represents the number of times edge $e \in E$ is traversed in period $t \in T$, y_{it} the number of vehicles that visit site $i \in V_0$ in period $t \in T$, s_{it}^p the amount of inventory of product $p = 1, 2$ at site $i \in V_0$ at the end of period $t \in T$, q_{it}^p the amount shipped from the supplier to retailer $i \in V$ in period $t \in T$, z_{it}^{12} the amount of product 1 used to substitute product 2, and z_{it}^{22} the amount of product 2 used to satisfy the demand for product 2.
The inventory routing problem with product substitution is formulated as follows:

$$\min \sum_{i\in V_0} \sum_{t\in T} (h_{it}^1 s_{it}^1 + h_{it}^2 s_{it}^2) + \sum_{t\in T} \sum_{e\in E} c_e x_e^t + \sum_{i\in V} \sum_{t\in T} f_{it} z_{it}^{12} \tag{49}$$

$$\text{s.t. } s^p_{0,t-1} + r^p_t = \sum_{i \in V} q^p_{it} + s^p_{0t} \qquad p = 1, 2, t \in T \tag{50}$$
$$s^1_{i,t-1} + q^1_{it} = d^1_{it} + z^{12}_{it} + s^1_{it} \qquad i \in V, t \in T \tag{51}$$
$$s^2_{i,t-1} + q^2_{it} = z^{22}_{it} + s^2_{it} \qquad i \in V, t \in T \tag{52}$$
$$z^{12}_{it} + z^{22}_{it} = d^2_{it} \qquad i \in V, t \in T \tag{53}$$
$$s^p_{i0} = \overline{s}^p_{i0} \qquad p = 1, 2, i \in V_0 \tag{54}$$
$$s^1_{it} + s^2_{it} + d^1_{it} + d^2_{it} \leq u_i \qquad i \in V, t \in T \tag{55}$$
$$x^t(\delta(i)) = 2y_{it} \qquad i \in V_0, t \in T \tag{56}$$
$$Qx^t(E(S)) \leq \sum_{i \in S}(Qy_{it} - q^1_{it} - q^2_{it}) \qquad S \subseteq V, t \in T \tag{57}$$
$$0 \leq x^t_e \leq 2 \qquad e \in \delta(0), t \in T \tag{58}$$
$$0 \leq x^t_e \leq 1 \qquad e \in E \setminus \delta(0), t \in T \tag{59}$$
$$s^p_{it} \geq 0 \qquad p = 1, 2, i \in V_0, t \in T \tag{60}$$
$$q^p_{it} \geq 0 \qquad p = 1, 2, i \in V, t \in T \tag{61}$$
$$0 \leq y_{it} \leq 1 \qquad i \in V, t \in T \tag{62}$$
$$0 \leq y_{0t} \leq m \qquad t \in T \tag{63}$$
$$x \text{ and } y \text{ integer.} \tag{64}$$

Objective (49) minimizes the total inventory holding, transportation and substitution costs during the planning horizon. Constraints (50)-(52) are inventory balance equations for products 1 and 2 at the supplier and retailers. Constraints (53) ensure that the demand for product 2 is satisfied on time using products 1 and 2. Initial stock levels are set by constraints (54) and maximum levels of inventories at the retailers are imposed by constraints (55). Constraints (56) are the degree constraints. Constraints (57) are capacity constraints and they eliminate subtours with positive delivery amounts. The remaining constraints are variable bounds and integrality constraints.

Inequalities and Algorithm

We propose four new classes of valid inequalities for the IRP with product substitution. The first class is the variable upper bound constraints for aggregate delivery amounts. The second class contains the common setup (l, S^1, S^2) inequalities based on the inequalities presented by [1] extending the well-known (l, S) inequalities to the two-item lot sizing problem with separate setups under one-way product substitution structure. The other classes are the lifted flow cover inequalities that are valid for our problem if there is a single vehicle. Additionally, to improve the LP relaxation bound, we adapt the inequalities proposed by [2] for the lot sizing prob-

lem with Wagner Whitin costs, the rounded capacity constraints proposed by [3] for the IRP with one product, and the connectivity constraints for the vehicle routing problems to our problem. We strengthen the Wagner Whitin inequalities and the capacity constraints by exploiting the logic of variable upper bound constraints.
The proposed branch-and-cut algorithm starts by solving a relaxation of the problem. The inequalities are added to the relaxation whenever they are violated. If there is a single vehicle, all valid inequalities are applied at the root node and the connectivity constraints are used to eliminate subtours; otherwise, all inequalities, except the lifted flow cover inequalities, are applied at the root node and the capacity constraints are used to eliminate subtours and capacity violations. We use MIP models, heuristics and exact algorithms to separate the inequalities.

Computational Experiments

We generate two sets of instances for the computational experiments. The first set employs instances from the IRP literature while the second set employs instances randomly generated with respect to the same data structure. The instance structure described by [4] is for a single product problem, with supplies and demands being time independent during the planning horizon. We introduce new parameters to incorporate the substitution option and time dependency into the instances.
First of all, we assess the effectiveness of the valid inequalities implemented within the branch-and-cut algorithm. We increase the number of instances that could be solved to optimality or improve the best bound found within the time limit using our approach. Secondly, we investigate situations where product substitution is an effective option and validate potential savings in total transportation and system costs, including transportation, holding and substitution costs. We show that product substitution can lead to noticeable savings in routing and holding costs under realistic scenarios, when demands and supplies fluctuate between periods and vehicle capacity is small.

References

[1] Hande Yaman. Polyhedral analysis for the two-item uncapacitated lot-sizing problem with one-way substitution. Discrete Applied Mathematics, 157, 3133–3151 (2009)

[2] Yves Pochet and Laurence A Wolsey. Polyhedra for lot-sizing with Wagner—Whitin costs. Mathematical Programming, 67, 297–323 (1994)

[3] Pasquale Avella, Maurizio Boccia, and Laurence A Wolsey. Single-item reformulations for a vendor managed inventory routing problem: Computational experience with benchmark instances. Networks, 65, 129–138 (2005)

[4] Claudia Archetti, Luca Bertazzi, Gilbert Laporte, and Maria Grazia Speranza. A branch-and-cut algorithm for a vendor-managed inventory-routing problem. Transportation Science, 41, 382–391 (2007)

The Consistent Production Routing Problem

Aldair Alvarez, Jean-François Cordeau, Raf Jans
HEC Montréal and GERAD, Montréal, Canada
{*aldair.alvarez, jean-francois.cordeau, raf.jans*}*@hec.ca*

Abstract

We introduce the consistent multi-plant production routing problem. In this problem, we are interested in finding minimum-cost production-routing plans that also meet specific consistency requirements. In this context, consistency is defined as the degree to which some specified features of the solution remain invariant over time. We consider four forms of consistency, namely: driver, source, product and plant consistency. In this study we present a mathematical formulation, an exact branch-and-cut algorithm, and a heuristic solution method for the problem. In a series of computational experiments we analyze the cost-consistency trade-off of the solutions, confirming that it is possible to impose consistency without excessively increasing the cost.

Introduction and Problem Definition

Production routing problems (PRPs) are planning problems that integrate several activities across the supply chain, such as lot-sizing, inventory management, and vehicle routing. Particularly, the multi-plant multi-product PRP (MPRP) considers a network with multiple production facilities (plants) and multiple customers that demand several products over a multi-period discrete planning horizon. The MPRP consists of determining a minimum-cost production and routing plan defining, for each period, the setup and production quantities of each product at each plant; the delivery quantities to each customer from each plant; and the vehicle routes required to deliver those quantities. This plan must satisfy all the production, storage and transportation capacities, ensuring that no out-of-stock occurs at any facility. In the MPRP each plant has its own fleet of homogeneous capacitated vehicles. The system incurs fixed setup and unit production costs at the plants, as well as unit inventory holding costs in every facility. Travel cost are incurred with the delivery routes.

In this study, we address a MPRP for which we are interested in finding minimum-cost production-routing plans that also meet specific *consistency* requirements. In

this context, consistency is defined as the degree to which some specified features of the solution remain invariant over time. The concept of consistency in logistics optimization has drawn significant interest in recent years, mainly because of its relevance in practical settings. On the one hand, consistency at the customer level, e.g., reducing the number of different drivers that visit each customer over a certain period, can increase their satisfaction due to familiarity and customization of the service, which can, in turn, positively impact the revenues of the company. From a server point of view, consistency can improve the efficiency of the operations and increase their productivity, potentially improving the quality of the processes. For a review of applications in the vehicle routing problem context please see [3]. Applications for inventory routing problems were presented by [1] and [2].
To the best of our knowledge, consistency has not been studied for PRPs. As such, we propose and study the consistent PRP in a setting with multiple production plants and multiple products. We refer to this problem as to the ConMPRP. In this paper, we model the ConMPRP extending the MPRP by considering four forms of consistency, as follows.

- *Driver consistency*: in this setting, it is desired that a limited number of different drivers visits each customer over the planning horizon. Solutions with this property can lead to enhanced service quality and efficiency since drivers visit the same customer locations more often.
- *Source consistency*: this form of consistency favors reducing the number of different plants from which a specific customer receives its orders over the entire horizon. The idea is to reduce the variability in the quality of the products dispatched to each customer by providing solutions in which each customer receives products from a limited set of plants.
- *Product consistency*: at each production plant, it is preferred to produce a limited number of different products throughout the planning horizon. The idea is to take advantage of limited process flexibility, i.e., reducing the set of products that each plant produces, to increase their productivity and quality.
- *Plant consistency*: in this configuration of the problem, it is preferred that each product is manufactured at a limited number of plants over the planning horizon. As a result, solutions with this feature might benefit from enhanced product quality variability.

In the ConMPRP, the different consistency requirements are addressed by defining a target maximum value for each of the features. These requirements are optimized simultaneously with the production and routing plan. Specifically, there is a target value defining the maximum number of different drivers that should visit a customer over the planning horizon. This value defines the decision-maker's tolerance in terms of driver consistency requirements for each individual customer. Additionally, there are analogous parameters defining the target maximum number of different plants that should deliver to a customer, products that each plant should produce, as well as plants at which each product should be manufactured, over the entire planning horizon.

In the ConMPRP, we impose these targets as soft constraints, given that enforcing them as hard constraints may be too restrictive and, in practice, a certain degree of violation might be acceptable. As such, the violations of the given targets are penalized in the objective function to favor the finding of solutions considering the consistency requirements. This type of approach is in line with those used in the works of [4] and [5], and provides flexibility to the system by giving the decision-maker the freedom to choose which consistency features to optimize and with what significance, according to specific situations.

Overview and Main Findings

We first modeled the problem with a vehicle flow-based mathematical formulation. An exact branch-and-cut algorithm was then presented. In this algorithm, we used a separation algorithm that solves a series of minimum $s-t$ cut problems to identify violated subtour elimination constraints. We further strengthened the formulation with several valid inequalities and defined specific branching priorities for the branch-and-cut algorithm (instead of using the solver default option).

We also proposed an iterated local search-based heuristic method to solve the problem. In our implementation, we handle the various decisions of the problem using different components, embedding the framework within a multi-start approach. In this context, we use a construction heuristic based on the solution of a production-distribution formulation, a local search heuristic that explores several routing neighborhoods, and a perturbation procedure that simultaneously changes multiple solution attributes. Also, we use an improvement routine based on a mixed-integer programming formulation, an intensification operator that employs a state-of-the-art metaheuristic, and a linear programming model to set the values of the continuous variables of the problem.

The computational experiments of this study were divided into two parts. The first one assesses the performance of the proposed solution methods while the second part focuses on analyzing the impact of taking the consistency features into account and the trade-offs arising from their inclusion. For our experiments, we generated a MPRP test set containing 2,208 problem instances, based on benchmark instances from the PRP literature.

In a first series of experiments we verified that the additional features enhance the performance of the exact algorithm with regards to the number of instances with feasible and optimal solutions as well as to the CPU time required to prove the optimality of the solutions. The experiments with the heuristic method showed that it performs robustly for the ConMPRP, the MPRP and the standard PRP. This assessment was carried out with respect to the exact method for the ConMPRP and the MPRP, while for the standard PRP we compared against eight state-of-the-art heuristic methods from the literature.

We also analyzed the trade-off between cost and consistency of the solutions in the ConMPRP context. To analyze the consistency features of the solutions, we performed experiments to observe the effect of using different values for the corresponding violation penalties. The idea is to verify the change in the cost and solution structure when altering the relative importance of the consistency requirements. The results showed that, in general, it is possible to achieve such trade-offs without compromising excessively on the cost component. This conclusion indicates that it is possible to improve the consistency properties of the production-routing plans at a (relatively) low cost if this is desirable for the decision maker. An exception to this general behavior is product consistency, for which it might be infeasible to perform the changes required to significantly improve the solution attributes that affect it. The results also revealed the significant impact of the first period when optimizing and measuring the consistency features we studied. In particular, they showed that low initial inventories at the customer locations can negatively impact the consistency features of the solutions.

References

[1] Coelho, L., Cordeau, J.-F., and Laporte, G. (2012). Consistency in multi-vehicle inventory-routing. *Transportation Research Part C: Emerging Technologies*, 24:270–287.

[2] Diabat, A., Archetti, C., and Najy, W. (2021). The fixed-partition policy inventory routing problem. *Transportation Science*, 55(2):353–370.

[3] Kovacs, A. A., Golden, B. L., Hartl, R. F., and Parragh, S. N. (2014). Vehicle routing problems in which consistency considerations are important: A survey. *Networks*, 64(3):192–213.

[4] Kovacs, A. A., Golden, B. L., Hartl, R. F., and Parragh, S. N. (2015). The generalized consistent vehicle routing problem. *Transportation Science*, 49(4):796–816.

[5] Smilowitz, K., Nowak, M., and Jiang, T. (2013). Workforce management in periodic delivery operations. *Transportation Science*, 47(2):214–230.

Author Index