15th International Workshop on Lot Sizing

The annual meeting of the EURO Working Group on Lot Sizing

University of Groningen & Erasmus University Rotterdam

Groningen, The Netherlands

August 27-29, 2025

Local Organizing Committee

Onur Kilic
Ruud Teunter
Wilco van den Heuvel
Albert Wagelmans

Steering Committee

Nabil Absi (France)
Kerem Akartunali (UK)
Christian Almeder (Germany)
Stéphane Dauzère-Pérès (France)
Silvio de Araujo (Brasil)
Raf Jans (Canada)
Hark-Chin Hwang (South Korea)
Safia Kedad-Sidhoum (France)
Florian Sahling (Germany)
Wilco van den Heuvel (Netherlands)
Laurence A. Wolsey (Belgium)

Supported by

Association of European Operational Research Societies (www.euro-online.org)
Netherlands Society for Statistics and Operations Research (www.vvsor.nl)
Erasmus Research Institute of Management (www.eur.nl/en/erim)
Econometric Institute – Erasmus University Rotterdam (www.eur.nl/en/ese/department-econometrics/about-econometric-institute)
Research Institute of the Faculty of Economics and Business – University of Groningen (www.rug.nl/research/feb-ri)

List of Participants

Name	University	Country
Ahoura Keshavarz	Universite de Lorraine	France
Albert Wagelmans	Erasmus University Rotterdam	Netherlands
Alf Kimms	University of Duisburg-Essen	Germany
Ayse Akbalik	Universite de Lorraine	France
Bahar Cennet Okumusoglu	TU Munchen	Germany
Camil Zarrouk	Ecole des mines de Saint Etienne	France
Celine Gicquel	Universit Paris Saclay	France
Chee Khian Sim	University of Portsmouth	UK
Christian Almeder	European University Viadrina Frankfurt (Oder)	Germany
Daryna Dziuba	European University Viadrina Frankfurt (Oder)	Germany
David Tremblet	HEC Montreal	Canada
Edwin Romeijn	Georgia Tech	US
Elodie Suzanne	Ecole des Mines de Saint Etienne	France
Eyup Ensar Isik	Yildiz Technical University	Turkey
Florian Sahling	RPTU University of Kaiserslautern-Landau	Germany
Franco Quezada	Universidad de Santiago de Chile	Chile
Herbert Meyr	University of Hohenheim	Germany
Katja Schimmelpfeng	University of Hohenheim	Germany
Kerem Akartunali	University of Strathclyde	UK
Leo Gilbert	IMT Atlantique	France
Lukas Bath	RPTU University of Kaiserslautern-Landau	Germany
Marcel Turkensteen	Aarhus University	Denmark
Margaretha Gansterer	University of Klagenfurt	Austria
Marion Lemke	HAW Landshut / KU Eichsttt-Ingolstadt	Germany
Maurcio Goncalves	State University of Sao Paulo	Brazil
Mourad Terzi	Ecole des Mines de Saint Etienne	France
Nabil Absi	Ecole des Mines de Saint Etienne	France
Natalia Jorquera Bravo	ENSTA Paris	France
Onur Kilic	University of Groningen	Netherlands
Raf Jans	HEC Montreal	Canada
Ruud Teunter	University of Groningen	Netherlands
Rym Ben Maktouf	Ecole des Mines de Saint Etienne	France
Safia Kedad-Sidhoum	Conservatoire National des Arts et Metiers	France
Sajjad Hedayati	Ghent University	Belgium
Silvio de Araujo	State University of Sao Paulo	Brazil
Stefan Helber	Leibniz Universitt Hannover	Germany
Stephan Koeppel	RPTU University of Kaiserslautern-Landau	Germany
Stephane Dauzere-Peres	Ecole des Mines de Saint Etienne	France
Wilco van den Heuvel	Erasmus University Rotterdam	Netherlands
Zakaria Saandi Idjabou	Universit Ibn Zohr	Morocco

Foreword

Dear lot sizing community, dear friends,

It is a great pleasure to welcome you to the 15th International Workshop on Lot Sizing. This is the second time that the workshop is organized in the Netherlands. As some of you will certainly remember, the 2012 edition was held in Rotterdam. This year we meet in the wonderful city of Groningen. The most well-known slogan for Groningen is “Er gaat niets boven Groningen” which translates to “Nothing tops Groningen”. This slogan emphasizes the unique and special qualities of the city and the province of the same name, making it a popular saying within Dutch culture. It highlights Groningen’s vibrant student life, historical architecture, and cultural scene.

We are happy to continue the tradition of the previous workshops to discuss high quality research in a relaxed atmosphere. As in the previous editions, the aim of the workshop is the following:

“The goal of the workshop is to cover recent advances in lot sizing: new approaches for classical problems, new relevant problems, integration of lot sizing with other problems such as scheduling, distribution or vehicle routing, presentation of case studies, etc. Additionally, the workshop aims to maintain an exchange of ideas between researchers and enhance fruitful collaboration.”

We would like to thank the Association of European Operational Research Societies (EURO), the Operations Research section of the Netherlands Society for Statistics and Operations Research (VVSOR), the Erasmus Research Institute of Management (ERIM), the Econometric Institute at Erasmus University Rotterdam, and the Research Institute of the Faculty of Economics and Business (FEBRI) at University of Groningen for their (financial) support in organizing the workshop.

We wish you a pleasant stay in Groningen and hope that you find the workshop inspiring and productive.

On behalf of the IWLS 2025 Organizing Committee,
Albert Wagelmans

Conference Program (1/3)

	Tuesday, August 26
19:00-21:30	**Welcome Drinks & Food** De Uurwerker (Uurwerkersplein 1, 9712 BH Groningen)

	Wednesday, August 27
8:30-9:00	Registration
9:00-9:30	Opening Session
9:30-10:30	**Session 1** – Lot-sizing with energy aspects (Chair: Wilco van den Heuvel)
	Solving electricity distribution planning problem with battery cycling considerations **Natalia Jorquera-Bravo**, Sourour Elloumi, Safia Kedad-Sidhoum, Agnès Plateau (p8)
	Energy-Oriented Batch Scheduling **Lukas Bath**, Florian Sahling (p12)
10:30-11:00	Coffee Break
11:00-12:30	**Session 2** – Lot-sizing and scheduling (Chair: Herbert Meyr)
	A construction heuristic for a lot-sizing problem with job-shop scheduling constraints **Daryna Dziuba**, Christian Almeder, Stéphane Dauzère-Pérès (p16)
	A Novel Decomposition-Based Exact Solution Approach for a Lot-sizing and Scheduling Problem **Eyüp Ensar Işık**, Z. Caner Taşkın, Semra Ağralı (p20)
	Multi-item lot sizing with fixed sequences on parallel machines **Ben Maktouf Rym**, Stéphane Dauzère-Pérès, Daniel Godard, Safia Kedad-Sidhoum (p25)
12:30-14:00	Lunch
14:00-15:00	**Session 3** – Lot-sizing with remanufacturing (Chair: Florian Sahling)
	Remanufacturing and refurbishment of pre-owned consumer electronic products under uncertainty **Mourad Terzi**, Nabil Absi, Xavier Schepler, Antoine Jeanjean (p29)
	Dynamic lot sizing model with remanufacturing and separate setup costs: Time complexity and optimality **Chee-Khian Sim** (p33)
16:15-18:00	**Boat tour** De Toeter (Turfsingel 6, 9712 KP Groningen)

Conference Program (2/3)

	Thursday, August 28
9:00-10:30	**Session 4** – Integrated lot-sizing problems (Chair: Christian Almeder)
	Solving a Cutting and Planning Problem with Reusable Leftovers **Léo Gilbert**, Maria I. Restrepo, Nadjib Brahimi, François Klein (p34)
	Robust Production Planning Under Inherent Supply Chain Disruptions Amin Javanmard, **Kerem Akartunali**, Ashwin Arulselvan (p38)
	Capacitated lot-sizing problem with stock-out based substitution **David Tremblet**, Raf Jans, Yossiri Adulyasak, Sonja Ursula Katharina Rohmer (p42)
10:30-11:00	Coffee Break
11:00-12:00	**Session 5** – Lot-sizing with joint decisions (Chair: Stefan Helber)
	Decentralized Lot Sizing: Smart Contracting and Collaboration Anton Klymenko, **Margaretha Gansterer** (p46)
	Joint economic lot sizing with simultaneous determination of the required amount of product-specific reusable transport items (RTIs) **Marion Lemke**, Heinrich Kuhn (p50)
12:00-12:30	Group Photo followed by Meeting Euro Working Group on Lot-sizing
12:30-14:00	Lunch
14:00-15:00	**Session 6** – Lot-sizing with carbon considerations (Chair: Nabil Absi)
	The trade-off between costs and carbon emissions from lot-sizing decisions **Marcel Turkensteen**, Wilco van den Heuvel (p54)
	Carbon-Aware Stochastic Lot-Sizing with Perishable Inventory **Bahar Cennet Okumuşoğlu**, Martin Grunow (p58)
16:30-18:00	**Guided City Tour** Flonk Hotel (Groningen Center) (Radesingel 50, 9711 EK Groningen)
19:00-22:30	**Conference dinner** Feithhuis (Martinikerkhof 10, 9712 JG Groningen)

Conference Program (3/3)

	Friday, August 29
9:00-10:30	**Session 7** – Multi-Stage lot-sizing (Chair: Stephane Dauzere-Peres)
	Dynamic Re-optimization of Timed Routes in Multistep Production Planning **Camil Zarrouk**, Nabil Absi, Quentin Christ, Stéphane Dauzère-Pérès, Renaud Roussel (p64)
	Decision Strategies on a Two-level Multi-Stage Stochastic Lot Sizing Problem considering Blending **Maurício Rocha Gonçalves**, Raf Jans, Alexandre de Araujo (p68)
	Stochastic Lot-Sizing and Scheduling for Quality-Differentiated Disassembly **Sajjad Hedayati**, Stijn De Vuyst (p73)
10:30-11:00	Coffee Break
11:00-12:30	**Session 8** – Lot-sizing under uncertainty (Chair: Albert Wagelmans)
	Lot-sizing under decision-dependent uncertainty: A probing-enhanced stochastic programming approach **Franco Quezada**, Céline Gicquel, Safia Kedad-Sidhoum, Bernardo Pagnoncelli (p77)
	The Production Routing Problem with Stochastic Demand and Service Levels Ali Kermani, Jean-François Cordeau, **Raf Jans** (p81)
	A simple approach for computing inventory policies based on function approximation **Onur Kilic**, Armagan Tarim (p85)
12:30-14:00	Lunch

Solving electricity distribution planning problem with battery cycling considerations

Natalia Jorquera-Bravo
UMA, ENSTA Paris, IP Paris, France.
CEDRIC, CNAM, 75003 Paris, France.
PDSPS, Department of Industrial Engineering, University of Santiago of Chile.
natalia.jorquera@ensta-paris.fr

Sourour Elloumi
UMA, ENSTA Paris, IP Paris, France.
CEDRIC, CNAM, 75003 Paris, France.
ENSIIE, 91025 Evry, France.
sourour.elloumi@ensta-paris.fr

Safia Kedad-Sidhoum
CEDRIC, CNAM, 75003 Paris, France.
safia.kedad_sidhoum@cnam.fr

Agnès Plateau
CEDRIC, CNAM, 75003 Paris, France.
agnes.plateau_alfandari@cnam.fr

Abstract

We address an electricity distribution planning problem in a self-consumption community where both generation and battery storage are shared among all members. Each user has an individual demand, but energy generation and storage are centrally managed. No energy exchanges between users are allowed. The goal is to minimize the total community cost, which encompasses energy procurement expenses and incorporates a constraint restricting battery cycling. We formulate the problem as a mixed-integer program and evaluate its performance under varying demand scenarios. Our results show that explicitly modeling battery cycling impacts operational decisions and overall system cost.

1 Problem statement

In this work, we study energy communities formed within buildings where residents collectively invest in photovoltaic (PV) panels and a shared electricity storage system. Each household can meet its electricity demand through energy supplied by the PV panels, the shared battery, or the main grid. Surplus electricity can either be stored in the battery or sold back to the grid; however, energy exchanges between users are not permitted. Each household is equipped with a smart meter that monitors consumption at each time period.

A central community manager is responsible for distributing the PV-generated electricity at each time step. Operational constraints include: (i) a user cannot simultaneously charge and discharge the battery, and (ii) a user cannot buy and sell electricity to the grid within the same time period.

We model the community as a microgrid comprising a set $\mathcal{H}$ of smart homes, a PV panel array, a shared battery system, and a connection to the main grid. The planning horizon consists of discrete time intervals $\mathcal{T}$ of duration δ hours, typically $\delta = 0.25$ (15 minutes). Each household $j \in \mathcal{H}$ has a known electricity demand $D_{j,t}$ at each time period $t \in \mathcal{T}$. The battery is characterized by charging and discharging efficiencies $\gamma_c, \gamma_d \in (0, 1)$, an initial state of charge S_0, capacity bounds $S^{\min}$ and $S^{\max}$, and maximum charge/discharge rates F_c and F_d. At each time step t, PV panels generate C_t^{PV}. Electricity can be purchased from the grid at a time-dependent price ν_t, discharged from the battery at a cost μ, or sold to the grid at a fixed price β, set by the French Energy Regulatory Commission. As is typical in France, the sale price β is strictly less than the purchase price ν_t for all $t \in \mathcal{T}$. We do not account for PV generation costs, assuming these are negligible due to the absence of fuel and operational labor, and considering maintenance as part of the upfront investment.

As a main extension of the study presented in [4], we analyze battery health by considering the number of charging cycles. A charging cycle begins when the battery switches from an active discharging state to a charging state, and it ends when it switches from an active charging state to an active discharging state. We present a mathematical formulation to address the electricity distribution planning problem in collective self-consumption communities and incorporates battery cycling dynamics. We illustrate the impact of modeling battery cycling on operational decisions.

2 Mathematical formulation

We propose a mixed-integer linear programming (MILP) model to address the electricity distribution planning problem, incorporating battery cycling dynamics. The model determines, for each house $j \in \mathcal{H}$ and time period $t \in \mathcal{T}$, the electricity supplied from three sources: photovoltaic generation $p_{j,t}$, battery discharge $y_{j,t}$, and the main power grid $i_{j,t}$.

Simultaneously, the model optimizes the amount of electricity each house stores in the battery $z_{j,t}$ (charging), sells back to the grid $g_{j,t}$, and the battery's state of charge s_t at the end of each period.

To capture the battery charging behavior over time, we introduce binary variables

v_t, v_t^{down}, and v_t^{up}, which indicate the charging status. Specifically, $v_t = 1$ if there exists a household $j \in \mathcal{H}$ such that $z_{j,t} > 0$, and $v_t = 0$ otherwise. The variable $v_t^{down} = 1$ denotes the start of a charging cycle, while $v_t^{up} = 1$ represents the end of a charging cycle, and thus the beginning of a discharging phase.

$$
\begin{cases}
\min_x \quad c(x) = \delta \sum_{j\in\mathcal{H}, t\in\mathcal{T}} (\mu y_{j,t} + \nu_t i_{j,t} - \beta g_{j,t}) + \eta^c \sum_{t\in\mathcal{T}} v_t^{up} & & \text{(1a)}\\
\text{s.t.} \quad \sum_{j\in\mathcal{H}} p_{j,t} = C_t^{PV} & \forall t \in \mathcal{T} & \text{(1b)}\\
p_{j,t} + y_{j,t} + i_{j,t} - z_{j,t} - g_{j,t} = D_{j,t} & \forall j \in \mathcal{H}, t \in \mathcal{T} & \text{(1c)}\\
s_t = s_{t-1} + \delta\gamma_c \sum_{j\in\mathcal{H}} z_{j,t} - \frac{\delta}{\gamma_d} \sum_{j\in\mathcal{H}} y_{j,t} & \forall t \in \mathcal{T} & \text{(1d)}\\
\sum_{j\in\mathcal{H}} y_{j,t} \leqslant F_d(1 - v_t) & \forall t \in \mathcal{T} & \text{(1e)}\\
\sum_{j\in\mathcal{H}} z_{j,t} \leqslant F_c v_t & \forall t \in \mathcal{T} & \text{(1f)}\\
S^{\min} \leqslant s_t \leqslant S^{\max} & \forall t \in \mathcal{T} & \text{(1g)}\\
v_t - v_{t-1} = v_t^{down} - v_t^{up} & \forall t > 1 & \text{(1h)}\\
v_t^{up} + v_t^{down} \leqslant 1 & \forall t & \text{(1i)}\\
v_1 = v_1^{down}; \qquad s_{|\mathcal{T}|} = S_0 & & \text{(1j)}\\
s_t,\ i_{j,t},\ g_{j,t},\ z_{j,t},\ y_{j,t},\ p_{j,t} \geqslant 0 & \forall j \in \mathcal{H},\ t \in \mathcal{T} & \text{(1k)}\\
v_t, v_t^{down}, v_t^{up} \in \{0,1\} & \forall t \in \mathcal{T} & \text{(1l)}
\end{cases}
$$

Equation (1a) defines a well accepted objective function used to compute the total cost of the community (see [2]). In this formulation, all equipment capacities are considered fixed and known in advance. Consequently, the model only includes operational and maintenance costs. The objective accounts for three components: (i) the maintenance costs related to battery discharging, (ii) the cost of electricity purchased from the main grid, and (iii) the revenue from selling surplus electricity back to the grid. Additionally, a cost per battery cycle is incorporated to reflect degradation effects. We define a representative cycle degradation cost η^c, calculated as $\eta^c = \frac{R}{n}$, where R denotes the battery replacement cost, and n is the number of cycles to failure under a typical 80% depth of discharge.

Equations (1b) represent the electricity generation of photovoltaic panels for each household. Constraints (1c) ensure the energy balance for each house at every time period, capturing the interaction between PV production, battery storage, and grid exchanges.

Battery operation is governed by Constraints (1d)-(1j). In particular, Equation (1d) models the state of charge of the battery at the end of each time period as a function of its previous state, the energy charged, and the energy discharged during that period. Constraints (1e) and (1f) impose upper bounds on the discharging and charging power, respectively. Following the methodology proposed in [1], Constraints (1h) are used to detect the transitions between charging and discharging cycles. Specifically, the binary variable $v_t^{down} = 1$ if a new charging cycle begins at time t, which occurs when the battery was not being charged in the previous period ($v_{t-1} = 0$) and is being charged in the current one ($v_t = 1$). Similarly, $v_t^{up} = 1$ if a charging cycle ends and a discharging cycle begins, i.e., $v_{t-1} = 1$ and $v_t = 0$. When $v_{t-1} = v_t$, no transition takes place, and both v_t^{down} and v_t^{up} are set to zero.

Without loss of generality, we define the initial condition at time $t = 1$ as follows: if any household is charging the battery at $t = 1$, then a charging cycle is assumed to start, and we set $v_1^{down} = 1$ and $v_1^{up} = 0$. Otherwise, if no charging occurs at $t = 1$, no cycle is initiated, and $v_1^{down} = v_1^{up} = 0$. These conditions are formalized in Constraints (1j).

In this talk, we will present some preliminary results on a real case study in France. These experiments demonstrate the impact of integrating battery health considerations into operational decisions and system efficiency.

Acknowledgments This work has been carried out at the Energy4Climate Interdisciplinary Center (E4C) of IP Paris and Ecole des Ponts ParisTech, which is in part supported by 3rd Programme d'Investissements d'Avenir [ANR-18-EUR-0006-02]. The authors also acknowledge the financial support from Programa de cooperación científica ECOS-ANID project number ECOS230013 and Programa Regional STIC-AMSUD project number AMSUD240008.

References

[1] Chiara Bordin, Harold Oghenetejiri Anuta, Andrew Crossland, Isabel Lascurain Gutierrez, Chris J. Dent and Daniele Vigo; A linear programming approach for battery degradation analysis and optimization in offgrid power systems with solar energy integration, Renewable Energy, 101, 407-430 (2017)

[2] Fausto Calderon-Obaldia, Jordi Badosa, Anne Migan-Dubois and Vincent Bourdin; A Two-Step Energy Management Method Guided by Day-Ahead Quantile Solar Forecasts: Cross-Impacts on Four Services for Smart-Buildings, Energies, 13 (2020)

[3] Di Zhang, Songsong Liu and Lazaros G. Papageorgiou; Fair cost distribution among smart homes with microgrid, Energy Conversion and Management, 80, 498-508(2014)

[4] Natalia Jorquera-Bravo, Sourour Elloumi, Safia Kedad-Sidhoum and Agnès Plateau; Ensuring fair allocation of renewable energy in microgrids for supply planning, IWLS(2023)

Energy-Oriented Batch Scheduling

Lukas Bath
RPTU University Kaiserslautern-Landau, Germany
lukas.bath@rptu.de

Florian Sahling
RPTU University Kaiserslautern-Landau, Germany
florian.sahling@rptu.de

Abstract

Efficient energy usage and the integration of renewable energy sources are essential for sustainable industrial manufacturing. We present a new mathematical model formulation for energy-oriented scheduling of batch processes. This model considers multiple products and parallel production resources with different energy-related production modes. The energy consumption profiles of the batch processes are mapped in 15-minute intervals, providing a detailed representation of time-dependent power demand. This enables peak consumption and associated costs to be reduced. The model also incorporates an on-site photovoltaic and energy storage systems. The objective is to determine a feasible production schedule that minimizes production and energy costs. Preliminary results show that the integration of energy-related production modes with a fine-mesh period grid may result in substantial cost savings.

1 Introduction

Energy consumption in the manufacturing industry is a central challenge in addressing climate change. Accordingly, reducing energy consumption through suitable methods is a key objective for companies. In this context, both total energy consumption and resulting peak loads are important and of particular relevance. For example, the installation of photovoltaic systems can decrease the dependency on external electricity procurement. The integration of energy storage systems further enhances grid independence and supports the transition toward climate neutrality. These changes have both economic and ecological benefits for the companies concerned in times of rising energy prices.

Overviews of energy-oriented planning problems are provided by [1] and [2]. The present work focuses on energy-oriented scheduling of batch processes. In this short-term planning context, detailed consideration of process-specific energy profiles is essential. The resulting energy requirements must be satisfied in addition to the product-specific demand. A mathematical model is formulated for an energy-oriented batch scheduling problem with limited peak power consumption and the integration of an energy storage system in a multi-product setting (EO-BSP-PP-ESS-MP).

2 Problem statement

The EO-BSP-PP-ESS-MP addresses the production of I products over a discrete planning horizon of T periods. By the end of the planning period, the demand for each product (d_i) must be completely satisfied. There are R production resources, each with a specific capacity of c_r.

Production takes place in batches, with a maximum of one batch processed per resource and period. Preemption is not allowed. Each batch is manufactured using one of M production modes (also referred to as operating modes), which represent combinations of energy consumption (e.g., eco or sprint mode) and filling level of the production resource (e.g., 75 % or 100 %). For example, the resulting operating modes may include eco_{75}, eco_{100}, $sprint_{75}$, and $sprint_{100}$. A batch of product i may require five periods in mode 1 and eight in mode 2. Production costs are denoted by bc_{imr} monetary units. Moreover, each batch is associated with an energy profile $ed_{imr\tau}$, which specifies the power consumption in production period τ for product i in mode m on resource r. An example of such a profile is shown in Figure 1.

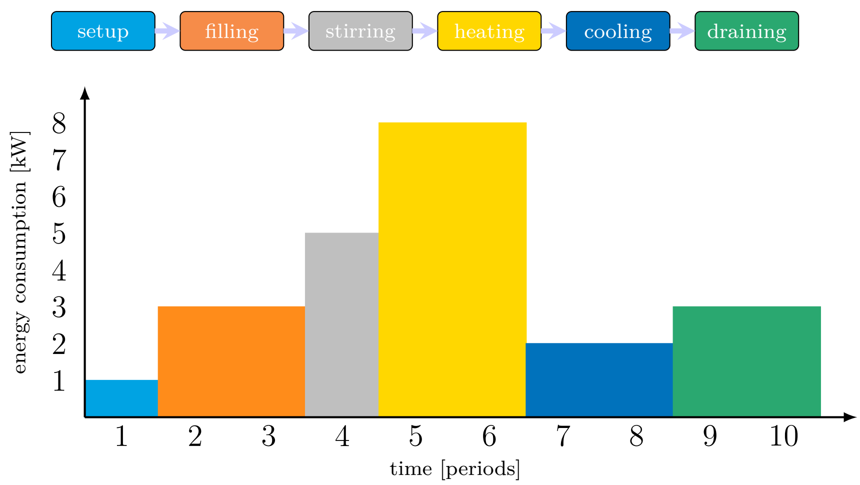

Figure 1: Exemplary energy consumption of a process

The example process consists of six sequential stages. For each stage, the required power is depicted over the duration of the respective stage. Machine assignments remain fixed throughout the process, i.e., each batch is fully processed on a single production resource. The production system described is supplemented by the energy-management-system in Figure 2.

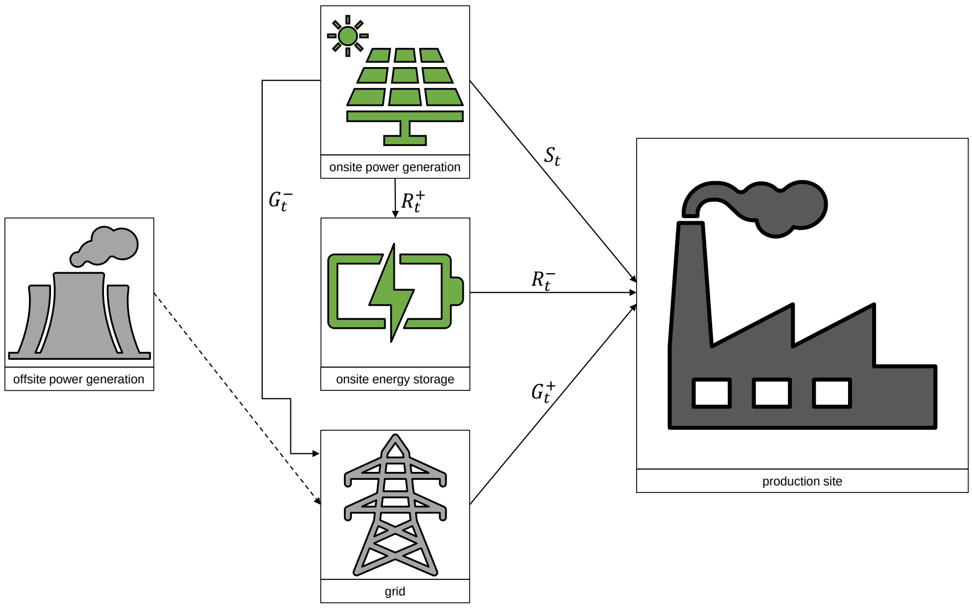

Figure 2: Energy flow of the production system

The energy requirements of a batch can be covered by purchasing energy from the power grid (G_t^+), using on-site photovoltaic generation (S_t), or discharging of the energy storage system (R_t^-). The storage system has a maximum capacity $S^{\max}$ and can either be charged or discharged during each period. Photovoltaic electricity can be used directly for production, stored in the energy storage system (R_t^+), or fed into the power grid (G_t^-).

Purchasing energy incurs period-specific costs (ec_t^{grid}), while feeding energy into the power grid yields revenue (er_t^{res}). In addition to these variable energy costs, the maximum power drawn from the grid ($PP^{\max}$) is priced using the power price (pc). This pricing structure is intended to incentivise avoiding peak production loads.
In the presentation, we will present preliminary numerical results.

Acknowledgement

This research is part of the project ”Smart batch processes in the energy system of the future” funded by the Carl-Zeiss-Stiftung (CZS) in the ”CZS Breakthroughs” program.

References

[1] Bänsch, K., Busse, J., Meisel, F., Rieck, J., Scholz, S., Volling, T., Wichmann, M.G. Energy-aware decision support models in production environments: A systematic literature review. Computers & Industrial Engineering, 159, 107456 (2021).

[2] Terbrack, H., Claus, T., Herrmann, F. Energy-oriented production planning in industry: A systematic literature review and classification scheme. Sustainability, 13(23), 13317 (2021).

A construction heuristic for a lot-sizing problem with job-shop scheduling constraints

Dziuba Daryna
European University Viadrina, Department of Supply Chain Management, Frankfurt (Oder), Germany
dziuba@europa-uni.de

Almeder Christian
European University Viadrina, Department of Supply Chain Management, Frankfurt (Oder), Germany
almeder@europa-uni.de

Dauzère-Pérès Stéphane
Mines Saint-Etienne, Univ. Clermont Auvergne, CNRS, UMR 6158 LIMOS
dauzere-peres@emse.fr

Abstract

The integration of lot-sizing and scheduling decisions usually ensures better capacity utilization and is more cost-effective. However, this integration introduces greater complexity than the classical hierarchical approach, in which lot-sizing decisions are made first and used as input for scheduling activities second. We propose a fast construction heuristic to address the capacitated lot-sizing problem that integrates job-shop scheduling constraints. Evaluating the available capacity in this problem differs substantially from evaluating it in classical lot-sizing problems, in which periods are independent in terms of capacity. We formalize this subproblem, which is critical for the proposed construction heuristic and propose several solution strategies. Finally, we present some numerical results illustrating the efficiency of the proposed construction heuristic and benchmarking it against other solution methods.

1 Problem description

We consider the multi-item capacitated lot-sizing problem with job-shop scheduling constraints. This problem allows for the planning of quantities to produce in a horizon discretized in periods and the detailed scheduling of these production quantities on machines. The goal is to minimize the total inventory and set-up cost while ensuring that the operational capacity constraints are satisfied. We propose a simpler and faster heuristic than Wolosewicz et al. (2015) to solve the aforementioned problem by extending the construction heuristic proposed in Dziuba and Almeder (2023).

In this problem, production quantities of different items must be optimized on a discretized planning horizon of H periods to meet demands at the end of each period and to minimize various costs. Let us denote by X_{it}, respectively d_{it}, the lot size, respectively, the demand, of the item i that must be completed, respectively, delivered, in period t. Each lot must be processed through a series of operations, called route, before being completed, and each operation requires a machine. Let us denote by o_{itk} the k-th operation of a lot of item i that is finished in period t. Therefore, operation o_{itk} must be finished before o_{itk+1} starts. o_{itK_i} denotes the last operation in the route of item i. Moreover, the sequence in which operations are processed on each machine is given. Let us denote by $\mathcal{S}$ the set of pairs of operations (o, o') such that o is sequenced before o' on a machine.

To allow for a compact mathematical model, the indices i, t, and k are used for all the parameters and variables related to the operation if applicable. Let C_{itk} denote the completion time of operation o_{itk}, and p_{itk}, respectively s_{itk}, denote the processing time per unit of item i, respectively the fixed setup time of item i. Moreover, let Y_{it} be the binary setup variable associated with the lot of item i that must be completed in period t, i.e. $Y_{it} = 1$ if $X_{it} > 0$ and is equal to 0 otherwise.

The following constraints ensure that the completion times respect the precedence relations given by the item routes and the machine sequences.

$$C_{it(k+1)} \geq C_{itk} + p_{it(k+1)} X_{it} + s_{it(k+1)} Y_{it} \qquad \forall i, t, k \leq K_i - 1 \tag{1}$$

$$C_{i't'k'} \geq C_{itk} + p_{i't'k'} X_{i't'} + s_{i't'k'} Y_{i't'} \qquad \forall i, i', t, t', k, k' \text{ where } (o_{itk}, o_{i't'k'}) \in \mathcal{S} \tag{2}$$

The feasibility of a plan is given, if the last operation in the item route finishes in the respective period, i.e.,

$$\sum_{l=1}^{t-1} capa_l \leq C_{itK_i} \leq \sum_{l=1}^{t} capa_l \qquad \forall i, t \tag{3}$$

Note that the processing of lot X_{it} can start before period t, i.e. operations o_{itk} such that $k < K_i$ can start before period t.

2 2-SCH for CLSP with job-shop scheduling constraints

The 2-step construction heuristic (2-SCH) is a greedy construction heuristic originally developed for a single-level multi-item capacitated lot-sizing problem (CLSP). Regardless of the application, the general framework of the method remains unchanged.

2-SCH transforms the original problem into two subproblems:

1. *Demand integration subproblem*, where a single demand element, i.e., a demand for a specific item and a specific period, must be integrated into an existing partial production plan $PP^n = \{X^n_{it}, Y^n_{it}, I^n_{it}, Schedule\}$ consisting of lot sizes, auxiliary decision variables for setups and inventory levels, and the corresponding schedule at minimal cost. 2-SCH solves this subproblem approximately by a case-based heuristic.

2. *Demand sorting subproblem*. Here the demand elements are sorted before being added step-by-step to the production plan by solving the previous subproblem. This is a sequencing problem with an implicit, complex objective function, where there is no direct mapping between similar sequences and similar costs. In 2-SCH, it is solved by a simple technique like static sorting rules.

The main adjustment to 2-SCH in the context of integrated lot-sizing and scheduling is the estimation of available capacity. Another substantial change is the replacement of the notion of overtime by the notion of tardiness. Finally, even though the lot sizing part of the problem remains in the discretized time horizon, the scheduling part lies in the domain of continuous time (capacity). This makes it necessary to simulate the decisions when several periods must be considered to allocate the production for one demand, i.e. when splitting a lot.

3 Capacity estimation

Available capacity needs to be assessed in order to determine the production quantity that can be produced without violating the capacity constraints. We consider different ways to compute the feasible production quantity as shown further.

Simple formula. The first strategy relies on the approach presented in Gomez Urrutia (2014) and used in Wolosewicz et al. (2015), that allows to determine the quantity of item i that can be produced in period t as $Q_{it} = \min_{k=1,\dots K} Q^{max}_{itk}$, where Q^{max}_{itk} is the maximal quantity that can be produced by operations up to the k^{th} operation in the route with available capacity using the slack time of operations. $Q^{max}_{itk} = \frac{slack_{itk} - s^{total}_{ik} \cdot (1 - Y^{n-1}_{it})}{p^{total}_{ik}}$, where the total setup time s^{total}_{ik} and the total processing time p^{total}_{ik} of operations in the route up to position k, and the existing slack of operation o_{itk}, $slack_{itk}$, is only reduced by the total setup time if there is no lot for item i in period t (i.e. $Y^{n-1}_{it} = 0$). This is a pessimistic estimate, which gives only a lower bound on the possible feasible amount that could be produced without violating capacity.

Improved formulas. In another strategy, we consider a so-called additional slack that can be used for production and setup in operation $o_{itk'}$ without postponing the start of operation o_{itk}, if $o_{itk'}$ precedes o_{itk} in an item route and $slack_{itk'} > slack_{itk}$. The additional slack can be calculated either by an exact algorithm that identifies the longest path from $o_{itk'}$ to o_{itk} and its length ($addSlack^{e}_{k'k}$), or by a conservative formula that is an approximation ($addSlack^{c}_{k'k} = max(0; slack_{itk'} - slack_{itk})$).

By accounting for the additional slack between all pairs of predecessors-successors in the item route, we can obtain an improved lower bound on the feasible production quantity.

Simulation. The last strategy allows to determine the accurate quantity that could be produced by iterative simulation. This simulation is a loop, where the quantity calculated by one of the strategies described above is added to the production plan and then the formula is applied again. The loop terminates when the calculated quantity is zero, i.e. there is no more available capacity. Such an iterative simulation might be time-consuming, although the number of iterations is expected to be reduced with the improved formula. Runtime could be managed by termination criteria such as stopping after a predefined number of iterations or when the calculated maximal quantity is below some threshold.

Table 1 reports the normalized capacity estimates (factors), defined as ratio of the amount obtained by a strategy to the amount obtained by simulation. We distinguish between the cases, in which capacity is estimated when an existing lot should be extended, and those, when a new lot should be created.

	Average factor ($^{***} - p \leq 0.001$)		
	Extend existing lot	Create new lot	Overall
Simple formula	0.54^{***}	0.49^{***}	0.51^{***}
Improved formula ($addSlack^{c}$)	0.58^{***}	0.54^{***}	0.56^{***}
Improved formula ($addSlack^{e}$)	0.59^{***}	0.55^{***}	0.57^{***}
Simulation (benchmark)	1.00	1.00	1.00

Table 1: Difference in capacity estimation across all strategies.

References

Dziuba, D., Almeder, C., 2023. New construction heuristic for capacitated lot sizing problems. European Journal of Operational Research 311, 906–920.

Gomez Urrutia, E.D., 2014. Optimisation intégrée des décisions en planification et ordonnancement dans une chaîne logistique. Phd thesis. École Nationale Supérieure des Mines de Saint-Étienne. Gardanne, France. Available at `https://theses.hal.science/tel-01089338`.

Wolosewicz, C., Dauzère-Pérès, S., Aggoune, R., 2015. A lagrangian heuristic for an integrated lot-sizing and fixed scheduling problem. European Journal of Operational Research 244, 3–12.

A Novel Decomposition-Based Exact Solution Approach for a Lot-sizing and Scheduling Problem

Işık, Eyüp Ensar
Department of Industrial Engineering, Boğaziçi University
eeisik@yildiz.edu.tr

Taşkın, Z. Caner
Department of Industrial Engineering, Boğaziçi University
caner.taskin@bogazici.edu.tr

Ağralı, Semra
Department of Industrial Engineering, MEF University
agralis@mef.edu.tr

Abstract

In this study, we focus on a single-machine single-level lot-sizing and scheduling problem with sequence-dependent setup costs. The problem is essential for production planning due to its effect on overall efficiency. To find efficient solutions, we develop a decomposition-based exact solution approach. Using this approach, we decompose the problem and use the integer L-shaped cuts to reach exact solutions. In the decomposed model, our master problem determines which products are produced in which period and in what quantity. On the other hand, our subproblem determines the sequence of production that is decided in the master problem. We are using the integer L-shaped cuts since our subproblem has binary variables. However, the integer L-shaped cuts are weak in their basic version. Therefore, we perform improvements on the basic model. We start by strengthening the cuts by calculating better lower bounds and through lifting. Afterwards, we disaggregate the cuts to develop another set of cutting planes. Our preliminary studies show that the algorithm spends most of its time solving the subproblem. Considering this, we introduce a dynamic programming-based algorithm to efficiently solve our subproblem. We make our computational tests on a benchmark dataset to compare the efficacy of the developed algorithm and the improvements. The results show that our decomposition-based exact solution approach outperforms exact solution approaches in the literature.

1 Introduction

Lot-sizing and scheduling problems are one of the well-known NP-hard problems in the literature [2]. The problem considers two of the well-studied problems, called the lot-sizing problem and the scheduling problem, simultaneously. To achieve effective and optimal production plans, making these decisions simultaneously is critical. Solving this problem optimally provides a balance between customer demand and costs [10]. By deciding the lot sizes, customer demands are satisfied while minimizing the production-related costs. Simultaneously, sequence decisions of the production lots are made to minimize the total setup cost.

In the literature, the lot-sizing and scheduling decisions are made hierarchically in many production processes [8]. However, there are also many studies that recognize the importance of making lot-sizing and scheduling decisions simultaneously. Therefore, several models and solution methods have been developed to solve this problem efficiently [2]. The General Lot-Sizing and Scheduling Problem (GLSP) uses macro and micro periods to address the impracticality of using smaller subperiods [4]. The problem is called "general" since it can be adapted to other problem variants. Therefore, GLSP has been studied more than other variants of the problem in the literature.

GLSP is well-studied in the literature. However, most of the studies work on meta or matheuristics [2]. To the best of our knowledge, there is no exact solution method other than the methods that use MIP formulations that are solved using commercial solvers.

2 Problem Description and Solution Method

In this study, we consider a single-machine, single-level lot-sizing and scheduling problem with sequence-dependent setup costs. The problem can be formalized using the GLSP model proposed by [4]. This model is the basic formulation for GLSP, but there are other formulations that are improved versions of the basic model. Most of these formulations are classified in [5]. The authors emphasize that both the network-flow reformulation of GLSP (GLSP$^{\mathrm{NF}}$) and the facility-location reformulation proposed by [1] (GLSP$^{\mathrm{CC}}$) provide better results.

We decompose the problem into two phases. In the first phase (master problem), lot-sizing decisions are made to minimize the inventory holding cost and estimated setup cost. In the second phase (subproblem), the sequences of the lots that are decided to be produced in the first phase are decided, and it can be considered as a single-machine scheduling problem with sequence-dependent setup times, release dates, and due dates. The subproblem can be formally defined as in the [3]. By decomposing the problem, we significantly decrease the number of variables in the master problem by eliminating micro-periods. The complexity of our master problem is $\mathcal{O}(|P||T|)$, while the basic model is $\mathcal{O}(|P|^2|T|)$. As can be easily observed, our subproblem has binary integer variables. Therefore, basic decomposition approaches (like Benders decomposition) cannot be used. Here, we adapt the integer L-shaped method proposed by [9] for two-stage stochastic programming integer programs. Using this method, we can drive our optimality cuts.

2.1 Algorithmic Improvements

Since our preliminary studies show that the basic algorithm is not effective for medium and large-sized instances, we perform improvements to the algorithm. We started with cut lifting by observation on the setup costs, which have a triangular property; therefore, it cannot be decreased by adding new productions. Using this observation, we reformulate the L-shaped cuts and prove that they are valid and stronger than the basic version of the cuts for triangular setup instances.

Another observation on integer L-shaped cuts is the impact of the lower bound value for estimated setup cost on its performance. To increase the performance of the cuts, we calculate two types of lower bounds: Global Lower Bound, which is the minimum setup cost of all products that have a demand if they are produced in only one period, and Local Lower Bound, which is calculated according to the feasible solution set of each iteration.

Afterwards, we disaggregate our cuts for each period and between periods to provide stronger cuts. To generate these cuts, we need an efficient way to calculate a lower bound on the setup cost for each period. We use a dynamic programming algorithm (*MinSetup*) that recursively computes the minimum setup cost of each iteration to calculate lower bounds. We extend this algorithm to calculate the minimum setup cost of the whole sequence. The extended version recursively divides the whole sequence and solves the left and right parts separately with fixed starting and ending products. It continues to do that until reaching a single period, then it uses the algorithm *MinSetup*.

3 Computational Study

To test the effectiveness of our decomposition algorithm (DA), we use benchmark instances from [7] and compare it with other exact solution approaches developed for GLSP. We also test the performance of the algorithmic improvements by comparing the basic DA with its improved versions. We use a subset of the benchmark instances that includes: $|P| \in \{3, 4, 5, 10, 15\}$, $|T| \in \{3, 4, 5\}$, $Cap \in \{0.6, 0.8\}$, $CapVar \in \{0.5\}$, and $CF \in \{50, 100\}$. We made our computational study using the IBM CPLEX 22.1.1 (64-bit) solver on C++ with a time limit of 600 seconds. We worked on a server with an Intel Xeon Silver 4214R 2.4 GHz CPU, 64 GB RAM, running Windows Server 2019. All data and source code are accessible at [6].

We start the comparison with the basic version of our DA and three MIP formulations in the literature. The results show that MIP models and DA can optimally solve small-sized instances. For larger-sized instances, the performance of all approaches decreases. However, $GLSP^{NF}$ has the best average solution time for small and medium-sized instances. For large-sized instances, $GLSP^{CC}$ has better performance, but it has a problem in producing feasible solutions for the largest sizes. Here, DA can achieve the best solutions among them.

We test the performance of our improvements in three phases. First, we focus on the impact of cut lifting and the lower bounds. Second, we also use disaggregated cuts. Finally, we use the dynamic programming approach to solve our subproblem. According to the results of the first phase, cut-lifting and lower-bound calculations significantly improved the performance of the algorithm. The number of cuts added to the MP is decreased, and the solution time and quality improved for all instances. Now the algorithm can solve most of the medium-sized instances to optimality. In the second phase, we can solve all medium-sized instances to optimality. For small-sized instances, solution times are increased, and for large-sized instances, optimality gaps are decreased. Finally, we obtain better solution times and optimality gaps for all instances when we use the dynamic programming approach to solve our subproblem.

At the end, we make a comparison with the best version of our algorithm and MIP formulations. For small and medium-sized instances, $GLSP^{NF}$, $GLSP^{CC}$, and our algorithm can find optimum solutions, but our algorithm has better solution times. For large-sized instances, our algorithm can find solutions with lower optimality gaps. We also investigate the impact of the proportion of the setup cost within the total cost on MIP formulations and our algorithm. Results show that when the proportion of the setup cost decreases, the performance of the solution approaches increases. However, our algorithm is affected more by this change, and a much greater improvement in solution times and quality is observed compared to other MIP formulations.

4 Conclusion

We study a single-machine, single-level integrated lot-sizing and scheduling problem with sequence-dependent setup costs. To efficiently solve this problem, we propose a decomposition-based exact solution approach. Since our subproblem has binary integer variables, we use the integer L-shaped-based cutting plane algorithm to solve the decomposed model. We also perform improvements on the algorithm by lifting the cuts, calculating better lower bounds, disaggregating the cuts, and developing a dynamic programming-based approach to solve the subproblem. We use a well-known benchmark dataset to test the performance of the algorithm. The results show that improvements significantly affect the performance of the algorithm, and the best version of our algorithm outperforms other exact solution approaches in the literature. Effective heuristics can be used to tackle more complex cases as a future research area. Adding stochastic demand structures to the problem is an additional approach that could be taken.

References

[1] Clark, A. R. and Clark, S. J. Rolling-horizon lot-sizing when set-up times are sequence-dependent. International Journal of Production Research, 38(10):2287–2307 (2000).

[2] Copil, K., Wörbelauer, M., Meyr, H., and Tempelmeier, H. Simultaneous lotsizing and scheduling problems: a classification and review of models. OR Spectrum, 39(1):1–64 (2017).

[3] Drexl, A. and Kimms, A. Lot sizing and scheduling — survey and extensions. European Journal of Operational Research, 99(2):221–235 (1997).

[4] Fleischmann, B. and Meyr, H. The general lotsizing and scheduling problem. Operations-Research-Spektrum, 19(1):11–21 (1997).

[5] Guimaraes, L., Klabjan, D., and Almada-Lobo, B. Modeling lotsizing and scheduling problems with sequence dependent setups. European Journal of Operational Research, 239(3):644–662 (2014).

[6] Işık, E. E., Taşkın, Z. C., and Ağralı, S. Github repository for 'A Novel Decomposition-Based Exact Solution Approach for Lot-sizing and Scheduling Problem' (2025).

[7] James, R. J. and Almada-Lobo, B. Single and parallel machine capacitated lotsizing and scheduling: New iterative mip-based neighborhood search heuristics. Computers & Operations Research, 38(12):1816–1825 (2011).

[8] Kwak, I.-S. and Jeong, I.-J. A hierarchical approach for the capacitated lot-sizing and scheduling problem with a special structure of sequence-dependent setups. International Journal of Production Research, 49(24):7425–7439 (2011).

[9] Laporte, G. and Louveaux, F. V. The integer l-shaped method for stochastic integer programs with complete recourse. Operations Research Letters, 13(3):133–142 (1993).

[10] Wagelmans, A., van Hoesel, S., and Kolen, A. Economic lot sizing: An o(n log n) algorithm that runs in linear time in the wagner-whitin case. Operations Research, 40(1-supplement-1):S145–S156 (1992).

Multi-item lot sizing with fixed sequences on parallel machines

Ben Maktouf Rym
Mines Saint-Étienne, Univ. Clermont Auvergne, CNRS, UMR 6158 LIMOS
DecisionBrain
rym.benmaktouf@emse.fr, rym.ben-maktouf@decisionbrain.com

Dauzère-Pérès Stéphane
Mines Saint-Étienne, Univ. Clermont Auvergne, CNRS, UMR 6158 LIMOS
dauzere-peres@emse.fr

Godard Daniel, Lerouge Mathieu
DecisionBrain
daniel.godard@decisionbrain.com, mathieu.lerouge@decisionbrain.com

Kedad-Sidhoum Safia
Conservatoire National des Arts et Métiers
safia.kedad_sidhoum@cnam.fr

1 Problem description and related works

The usual solution approach to solve production and scheduling problems involves two sequential optimization problems. First, a lot sizing problem is solved to determine a production plan; for each time period of a planning horizon and for each product, a production quantity (a lot) is decided. The objective is to satisfy the demands occurring over the horizon while optimizing the inventory levels, i.e., levels which comply with both safety stocks and maximum inventory levels. Then, a detailed scheduling problem is solved where lots, often called jobs, are scheduled on a set of parallel machines to minimize one or several criteria. Solving sequentially lot sizing and scheduling problems often results in inconsistencies: In the final schedule, jobs may not be scheduled within the time period assigned to the corresponding lots in the lot sizing problem. The demands are satisfied at due dates occurring during the time horizon which results in sub-optimal inventory levels. One way to limit these inconsistencies is to define and solve integrated lot sizing and scheduling problems.

Integrated lot sizing and scheduling problems have been reviewed and classified in [2], where a generalization of simultaneous lot sizing and scheduling is introduced. An integrated lot sizing and job-shop scheduling problem is addressed in [3] and [4], using approaches that solve a lot sizing problem for a fixed sequence of operations on the machines. In contrast to the literature, which relies on discrete time periods, we consider in the present work a continuous time horizon.

This work studies a multi-item capacitated lot sizing problem with fixed sequences on parallel machines. As an input to this problem, we assume that we have a reference schedule, obtained for instance by applying the sequential lot sizing and scheduling approach mentioned above. This reference schedule is used to derive information for our problem, namely the sequences of jobs on each machine and for each item and the initial production quantities (lot sizes). The aim is to schedule the jobs, i.e decide the start and completion times of each job as well as their associated production quantities in order to minimize the inventory costs, which include the violation cost of the maximum inventory levels, the backlog cost and the violation cost of the safety stocks. A first version of this work has been presented in [1], which is extended here by precisely evaluating the inventory levels and developing an iterative solution approach. More specifically, for each kind of inventory cost, the total costis the sum of the costs related to the product inventory levels measured at different times: The start and completion times of jobs as well as the end of the scheduling horizon.

This paper describes the mathematical model of the Capacitated Lot Sizing Problem with Fixed Sequences, denoted (CLSPFS), and presents a solution approach and numerical results for industrial instances.

2 Modeling the Capacitated Lot Sizing Problem with Fixed Sequences (CSLPFS)

The (CLSPFS) considers a continuous planning horizon, where a set of products (items) $\mathcal{I}$ are produced on a set of parallel identical machines $\mathcal{M}$. Each product (item) is associated to one or several jobs in a set of jobs $\mathcal{J}$. Each product $i \in \mathcal{I}$ is characterized by a maximum inventory level, a safety stock and backlog costs. In addition, for each product, demand quantities must be satisfied at various due dates in the planning horizon. From such a demand distribution, we define the accumulated demand $D_i(t)$ of product i at time t, which is a piecewise constant function of time. Each job $j \in \mathcal{J}$ is assigned to a machine, based on the reference schedule. Job j has a unitary processing time p_j and requires a setup time $s_{j'j}$ which depends on the job j' performed right before j.

Based on the reference schedule, we can define, for each machine $m \in \mathcal{M}$, $\mathcal{S}_m^M$ the fixed sequence of jobs on machine m, i.e., the order in which the jobs assigned to machine m are executed. We can also use the reference schedule to define, for each product $i \in \mathcal{I}$, $\mathcal{S}_i^P$ the fixed sequence of jobs relative to product i, which specifies the order in which the jobs of i are completed on the different machines.

The completion time of each job j is associated with variable C_j and the production quantity of j with variable X_j. Given the fixed sequences of jobs per product, we can also define the accumulated production variable Q_j for each job j. Finally, we introduce inventory variables which are used in the objective function to evaluate the inventory costs. Specifically, for each product $i \in \mathcal{I}$, we consider three inventory levels: (1) At the start times of the jobs producing i $(I_{ij}^{start})_j$, (2) At their completion times $(I_{ij}^{end})_j$, and (3) At the end of the horizon I_{if}. Ideally, we would rather consider the inventory levels at the demand due dates rather than at the start times of the jobs, as demand due dates correspond to inventory level changes. However, modeling these levels adds too much complexity to the problem. This is why they are estimated by considering inventories at the start time of jobs.

The constraints of the (CLSPFS) include precedence constraints for the fixed sequences per machine and per product, and the modeling of the accumulated production quantities and accumulated demands. In addition, Constraints (1), (2) and (3) define the inventory levels of job j of product i using the initial inventory I_{i0}.

$$I_{ij}^{end} = I_{i0} + Q_j - D_i(C_j) \qquad \forall i \in \mathcal{I}, \forall j \in \mathcal{J}_i^P \tag{1}$$

$$I_{ij}^{start} = I_{i0} + Q_j - X_j - D_i(C_j - p_j X_j) \qquad \forall i \in \mathcal{I}, \forall j \in \mathcal{J}_i^P \tag{2}$$

$$I_{if} = I_{i0} + Q_{u_i} - D_i(h) \qquad \forall i \in \mathcal{I} \tag{3}$$

Finally, constraints limiting the makespan of each machine to its value in the reference schedule are added. These constraints ensure that schedules with unnecessary long makespans and machine idle times are not determined.

Given the industrial priorities of the inventory violation costs and the goal of reducing machine idleness, a lexicographical objective function has been implemented with the following priorities: (1) The maximum inventory violation cost is minimized, (2) The backlog cost is minimized, (3) The safety stock violation cost is minimized, and (4)The total idle time of machines is minimized.

3 Solution approach and numerical results

The (CLSPFS) is implemented as part of an iterative solution approach summarized in Figure 1. Given a reference schedule, a preprocessing step extracts the fixed sequences and the reference makespans. Then, the model is solved and a new schedule is obtained. If this schedule includes lot sizes that are equal to one unit, the corresponding jobs are deleted. The process is repeated, i.e., the (CLSPFS) is solved, until no jobs are deleted from the schedule.

The proposed solution approach has been applied on 20 industrial instances varying in complexity and size. The quality of the obtained schedules is evaluated using two different cost functions: (1) An "**optimized**" cost obtained by measuring the inventory levels at the start and completion times of jobs as well as at the end of the horizon, and (2) A "**post-processing**" cost obtained by measuring the quantities in

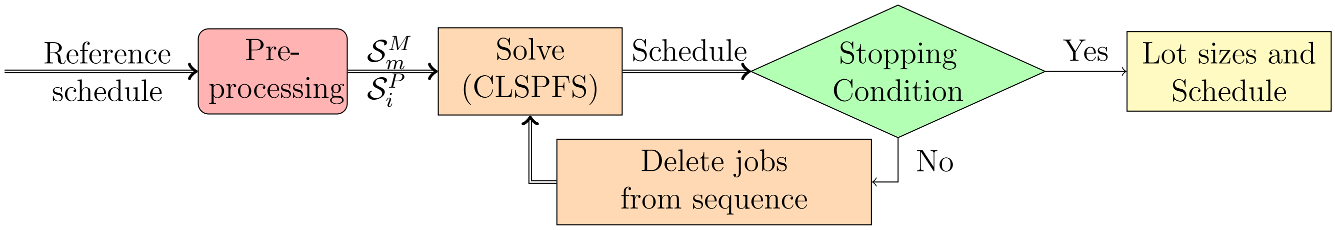

Figure 1: Iterative solution approach

stock at at the completion times of jobs and at the due dates of demands as well as at the end of the horizon.

The numerical results in Table 1 show that there is a significant gap between the "optimized" inventory costs and the "post-processing" inventory costs, in particular the backlog and safety stock violation costs. Note that the first iteration of the solution approach improves all costs except the backlog costs. However, at the final iteration, all costs are reduced. Also, the production quantities increase when optimizing the inventory levels.

	Cost type	**Reference schedule**	**First iteration**	**Final iteration**
Maximum inventory violation cost (Prio. 1)	Optimized	22,997,162	15,799,411 (-31.3%)	15,815,7031 (-31.2%)
	Post Processing	35,506,642	18,464,980 (-48.0%)	18,525,103 (-47.8%)
Backlog cost (Prio. 2)	Optimized	4,738	6,060 (27.9%)	4,579 (-3.4%)
	Post Processing	47,636	52,161 (9.5%)	42,445 (-10.9%)
Safety stock violation cost (Prio. 3)	Optimized	27,598,401	23,844,002 (-13.6%)	22,982,766 (-16.7%)
	Post Processing	81,566,163	78,998,596 (-3.2%)	75,844,428 (-7.0%)
Average idle time (min) (Prio. 4)		11,846	8,831 (-25.5%)	6,930 (-41.5%)
Average machine makespan (min)		29,792	28,542 (-4.2%)	28,459 (-4.5%)
Production quantity		154,443	169,060 (9.5%)	171,995 (11.4%)
Deletable jobs		0%	11%	13%

Table 1: Numerical results

References

[1] Rym Ben Maktouf, Stéphane Dauzère-Pérès, Daniel Godard, and Safia Kedad Sidhoum. Lot sizing and scheduling with fixed sequences and setup times. In *International Workshop on Lot-Sizing 2024*, 2024.

[2] Karina Copil, Martin Wörbelauer, Herbert Meyr, and Horst Tempelmeier. Simultaneous lotsizing and scheduling problems: a classification and review of models. *OR spectrum*, 39:1–64, 2017.

[3] Stéphane Dauzère-Pérès and Jean-Bernard Lasserre. Integration of lotsizing and scheduling decisions in a job-shop. *European Journal of Operational Research*, 75(2):413–426, 1994.

[4] Edwin David Gómez Urrutia, Riad Aggoune, and Stéphane Dauzère-Pérès. Solving the integrated lot-sizing and job-shop scheduling problem. *International Journal of Production Research*, 52(17):5236–5254, 2014.

Remanufacturing and refurbishment of pre-owned consumer electronic products under uncertainty

Terzi Mourad, Absi Nabil
Mines Saint-Etienne, Univ. Clermont Auvergne, CNRS, UMR 6158 LIMOS, CMP, Department of Manufacturing Sciences and Logistics, Gardanne, France
{mourad.terzi,absi}@emse.fr

Schepler Xavier,
Generix Group, 20 Pl. de la Défense, Puteaux, France
xavier.schepler@gmail.com

Jeanjean Antoine
OPT2A – Optimization Augmented Analytics, Bordeaux, France
antoine.jeanjean@recommerce.com

Abstract

This study addresses the tactical planning problem for reverse flows of pre-owned consumer electronics under uncertain returns. Motivated by industrial practice at Recommerce Group, the system combines product heterogeneity, a two-stage refurbishment process with batch remanufacturing, FIFO rules, and multi-class, multi-quality demand. The goal is to develop a planning model that captures the stochastic quantity-quality mix of incoming products while meeting demand expressed by product classes. The problem is formulated as a two-stage stochastic mixed-integer linear program: first-stage decisions route returns to refurbishment centers and launch remanufacturing batches, whereas second-stage recourse flows satisfy demand across scenarios drawn from probability distributions of return qualities. Sensitivity analysis highlights the joint impact of batch size and demand granularity on profitability under return variability.

1 Introduction

In the shift toward a circular economy, managing end-of-life product logistics presents significant c hallenges d ue t o t he c omplexity o f o perations l ike c ollection, storage, repair, and redistribution. This study builds on the earlier work of Schepler et al. [1], who introduced a planning model for managing flows and resources involved in the collection, processing, refurbishment, and remanufacturing of used electronic products. Inspired by industrial Recommerce practices, the model accounts for several features: the diversity of collected product types, processing in refurbishment centers, the distinction between refurbishment and remanufacturing processes, and demand expressed by product classes and quality levels. The model is formulated as a mixed-integer linear program and solved using a relax-and-fix heuristic.

The main limitation in the original formulation lies in its deterministic nature, whereas in practice, the quality of collected products is subject to uncertainty, driven by external factors such as the duration of the prior use. In this work, the initial model is extended by incorporating the products' stochastic quality. The objective is to develop a stochastic planning model that accounts for variability in incoming flows while ensuring demand by product class is satisfied. In the following section, the stochastic planning problem for trade-in, refurbishment, and remanufacturing of used electronics, hereafter referred to as *SRRP*, is described.

2 *SRPP* description

Used consumer electronics are collected through trade-in channels such as retail stores. These products are grouped into types (e.g., iPhone 7 32GB, various colors) and vary in quality, defined by their functional and cosmetic condition. Devices like smartphones and tablets are classified into categories such as 'Functional – Like new', 'Damaged', or 'To recycle'. Other collection channels include international sourcing and B2B buyback, which allow ordering lots of products from companies. Compared to trade-in, these activities are secondary for Recommerce Group. While their flows can be represented in the model, they fall outside the scope of this study.

Collected products are first consolidated in stores, then sent to refurbishment centers where they are quickly inspected, cleaned, tested, and possibly lightly repaired. Processing follows the First In, First Out (FIFO) rule. Some products are later sent in batches to remanufacturing centers, where subcontractors impose batch size constraints. The remanufacturing process, including transport, takes from a few to several weeks.

Demand is specified by product classes defined by their quality, with each class grouping products sharing similar features (e.g., all variants of the iPhone 7). Products are sold through various channels. Fully functional, good-quality items are sold individually to consumers or retailers, while others are sold in batches. Some products are directed to refurbishers, whereas broken or low-demand items are sold to remanufacturers or international brokers for parts recovery or resale in international markets.

The objective is to maximize total profit, defined as the difference between sales revenue and overall costs. Achieving this requires efficient logistics to minimize delays in collection, processing, and resale.

3 *SRRP* formulation

This section presents a general formulation of the *SRRP*. It specifically details the notations, including decision variables and parameters, outlines the objective and main constraints, and concludes with the formulation addressing the stochastic quality of collected products.

3.1 Notations

The following sets are used to define the initial model: (i) $\mathcal{I} = \{1, \ldots, I\}$: set of collection channels, (ii) $\mathcal{M} = \{1, \ldots, N\}$: set of product types, (iii) $\mathcal{R} = \{1, \ldots, R\}$: set of refurbishment centers, (iv) $\mathcal{T} = \{1, \ldots, T\}$: discrete time horizon, and (v) $\mathcal{J} = \{1, \ldots, J\}$: set of product qualities.

The decision variables of the *SRRP* include: (i) $X_{mirt} \in \mathbb{R}_+$: quantity of products of type m collected at channel i during period t transferred to test center r, (ii) $Y_{mirtt'} \in \mathbb{R}_+$: quantity of m arrived during t collected at channel i and tested during $t^{'}$ in r, (iii) $I^{-}_{mirtt'} \in \mathbb{R}_+$: inventory level of untested m collected at channel i during t in r at the end of $t^{'}$, (iv) $W_{mjt} \in \mathbb{R}_+$: quantity of m in quality j sent to remanufacturing at t, (v) $B_{rtt'} \in \{0, 1\}$: 1 if FIFO rule allows testing during $t^{'}$ in r products arrived during t, 0 otherwise, (vi) $V_t \in \{0, 1\}$: 1 if a remanufacturing batch is sent during t, 0 otherwise, (vii) $\tilde{I}^{+}_{mjt} \in \mathbb{R}_+$: inventory level of tested m in quality j at the end of t and, (viii) $\tilde{Z}_{mjt} \in \mathbb{R}_+$: quantity of m in quality j sold during t. Note that the last two variables are denoted with a tilde to indicate their dependence on the realization of the stochastic quality of the products.

3.2 *SRRP* formulation: objective and constraints

The objective of the *SRRP* is to maximize the total profit while satisfying several types of constraints. These include flow conservation across different stages such as collection, pre-processing storage, post-processing storage, remanufacturing, and resale. Additionally, the model accounts for the capacity limitations of test centers, enforces batch remanufacturing operations, and ensures that product processing follows a First-In-First-Out (FIFO) order.

The problem is formulated as a two-stage stochastic mixed-integer linear program to account for uncertainty in product quality:

- First-stage decisions (before the realization of product quality): includes variables X_{mirt}, $Y_{mirtt'}$, $I^{-}_{mirtt'}$, W_{mjt}, $B_{rtt'}$ and V_t.

- Second-stage decisions (after the realization of product quality): includes variables $\tilde{I}^+_{mjt}$ and $\tilde{Z}_{mjt}$.

3.3 Stochastic quality of the collected products

As stated earlier, this study considers the quality of collected products as a stochastic parameter. This assumption is motivated by practical considerations. In fact, while an initial quality assessment is typically provided at the time of collection, based on visual inspection or customer-declared condition, the actual quality is only confirmed after testing is performed at refurbishment centers. This discrepancy between declared and actual quality introduces uncertainty regarding the distribution of products across quality levels. In this context, the parameter $\tilde{\tau}_{mijt}$ represents the proportion of products of type $m \in M$ and quality level $j \in J$ that are collected through channel $i \in I$ during period $t \in T$. As the true quality of products is only revealed after refurbishment, $\tilde{\tau}_{mijt}$ is modeled as a random variable, capturing the variability and unpredictability in product quality outcomes.

4 Conclusion

This study addresses the tactical planning problem for reverse flows of pre-owned consumer electronics under uncertain return qualities. It extends the work of Schepler et al. [1], and is inspired by real-world industrial practices at Recommerce Group. A key contribution lies in the integration of stochastic product quality, motivated by the practical observation that true quality is only revealed after diagnostic testing at refurbishment centers. A scenario-based modeling approach is used to capture this uncertainty. The problem is formulated as a two-stage stochastic mixed-integer linear program: first-stage decisions involve routing returns to refurbishment centers and launching remanufacturing batches, while second-stage recourse flows satisfy demand across scenarios reflecting return quality variability. As a next step, the focus will be on solving the stochastic model using advanced solution techniques such as Benders decomposition.

References

[1] Schepler, Xavier, Nabil Absi, and Antoine Jeanjean. Refurbishment and remanufacturing planning model for pre-owned consumer electronics. International Journal of Production Research 62.7, 2499-2521 (2024).

Dynamic lot sizing model with remanufacturing and separate setup costs: Time complexity and optimality

First Author: Chee-Khian Sim
Affiliation: University of Portsmouth
email: chee-khian.sim@port.ac.uk

Abstract

We introduce the dynamic lot sizing model with remanufacturing and separate setup cost for manufacturing and remanufacturing. The optimal policy for the model when there is joint setup cost for manufacturing and remanufacturing can be obtained in polynomial time [1], while when the setup cost for manufacturing and remanufacturing is separate, obtaining the optimal policy efficiently is considerably more difficult. We polynomially reduced an instance of an NP-complete problem, the Partition Problem, to an instance of our model with separate setup cost for manufacturing and remanufacturing. This indicates that inherently, finding an optimal policy for our model is computationally difficult. By formulating our model as a dynamic program, we show that its optimal policy can be found with pseudo polynomial time complexity. We then propose a class of feasible policies for our model that satisfies two properties, one being the zero-inventory property, and show that an optimal policy in this class of policies can be found with polynomial time complexity. This is shown through the dynamic program formulation to find the policy. Next, we investigate the closeness of this policy to the optimal policy in terms of total system cost and show theoretically that the policy can be close to optimal. We implement our dynamic programs to find the optimal policy and the feasible policy, and present numerical results comparing the feasible policy with the optimal policy on instances of the model.

Results presented are based on the submitted work [2].

References

[1] R.H. Teunter, Z.P. Bayinder & W. van den Heuvel, *Dynamic lot sizing with product returns and remanufacturing*, International Journal of Production Research, Vol. 44, pp. 4377-4400 (2006).

[2] C.-K. Sim, *Time complexity and optimality of inventory and production policies for a dynamic lot sizing model with remanufacturing and separable setup costs*, submitted, 2025. https://optimization-online.org/?p=29925

Solving a Cutting and Planning Problem with Reusable Leftovers

Léo Gilbert
Reeverse Systems, 44000 Nantes, France
leo.gilbert@reeverse-systems.fr

María I. Restrepo
IMT Atlantique, LS2N, UMR CNRS 6004, 44000 Nantes, France
maria-isabel.restrepo-ruiz@imt-atlantique.fr

Nadjib Brahimi
Rennes School of Business, 35200 Rennes, France
nadjib.brahimi@rennes-sb.com

François Klein
Reeverse Systems, 44000 Nantes, France
francois.klein@reeverse-systems.fr

Abstract

The cutting process is present across many industrial sectors, yet remains poorly optimized. Today, approximately 30% of raw material is discarded. This waste, with both ecological and financial consequences, motivates our work: to reduce unused material in industry while maintaining industrial viability. To achieve this, we combine two classical operations research problems: the *Cutting Stock Problem* (CSP) and the *Lot Sizing Problem* (LSP).

1 Industrial and Scientific Context

Cutting operations are widespread across various industrial sectors, yet they often remain under-optimized. As a result, nearly 30% of raw material is typically discarded, leading to significant environmental and financial repercussions. Driven by these societal and operational challenges, this work aims to reduce industrial material waste while ensuring practical feasibility. To this end, we address a problem that integrates two classical operations research formulations: the *Cutting Stock Problem* (CSP) and the *Lot Sizing Problem* (LSP).

This research is conducted within a CIFRE doctoral framework in collaboration with Reeverse Systems, a company specializing in software solutions for raw material valorization. Our study is rooted in concrete industrial use cases provided by Reeverse's clients.

In recent years, the classical CSP has evolved to include the notion of *usable leftovers*, formalized as the *Cutting Stock Problem with Usable Leftovers* (CSP-UL) [1]. These models shift the focus from minimizing total waste to minimizing non-reusable material, as outlined in [2]. In this context, some remnants from cutting operations are reclassified as usable objects and can be reused in current or future production.

We focus particularly on scenarios where cutting decisions are embedded within a broader production planning framework. This gives rise to the *Integrated Lot Sizing and Cutting Stock Problem* (LSP-CSP), surveyed comprehensively in [3]. Solving these problems jointly rather than sequentially better reflects industrial constraints and yields solutions that are both feasible and cost-effective.

Further, combining both extensions—namely usable leftovers and integrated planning—leads to the LSP-CSP-UL model. While not the first attempt at such integration, the formulation proposed in [4] marks a significant step toward a more general framework. However, this approach still exhibits limitations in both modeling and computational performance.

2 Modelling

We aim to extend this modeling framework by incorporating setup costs and smoothing constraints for production planning. Moreover, the model is tailored to industrially relevant scenarios: a 2D cutting stock problem involving irregular item shapes with free rotation.

The proposed model seeks to determine which cutting patterns to apply on which objects $j \in J$, across a discrete planning horizon T, to meet known demand. Specifically, we decide the number $x_{t,j,k}$ of pattern instances $k \in K$ to execute on object j at time t, assuming the set of feasible patterns K is known in advance.

The MILP is composed of an objective function and of several constraints.

The objective function consists of the sum of raw material costs, storage costs (during the planning horizon and long-term) for both items and objects, setup costs, and production smoothing penalties. The function intends to guide us toward feasible solutions that minimize both production and raw material costs, not to perfectly represent real-world costs.

A first set of constraints represent the inventory level of items at each period in the planning horizon. These constraints guarantee that the demand is met, in different ways.

- The items can already be in stock.
- The items can be produced internally. The number of manufactured items are calculated as the product of the number of cuts and the number of items produced per cut, summed over all possible cuts at this time that produce the given item.
- The items can also be "procured externally" i.e., purchased, when not produced internally by the company. This procurement is a notional representation, as it may simply account for a production shortage.

Stock management constraints apply to both items and raw material units (new or reused). The core aspect of the LSP-CSP-UL model lies in the dynamic inventory of usable leftovers: each cutting operation may produce new reusable objects, which in turn can be consumed in subsequent operations.

Thus, a second set of constraints models the stock level of objects, in a way specific to the LSP-CSP-UL framework. These raw materials are transformed into finished products. It governs the reintroduction in stock of usable leftover, generated during the cutting process.

We also integrate setup costs to generate industry-relevant production schedules. The objective is to balance fixed costs (setup) against variable costs (raw materials, storage, and operations).

The setup constraints on the cutting pattern and the objects determine whether the setup variables are active. These constraints govern the behaviour of the setup variables, which must activate when a type of object or cutting pattern is used at a given time t. Thus, whenever one of the x_k variables becomes active, the corresponding setup variable is also activated. Notice that this constraint weakens the linear programming relaxation of the model.

3 Solution Methods

To solve realistic problem instances, we developed two solution strategies. The first method operates on a pre-defined library of industrial nesting patterns. This allows us to assess the benefits of process optimization with minimal operational change. We employ basic matheuristics to reduce computation time, enabling fast feedback loops with industrial partners to refine the model and align with practical needs.

The second strategy aims to go beyond existing patterns by generating new ones via column generation. This is computationally challenging, especially in the context of our 2D variant of the CSP. Reaching convergence often requires thousands of patterns, and generating a high-quality, feasible pattern can take several seconds using state-of-the-art nesting software[1].

4 Preliminary Results

The first solution approach has already been implemented in an industrial environment. It yields approximately a 4% reduction in raw material usage, with computation times under a minute for a one-month planning horizon involving 300 patterns, 50 items, and a similar number of objects. We expect to double this gain using the second approach, albeit at a significantly higher computational cost.

5 Introduction

References

[1] Cherri, A., Kantor, D., Oliveira, J. F. (2014). Cutting stock problem with usable leftovers: literature review. *International Journal of Production Research*, 52(2), 395–411.

[2] Senergues, C. (2023). A State of the Art on the Cutting Stock Problem with Usable Leftovers. *HAL*, hal-04010536.

[3] Melega, G. M., Morabito, R., & Lodi, A. (2018). Classification and literature review of integrated lot-sizing and cutting stock problems. *European Journal of Operational Research*, 264(1), 1–20.

[4] Do Nascimento, G. M., et al. (2022). Integrated Lot Sizing and Cutting Stock Problem with Usable Leftovers: A Comprehensive Formulation and Solution Approach. *Computers & Operations Research*, 139, 105653.

[1] Alma: 2D nesting software

Robust Production Planning Under Inherent Supply Chain Disruptions

Amin Javanmard
Dept. of Management Science, University of Strathclyde, Glasgow, UK
amin.javanmard@strath.ac.uk

Kerem Akartunalı
Dept. of Management Science, University of Strathclyde, Glasgow, UK
kerem.akartunali@strath.ac.uk

Ashwin Arulselvan
Dept. of Management Science, University of Strathclyde, Glasgow, UK
ashwin.arulselvan@strath.ac.uk

Abstract

This study explores a robust optimization approach to capacitated lotsizing under inherent supply chain disruptions, including bursty demand patterns and uncertain transition dynamics. An iterative adversarial framework assesses worst-case scenarios, enhancing production planning resilience. Various uncertainty representations are considered to model fluctuations and transitions. The framework integrates a structured demand process and optimizes decisions under uncertain conditions. Future research will refine uncertainty modelling, incorporate advanced forecasting techniques, and improve computational efficiency. This work contributes to resilient supply chain management by supporting adaptability in unpredictable environments.

1 Introduction

Supply chain disruptions are an unavoidable challenge in production planning. Unlike external disasters, inherent disruptions arise from within the supply chain itself, such as bursty demand patterns and unpredictable variations. These uncertainties necessitate a robust decision-making framework that ensures operational resilience while optimising costs.

We define supply chain resilience as the ability of the supply chain to plan for, respond to and recover from disruptions in a timely and cost-effective manner [3].

There are various proactive and reactive resilience strategies [2], such as contracting with back-up suppliers, using safety stocks, or employing flexible manufacturing approaches, among others.

We explore different approaches to handling inherent supply chain disruptions. We develop a novel state-based robust optimisation model designed to create production plans that are resilient to plausible, high-impact adverse scenarios. A key objective is to compare this model with a baseline approach that does not explicitly account for inherent disruptions. We provide brief remarks for the problem on hand next, and discuss further details including preliminary results in the talk.

2 Deterministic Lot-Sizing Model

To benchmark our approach, we first consider a baseline deterministic model with backlogging and constant capacities (Formulation 1):

$$
\begin{aligned}
\text{Minimize} \quad & \sum_{t=1}^{T} pX_t + \sum_{t=1}^{T} qY_t + \sum_{t=1}^{T} hS_t + \sum_{t=1}^{T} bR_t && [1], \\
\text{Subject to:} \quad & S_{t-1} - R_{t-1} + X_t = D_t + S_t - R_t && \forall t \;\; [1a], \\
& X_t \leq CY_t && \forall t \;\; [1b], \\
& X_t, R_t, S_t \in \mathbb{R}_+^N && \forall t \;\; [1c], \\
& Y_t \in \{0,1\}^n && \forall t \;\; [1d].
\end{aligned}
$$

The objective [1] minimizes production, setup, inventory, and backlog costs. Constraints [1a] ensure flow balance, [1b] enforce capacity limitations, and [1c] and [1d] impose nonnegativity and integrality. Introducing a new variable H_t to represent the overall holding and backlog costs simplifies the structure (Formulation 2):

$$
\begin{aligned}
\text{Minimize} \quad & \sum_{t=1}^{T} pX_t + \sum_{t=1}^{T} qY_t + \sum_{t=1}^{T} H_t && [2], \\
\text{Subject to:} \quad & h\sum_{i=1}^{t} (X_i - D_i) \leq H_t && \forall t \;\; [2a], \\
& -b\sum_{i=1}^{t} (X_i - D_i) \leq H_t && \forall t \;\; [2b], \\
& X_t \leq CY_t && \forall t \;\; [2c], \\
& X_t, H_t \in \mathbb{R}_+^N && \forall t \;\; [2d], \\
& Y_t \in \{0,1\}^n && \forall t \;\; [2e].
\end{aligned}
$$

Here, constraints [2a] and [2b] replace the flow balance constraints [1a], ensuring H_t captures both holding and backlog costs. This reformulation reduces variables and constraints, enhancing efficiency but at the cost of direct inventory and backlog interpretation.

3 Robust Lot-Sizing Model

To handle uncertainty, we introduce a two-stage robust optimization framework that iteratively solves two interconnected problems, a Master Problem (MP) and an Adversarial Problem (AP), broadly in a similar fashion as in [1].

3.1 Master Problem (MP)

The MP (Formulation 3) determines a single, robust production plan that performs well against the set S of known challenging scenarios (d^s, F^s) identified so far. Specifically, it minimizes Θ, which represents the maximum possible total cost calculated across all scenarios $s \in S$.

$$
\begin{aligned}
\text{Minimize} \quad & \Theta && && [3],\\
\text{Subject to:} \quad & \Theta \geq \sum_{t=1}^{T}(p_t X_t + q_t Y_t) + \sum_{t=1}^{T} H_t^s && \forall s \in S && [3a],\\
& Inv_t^s = \sum_{i=1}^{t} X_i - \sum_{i=1}^{t}\sum_{l=1}^{N} F_{il}^s d_{il}^s && \forall s \in S, \forall t && [3b],\\
& H_t^s \geq h.Inv_t^s && \forall s \in S, \forall t && [3c],\\
& H_t^s \geq -b.Inv_t^s && \forall s \in S, \forall t && [3d],\\
& X_t \leq CY_t && \forall t && [3e],\\
& X_t \geq 0, H_t^s \geq 0, \Theta \geq 0 && \forall s \in S, \forall t && [3f],\\
& Y_t \in \{0,1\} && \forall t && [3g].
\end{aligned}
$$

Constraint [3a] links the total cost of each scenario to the overall objective Θ. In [3b], Inv_t^s represents the net inventory at the end of period t under scenario s, calculated based on the production plan X and the specific demand sequence (d_s, F_s) associated with scenario s. Constraints [3c] and [3d] ensure H_t^s captures the maximum of the holding and backlogging costs. By considering all known worst cases $s \in S$ simultaneously, the MP seeks a production plan X resilient to a range of potential futures and provides a reliable lower bound ($LB = \Theta$) on the true optimal robust cost.

3.2 Adversarial Problem (AP)

Given a candidate production plan X from the MP, the AP seeks to find the single most challenging scenario that maximises the resulting sum of holding and linear backlog costs by choosing:

1. A specific state path, F, which is a sequence of active states over the time horizon.

2. The realized transition probabilities, $\tilde{P_{ll'}}$, from within their defined uncertainty set.

3. The realized demand values, $\tilde{d_{tl}}$, for each state and period from within their uncertainty set.

The crucial constraint is that the chosen path F must be plausible, i.e., its total probability must be greater than or equal to a threshold α. The worst-case scenario found by the AP is then added to the set of scenarios considered by the Master Problem in the next iteration. This iterative process refines the production plan until the gap between the lower and upper bounds converges.

We will discuss further technical details of AP (including its full formulation) in the talk, along with some observations and insights from our preliminary tests. We will also explore several research perspectives going forward.

References

[1] D. Bienstock and N. Özbay. Computing robust basestock levels. Discrete Optimization, 5(2):389–414, 2008.

[2] A. Mohammed, N. Zubairu, M. Yazdani, A. Diabat, and X. Li. Resilient supply chain network design without lagging sustainability responsibilities. Applied Soft Computing, 140:110225, 2023.

[3] Y. Sheffi and J.B. Rice Jr. A supply chain view of the resilient enterprise. MIT Sloan Management Review, 2005.

Capacitated lot-sizing problem with stock-out based substitution

David Tremblet
HEC Montréal
david.tremblet@hec.ca

Raf Jans
HEC Montréal
raf.jans@hec.ca

Yossiri Adulyasak
HEC Montréal
yossiri.adulyasak@hec.ca

Sonja Ursula Katharina Rohmer
HEC Montréal
sonja-ursula-katharina.rohmer@hec.ca

Abstract

We consider an integrated lot-sizing and assortment planning problem with stock-out based substitution. In each period, an assortment of items is proposed to a set of customers. Each customer has a preference list of items and chooses its preferred item among the set of available ones. We suppose that the market size and the rates at which the items are depleted are known at each period, but the customer types and the order in which each customer arrives at the shop is unknown. Each item is sold at a certain price, and our objective is to maximize the profit of each item with respect to the standard costs in lot-sizing, including production, setup, and holding costs. We propose a MILP that solves the deterministic case where each customer type arrives homogeneously in the store. We also investigate the integration of backlog in the model, corresponding to the case where some customer type has a probability of coming back to the store in the next period.

1 Introduction

Existing lot-sizing models allow some items to be substituted to satisfy the demand. However, these models assume that the retailer makes the substitution choices to offer to customers, which represents an optimistic scenario. Such substitutions usually involve upgrades, i.e. replacing a lower-quality product with a higher-quality one [4]. In practice, demand in retail stores is strongly influenced by the assortment of items proposed in the store, and customers only consider substituting their preferred product when it is out of stock. This type of substitution is known as stock-out-based substitution [3], i.e., customers base their choice on the assortment available at the time they visit the store. To address this, some studies incorporate customer choice behavior in response to varying product assortments [1, 2]. However, these approaches often rely on logit-based models to represent the demand, which leads to nonlinear formulations that are difficult to solve.

We investigate a lot-sizing problem with customer-driven demand substitution, where the substituted items are chosen by the customer based on the item availability. The most closely related work is by [5], which considered a single-period assortment planning problem under customer-driven demand substitution and formulated a mixed-integer linear program (MILP) to solve the problem. However, their model is limited to a special scenario of demand substitution and does not fully capture the substitution behavior of individual customer. In contrast, our approach focuses on the customer substitutions in response to stock-outs, offering a more comprehensive approach to substitution.

2 Problem description

We propose a model that maximizes the profit of a multi-item capacitated lot-sizing problem with stock-out based substitution. In this problem, a set of N items is considered and sold to the customers through a dedicated shelf in a store. We assume, without loss of generality, that item 0 corresponds to the no-purchase option. In each period $t \in \{1, \ldots, T\}$, a market size of τ_t customers come to the store. Each customer is defined by a customer type, represented by a sequence of items ordered by preference. A customer visiting the store selects the most preferred product among those available in assortment $\mathcal{A} \subseteq \{1, \ldots, N\}$, or leaves the store without purchasing if they could not find a suitable product. We assume that the customers arrive homogeneously at the store, and each item is depleted with respect to an assortment depletion matrix M, where each row represents the assortment proposed to the customer and each column gives the depletion rate of the corresponding product. Given assortment $\mathcal{A}_i \subseteq \{1, \ldots, N\}$ index by $i \in \left\{1, \ldots, 2^N\right\}$, the customers deplete each product $j \in \mathcal{A}_i$ at rate M_{ij} until a stock-out occurs for one item $k \in \mathcal{A}_i$. The process is then repeated with assortment $\mathcal{A}_i \setminus \{k\}$, and terminates when there is either no customer or no more items in the assortment. The goal of this problem is to find the quantities of items to produce, to stock, and to put on the shelf, while maximizing the difference between the profit of each item and the production, setup, and holding costs. We assume that the substitution cost is negligible.

We provide a small example to better understand the problem. Suppose we have a set of three customer types $C = \{(1,2,3,0), (1,2,0), (3,1,2,0)\}$ with the given arrival probabilities:

$$\mathbb{P}[c = (1,2,3,0)] = 0.45 \qquad \mathbb{P}[c = (2,1,3,0)] = 0.45 \qquad \mathbb{P}[c = (3,1,2,0)] = 0.1$$

For example, for a market size of $\tau_t = 100$ in period $t \in \{1, \ldots, T\}$, a demand of 45 is associated with customer type $(1,2,3,0)$.

Assortment	Depletion rate			
$\mathcal{A}_k$	1	2	3	0
$\{1,2,3\}$	0.45	0.45	0.1	0.0
$\{1,2\}$	0.55	0.45	0.0	0.0
$\{1,3\}$	0.9	0.0	0.1	0.0
$\{2,3\}$	0.0	0.9	0.1	0.0
$\{1\}$	1.0	0.0	0.0	0.0
$\{2\}$	0.0	1.0	0.0	0.0
$\{3\}$	0.0	0.0	0.55	0.45
$\{\}$	0.0	0.0	0.0	1.0

Table 1: Depletion of items for each assortment

Table 1 gives the assortment depletion of each item for the above example. In this table, we observe that item 3 is depleted at a much higher rate only when items 1 and 2 are unavailable. Thus, by only considering the primary demand of the product, item 3 is unlikely to be produced in large quantities. However, producing item 3 may become advantageous under certain conditions, such as limited production capacities for each item.

To solve this problem, we propose an MILP that uses the assortment depletion to compute the substitution behavior of each customer.

3 Preliminary experiments

In this section, we compare our approach, denoted as LS-SBS, with standard capacitated lot-sizing formulations on the instance proposed in Section 2. For this comparison, we consider a capacitated lot-sizing problem, denoted as CLSP, where each item j has a demand equal to the number of customers with item j as their first choice. For each model, the quantities produced at each period are proposed to the customers, and we simulate the depletion of each item by homogeneously depleting the products with respect to the preference of customers.

Model	T = 5	T = 10	T = 20
CLSP	8091.6	21616.5	42470.8
LS-SBS	9226.2	23591.9	46267.5

Table 2: Comparison of the average profit obtained by each model

Table 2 shows the average results of the two models on three instance types with different horizon lengths. The LS-SBS model shows a higher profit compared to the CLSP, which demonstrates the importance of considering the customer choice when incorporating retail decisions into lot-sizing problems.

In future work, we aim to incorporate time-dependent decisions into the model. One promising direction is to account for backlogging in the customer choice process. Specifically, we consider scenarios where some customers may come back to the store in a later period if they could not find a suitable substitute during their initial visit. Another research avenue involves examining customer behavior in the context of perishable products. For example, certain customers might opt not to purchase an item if its expiration date is too close, and we plan to incorporate this preference into future models.

References

[1] Fabian Dunke and Stefan Nickel. A choice-based approach to dynamic capacitated multi-item lot sizing with demand uncertainty. *Applied Mathematical Modelling*, 137:115705, January 2025.

[2] Ahmed Ghoniem and Bacel Maddah. Integrated retail decisions with multiple selling periods and customer segments: Optimization and insights. *Omega*, 55:38–52, September 2015.

[3] Dorothée Honhon, Vishal Gaur, and Sridhar Seshadri. Assortment planning and inventory decisions under stockout-based substitution. *Operations Research*, 58(5):1364–1379, October 2010.

[4] Fuying Jing and Yinping Mu. Dynamic lot-sizing model under perishability, substitution, and limited storage capacity. *Computers & Operations Research*, 122:104978, October 2020.

[5] Eda Yücel, Fikri Karaesmen, F. Sibel Salman, and Metin Türkay. Optimizing product assortment under customer-driven demand substitution. *European Journal of Operational Research*, 199(3):759–768, December 2009.

Dezentralized Lot Sizing: Smart Contracting and Collaboration

Klymenko, Anton
University of Klagenfurt
anton.klymenko@aau.at

Gansterer, Margaretha
University of Klagenfurt
margaretha.gansterer@aau.at

Abstract

Collaborative lot sizing usually relies on a centralized decision-maker to reap the benefits of joint operations. While this often leads to optimal results, concerns about transparency and dependency prevent agents from working together in practice. We aim to explore an alternative approach to dezentralisation in lot sizing using smart contract auctions. We develop and compare two models: a centralized optimizer that achieves the best possible objective value and a dezentralised mechanism, where agents submit self-generated bids. An off-chain combinatorial algorithm determines the optimal allocation, and the results are published on-chain via smart contracts to ensure transparency and auditability. We also investigate the role of smart contracts in contracting for lot sizing. Preliminary results suggest that while the decentralized model worsens the objective value, increased incentives for collaboration can compensate for this trade-off. Based on a computational study, the potential of smart contracts to enable dynamic real-time pricing and enhance transparency is analyzed.

1 Introduction

In today's globalised economy, the efficiency of su pply ch ains is pa ramount. Joint lot sizing – an operational strategy where multiple parties coordinate their inventory management, and production planning to reduce costs – allows for the achievement of the optimal result for everyone [Gansterer et al., 2021, Albrecht, 2010]. The agents are, however, reluctant to share the information that would make such a collaboration possible. The introduction of smart contracts raises the question of their usability in collaborative lot sizing. Smart contracts provide a solid foundation for privacy and fairness by minimising the need for intermediaries in digital transactions. This research investigates the feasibility of using smart contract auctions in decentralized collaborative lot sizing, evaluating whether efficiency can be maintained despite decentralization. We employ a mixed integer linear program (MILP) for the centralized model and develop an iterative bundled-cost hybrid smart contract framework for the decentralized approach. Computational experiments investigate the "cost of decentralization", convergence speed, and transparency outcomes, providing insights for future theoretical advancements.

2 Problem Definition and Experimental Setup

We adopt a two-tiered method in our study. The first layer uses a centralized model to create the best possible theoretical situation. The second layer shows the opposite: a new, privacy-protecting decentralized market mechanism. This method allows to directly assess the "cost of decentralization", which is the loss of efficiency that comes with getting benefits like privacy and trust. We focus on a capacitated lot sizing problem involving several agents and periods. In this case, several production agents need to work together to meet a common outside demand for every period, where each of these agents has its own private information.

Firstly, we create a centralized lot sizing model (C-LS) as benchmark for the proposed decentralized approach. This model depicts the theoretical ideal, wherein a singular, fully informed decision maker exists. The problem is put up as an MILP that minimizes total cost including setup, production, and holding cost, while also considering inventory balance equations and capacity limits. However, it is unlikely that it would be used in real life, as companies are very worried about transparency, dependency, and the possible loss of their competitive edge due to information revelation.

Secondly, we propose a bundle-based incremental cost auction mechanism that allocates demand while preserving agent privacy. The mechanism operates iteratively, with agents submitting bids based on their marginal costs derived from solving local optimisation problems. Instead of allocating demand to each period separately, we group consecutive periods into "bundles" to capture setup cost economies. A neutral auctioneer coordinates the process without accessing agents' private cost data. Each agent's subproblem is a capacitated multi-level lot sizing problem.

The algorithm for the decentralized system operates through three primary phases. First, during bid generation individual agents solve their local lot sizing problems and extract shadow prices from their demand constraints. Next, in the bundle allocation phase a neutral auctioneer allocates each bundle to the agent who has the lowest incremental cost. The entire procedure concludes with the convergence check, which terminates the process as soon as the allocations stabilise. While convergence is not theoretically guaranteed, the finite number of possible allocations and the monotonic improvement property ensure termination. In practice, all tested instances converged within 50 iterations.

A purely on-chain implementation where the entire auction is executed on a public blockchain is currently infeasible as the auctioneer has to solve the Winner Determination Problem (WDP), which is NP-hard. The computational complexity of this task would lead to prohibitive transaction costs and unacceptable execution delays on-chain. This practical limitation necessitates a hybrid architecture that separates computationally intensive tasks from on-chain enforcement.

The workflow employs a multi-phase commit-reveal system to ensure fairness for participants and to prevent issues such as bid-sniping. First, agents confirm their bids function on the blockchain, where a security deposit in an escrow account is kept. This step makes the bid legally binding and private. Once the bidding period expires, agents submit their bid details. The smart contract checks this information against the previous promise. After a successful revelation, the auctioneer picks the winner off-chain by solving the WDP. After that, results are published and final prices are send for clearing and allocating to the blockchain, where they will reside. Furthermore, there's a function to handle money transactions automatically, distributing possible winnings and returning deposits to real agents.

3 Computational Results

The model is tested against an extensive set of instances, which contain small (s) and medium (m) instances for 2 or 5 agents with different types of bills-of-materials.

The Price of Decentralization is computed as PoD = $\frac{\text{Cost}_{\text{DA-LS}} - \text{Cost}_{\text{C-LS}}}{\text{Cost}_{\text{C-LS}}} \times 100\%$. We define $Cost_{\text{DA-LS}}$ as the total cost under decentralized allocation and $\text{Cost}_{\text{C-LS}}$ as the optimal centralized cost.

Instance Category	**Price of Decentralization**
s A2 BASE	−9.54%
s A5 BASE	−15.35%
m A2 BASE	−7.43%
m A5 BASE	−10.46%

Table 1: Relative increase in total cost under decentralized coordination (price of decentralization).

At first glance, the results seem to show a significant efficiency loss . However, this gap must be interpreted in light of the well-documented barriers to centralized planning. Our DA-LS mechanism directly tackles these difficulties by requiring agents to share only their incremental costs (bids) instead of their proprietary cost data. Therefore, the "price of decentralization" is not a weakness but a necessary investment for creating a workable, privacy-respecting cooperative framework, which is a key goal in the search for effective coordination mechanisms. Just like the centralized model, based on the computations, the decentralized planning achieves the service level of 100%, while, if the agents planned separately, trying to cover the whole demand in all instances, it would drop on average to 92%, in some cases lower than 85%. From computational perspective, the convergence in the iterative algorithm is reached within 3.19 iterations on average. The bundle-based incremental cost mechanism, while not strictly incentive-compatible, exhibits properties that discourage strategic manipulation as an agent a is incrementally truthful if they report their true incremental cost $IC_a(B) = c_a(d_a + D_B) - c_a(d_a)$ for a bundle B. While the proposed mechanism is not theoretically strategy-proof several integrated design features promote truthful bidding by making strategic manipulation computationally intractable. The complexity of the multi-period, multi-product problem itself makes it difficult for an agent to calculate an optimal manipulation. The iterative nature of the auction and the use of bundled periods create significant interdependencies. For example, setup costs introduce "complementarity effects", which provide incentives for agents to win consecutive periods. Furthermore, the system limits information revelation; agents can only see the final allocation outcomes, not the specific bids submitted by their competitors. This is enforced by the smart contract's commit-reveal structure, which locks in bids and prevents agents from adjusting their submissions after observing competitor behaviour. Finally, the iterative process is robust enough to converge toward a solution even if minor strategic misreporting occurs.

Comprehensive game-theoretic analyses and empirical testing of strategic behaviour remain important areas for future work. Our current results, however, demonstrate the mechanism's effectiveness under the assumption of truthful reporting, which is reasonable given the complexity of profitable deviation in this setting.

References

[Albrecht, 2010] Albrecht, M. (2010). Coordination mechanisms for supply chain planning. In *Supply Chain Coordination Mechanisms*, pages 35–61. Springer Berlin Heidelberg, Berlin, Heidelberg.

[Gansterer et al., 2021] Gansterer, M., Födermayr, P., and Hartl, R. F. (2021). The capacitated multi-level lot-sizing problem with distributed agents. *International Journal of Production Economics*, 235:108090.

Joint economic lot sizing with simultaneous determination of the required amount of product-specific reusable transport items (RTIs)

Marion Lemke
Landshut University of Applied Sciences/TZ PULS, Germany;
Catholic University Eichstätt-Ingolstadt, Germany (PhD-student)
marion.lemke@haw-landshut.de

Heinrich Kuhn
Catholic University of Eichstätt-Ingolstadt, Germany
heinrich.kuhn@ku.de

Abstract

Increasing awareness of sustainability and regulatory pressure drive the use of product-specific returnable transport items (RTI) instead of disposable packaging. However, little attention has been paid to decision support for the joint optimization of economic production and order lot sizes, with the simultaneous calculation of RTI fleet sizes, considering different modes of ownership in a multi-tier supply chain.

This research models a three-tier supply chain consisting of a component supplier, a manufacturer, and a customer. The aim is to determine the joint economic production and delivery lot sizes, while simultaneously determining the required RTI fleet sizes for both the supplier's and the manufacturer's products, considering different cost components for RTI (rent, investment, or a combination thereof). In addition, capacity constraints must be respected. To achieve these objectives, a MINLP is developed and solved. Numerical examples are provided to demonstrate the model's behaviour.

1 Introduction

Driven by sustainability considerations and also called for by regulations such as the EU's Packaging and Packaging Waste Regulation [1], companies should integrate reusable packaging into their logistics processes [6]. In this study, we focus on product-specific returnable transport items (RTIs). The costs of acquiring them result from either a purchase (investment), continued rental charges, or a combination of both. Figure 1 shows an example of an RTI for which a combination of rental and investment costs is incurred.

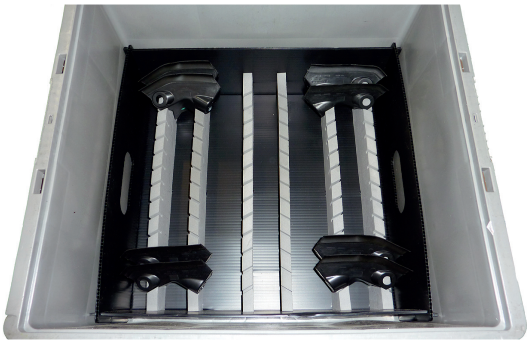

Figure 1: Example of product-specific RTI with combined cost types: rent for outer box and investment for tray with foam rails (photo taken by the author).

The responsible company within a supply chain has to determine and procure the amount of packaging required in the system, which is referred to as the fleet size. Product-specific RTIs are commonly purchased for the whole production life cycle of the respective product. The fleet size depends on the policy for returning empty RTI, expected fleet shrinkage, and the production lot size. Commonly, the production lot size is defined, and subsequently, the fleet size is determined. However, if the fleet size is insufficient, the planned production lot sizes cannot be realized. On the other hand, excess RTIs are inefficient. Therefore, the fleet size should be as small as possible. Furthermore, it can be more cost-effective to determine production lot size and fleet size simultaneously [3].

On the other hand, as shown by multiple authors, coordinating lot sizing decisions in supply chains increases competitiveness. However, the challenge of coordinating lot sizes in a supply chain while determining the RTI fleet sizes simultaneously has received little attention in the literature [4]. Therefore, we model a multi-tier supply chain in a deterministic setting to simultaneously determine the cost-effective production and delivery lot sizes, as well as the RTI fleet sizes, while considering different return policies and expected shrinkage of fleets. For demonstration purposes, this model is then applied to a three-tier supply chain, and the resulting MINLP is solved for several real-life instances. The results confirm expectations for better cost performance when coordinating fleet and lot size within supply chains.

2 The model

We base our model on the MLCLSP [5] with a finite planning horizon T. The objective is to minimize the total lot-dependent costs, which consist of setup and order costs, holding costs, transportation costs, and costs for RTIs. Additional loop days and processing steps, such as cleaning, are not considered. We further assume the following: a linear product structure, constant demand per period at the final customer; no backlog or repacking is allowed; lead times and return quantities are deterministic; production setup always occurs at the beginning of a period; open lot sizing; and all-parts-available policy.

The decision variables are the production lot sizes q_i, the production schedule $\gamma_{i,t}$, the delivery lot sizes and schedules $q_{i,t}^{\text{del}}$, the return policies ε_i and the RTI fleet sizes ρ_i. Figure 2 illustrates an exemplary, three-tier supply chain, highlighting the decision variables.

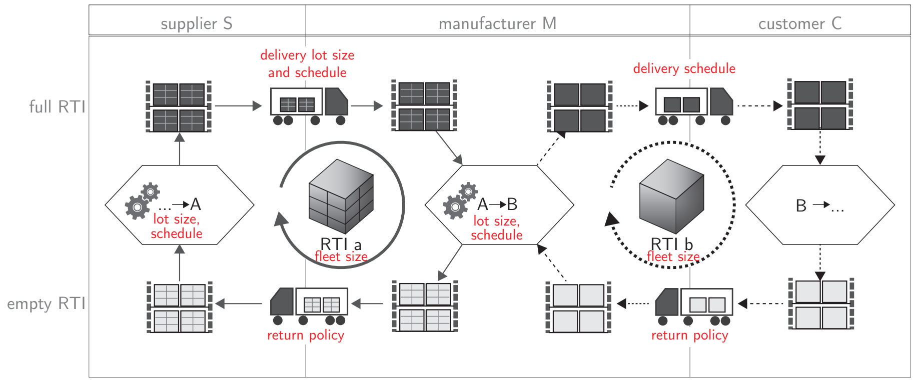

Figure 2: Example of a three-stage supply chain, which also illustrates the decision variables

Further, production and storage capacity (separately for full and empty RTIs) need to be observed, and stock levels need to be non-negative (for products and for RTIs). Furthermore, no parallel production is allowed, and component stock must be sufficient before starting a production order. Additionally, we model the expected fleet shrinkage as a linear function of the respective lot sizes.

3 First results

This model is a MINLP with a non-convex objective function. We employ fix-and-optimize heuristics, combined with decomposition strategies, approximations, and linearization, to solve various instances from practice and conduct the first sensitivity analyses using the commercial solver Baron [7].

The initial results indicate that coordinated lot and fleet sizes within a supply chain increase cost efficiency, thereby confirming the expectations. However, these results also suggest that this effect weakens if the RTI used between the manufacturer and the customer is more expensive than the RTI used between the supplier and the manufacturer. They also demonstrate that the choice of return policy can have a significant impact on the cost situation. Therefore, supply chain managers should also consider optimizing the configuration of RTI systems when determining production lot and RTI fleet sizes.

Acknowledgment

This research was supported by the project PR|IN|CE (Process Innovation Center) funded by the Bavarian State Ministry of Science and Art (FKZ H.5-H3361.LA/3/4).

References

[1] European Parliament, Council of the European Union, Regulation 2025/40: PPWR, available online: http://data.europa.eu/eli/reg/2025/40/oj (2025)

[2] Hellström, Daniel and Johansson, Ola, The impact of control strategies on the management of returnable transport items, Transportation Research Part E: Logistics and Transportation Review, 46 (6), 1128–1139 (2010)

[3] Lemke, Marion, Production Lot Size Optimization with simultaneous fleet size calculation for returnable transport items, Working Paper HAW Landshut, Germany, July 2025

[4] Mahmoudi, Monirehalsadat and Parviziomran, Irandokht, Reusable packaging in supply chains: A review of environmental and economic impacts, logistics system designs, and operations management, International Journal of Production Economics, 228, 107730 (2020)

[5] Tempelmeier, Horst, Production Analytics: Modelle und Algorithmen zur Produktionsplanung unter deterministischen und stochastischen Bedingungen. Horst Tempelmeier and BoD – Books on Demand, 7th edition (2023)

[6] Zaviša, Žilvinas, The Role of Reusable Packaging Implementation Toward More Sustainable Supply Chain, Management Theory and Studies for Rural Business and Infrastructure Development, 46 (2), 190–201 (2024)

[7] Zhang, Yi and Sahinidis, Nikolaos V., Solving continuous and discrete nonlinear programs with BARON, Computational Optimization and Applications (2024)

The trade-off between costs and carbon emissions from lot-sizing decisions

Marcel Turkensteen
Department of Economics and Business Economics, Aarhus University
matu@econ.au.dk

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

Abstract

We consider the trade-off between costs and carbon emissions of the following lot sizing decision. A company orders at regular intervals from a supplier, and the ordered goods are transported by truck. The size of the order determines the degree of utilization of the truck as well as the amount of inventory space needed, both of which drive carbon emissions and costs. We conduct a survey of empirical studies in order to establish the possible marginal emissions from holding inventory and performing a shipment with a truck.

In our experiments, we vary relevant factors related to the logistics situation, such as product characteristics, logistics considerations, and different demand scenarios. We measure the so-called efficient frontier, measured by all efficient solutions obtained by the weighted method (meaning that each efficient solution is optimal for some non-negative carbon cost value). The result is that it is often costly to reduce carbon emissions from the cost optimal solution (compared to carbon prices in the market), unless different vehicle types are available, transportation costs are low, or demand fluctuates.

1 Introduction

Recently, there has been much attention in Operations Research and Operations Management on the minimization of carbon emissions due to logistics decisions. Road freight transportation is an important source of these emissions. As of 2016, road freight transportation contributes to 4.95% of greenhouse gas emissions in the European Union and warehouses to 0.55%. Many initiatives are taken to reduce the environmental impact of transport, such as the Smartway initiative of the American Environmental Protection Agency (EPA) (also active in Canada) and the European transport initiatives such as the Transport and Environment initiative. These programs improve the environmental impact of transport through measurement and reporting of impacts and the dissemination of best practices (e.g., smart routing and economic driving). Furthermore, the European Union has further planned legislation in their European Green Deal, formalized in the Fit for 55 framework, which aims to reduce emissions by 55% up to 2035.

One way to reduce the transport emissions is to improve the utilization of a vehicle, since a fully loaded vehicle emits much less carbon per item or per ton than a poorly utilized vehicle. However, if one wishes to use well-utilized vehicles, it may result in low shipment frequencies and high inventory costs (and possibly high emissions from holding inventory). The trade-off between the potentially low carbon emissions from infrequent shipments versus the high inventory holding costs is relevant for a logistics provider who wishes to reduce carbon emissions.

An intuitive way of modeling this trade-off is using the Economic Order Quantity (EOQ) model. In this model, the decision is on how much and when to order. The EOQ model presumes a constant demand over an infinite time horizon and a continuous review setting, i.e., orders can be placed at any point in time. Many studies on EOQ models have considered the environmental impact in the form of carbon emissions and included measures such as caps on the amount of carbon emitted and carbon taxes. However, the assumptions make EOQ models quite stylized, in particular the assumption that demand is constant over time.

In this research, we choose to use *Economic Lot-Sizing (ELS)* models for the trade-off. Mostly, the focus in green lot-sizing is on the development of efficient algorithms for the formulated problems. One shortcoming of the mentioned ELS studies is that they only introduce generic emission parameters related to transport (per shipment and sometimes per item transported), and to inventory (per item stored at the end of a given period, or a fixed emission quantity for having any inventory at all). A survey on green inventory management by [1] confirms that the focus is mathematical modeling approaches, rather than on the computation of emission parameters. However, little effort is made to determine realistic values of these parameters. In contrast the previously mentioned study, [2] use lot-sizing models to measure the impact of measures such as carbon caps, taxes, and trading schemes. To the best of our knowledge, this is the only ELS paper that focuses on the trade-off between emissions savings and costs. However, the chosen parameter values are fictitious. Therefore, a conclusion such as "emission caps could be achieved more cost-effectively by adjusting operational decisions than by investing in costly more energy-efficient technology" [2] could be the result of the parameter choice.

ELS models provide a realistic way of modeling the trade-off between costs and carbon emissions in specific but relevant situations. However, current green lot-sizing models are often rather stylized. That makes it difficult to assess the trade-off: Is it actually worthwhile to reduced the number of shipments in an ELS setting to reduce carbon emissions? The aim and contribution of this research is to fill this gap by answering the following question: *Under which conditions is it sensible for a decision-maker, e.g., a logistics provider, to reduce carbon emissions related* *to shipping decisions in Economic Lot-Sizing (ELS)?* To that end, we consider the following sub-questions:

- What is the shape of the trade-off between cost and carbon emissions in an ELS setting?
- What are relevant parameters that determine this trade-off?

To answer these questions, we use the following methodology. We formulate a new ELS model, which includes realistic emission parameters. The emissions arise from transporting a given amount of items by truck and from holding items on inventory as well as possible. The model needs realistic ranges of values to ensure that it is valid for its purpose, assessing the trade-off between carbon emission and costs under realistic conditions. In order to find the trade-off curve, we use techniques and measures from bi-objective optimization. In the numerical experiments, we identify and vary the factors, such as product attributes, which affect the carbon emissions and/or costs and thereby, the trade-off between these objectives.

We highlight that our contribution is the construction and usage of a model for generating insights rather than the solution of a given model or the development of a novel solution approach. Our contribution, based on answering the above questions, is relevant as it provides practitioners and researchers with insights and tools to perform a realistic trade-off assessment and to identify situations where carbon emissions can be reduced at a low cost. An additional benefit of our experiments is that we can identify the relative relevance of inventory and transport-based emissions. In addition to the obtained insights, as a by-product, we have conducted a systematic literature study to assess carbon emission data from transporting a given amount of items and from holding items on inventory, which may be useful for other studies as well.

2 The problem setting

The setting which we consider is as follows. A company orders materials from a supplier on a periodic basis. We assume that the demand quantities within the horizon are known and deterministic. Emissions and set-up costs of production, e.g., related to starting a machine, are not considered in this paper.

For the transportation of the shipment, dedicated Full Truck Load vehicles are available. Note that if Less Than Truck Load vehicles are used, it is not clear that larger shipment sizes lead to reduced carbon emissions, since it depends on the other load traveling on the truck. The delivery is in the form of *direct shipments* between a supplier and an OEM. So-called *milk run* shipments, in which deliveries from multiple suppliers are collected, are not considered in this study. The inclusion of such shipments, in particular when the timing of deliveries from multiple suppliers can vary, can result in a highly intractable Inventory Routing Problem or Joint Replenishment Problem formulation.

We assume that decisions on the transportation route are independent of the delivery decisions. Only one route is considered, and the vehicle follows the driving conditions on the route. As a consequence, the choice between different routes and speeds is not a decision within our case, but a choice taken separately. This modeling choice will be motivated in more detail.

In our study, we consider the trade-off between costs and carbon emissions from an *operational and tactical* perspective. The resources and the supply network are available, and the question is how to use the resources. In this situation, an ELS type model allows us to mimic the decisions on how much to ship and how much to put on inventory in a specific time period and their consequences in detail. This offers a natural starting point for analysis. The study can be extended to strategic decisions, which usually deal with capacity and concern the composition and size of the fleet of vehicles, the reservation of capacity of different transportation modes, and the size of the storage space. The scope of our study differs from that of many EOQ studies, in which decisions under consideration are generally at a strategic level, such as investments in new technologies and the selection of supply.

To summarize, for a tactical horizon (typically up to a year ahead), we determine the periods in which we place orders. From this decision, we obtain the order and inventory quantities in each period. These quantities in turn determine the costs and the carbon emissions. The goal is to obtain solutions with different costs and carbon emissions and find the shape of trade-off between the two. Further details can be found in [3] on which this abstract is based.

References

[1] Marklund, J., Berling, P. Green inventory management. In: Sustainable supply chains. Springer, pp. 189–218, 2017.

[2] Benjaafar, S., Li, Y., Daskin, M. Carbon footprint and the management of supply chains: Insights from simple models. IEEE Transactions on Automation Science and Engineering 10 (1), 99–116, 2013.

[3] Turkensteen, M., Van den Heuvel, W. The trade-off between costs and carbon emissions from economic lot-sizing decisions. INFOR: Information Systems and Operational Research, 61(2), 169–198, 2023.

Carbon-Aware Stochastic Lot-Sizing with Perishable Inventory

Bahar Cennet Okumuşoğlu
TUM School of Management, Technical University of Munich
bahar.okumusoglu@tum.de

Martin Grunow
TUM School of Management, Technical University of Munich
martin.grunow@tum.de

Abstract

The food sector in the European Union aims to reduce greenhouse gas emissions by incorporating various policies during planning stages. In this regard, demand substitution is a highly relevant practice for lowering emissions and addressing stock-outs in the food industry. In this work, we consider a vertically integrated food producer operating a capacitated production system for multiple perishable products with substitutable demands. Due to environmental regulations, the producer must also ensure that the periodic carbon emissions do not exceed predefined environmental limits. We assume that products have fixed shelf-lives and constant production lead times. We formulate this problem as a multi-stage stochastic program with an infinite time horizon. To deal with the high-dimensional state space, we propose an approximation based on a two-stage stochastic program and solve it in a rolling horizon procedure. Our preliminary numerical results highlight the value of our proposed formulation for production and inventory planning of perishables and product substitution.

1 Introduction

In the European Union, various industries make significant efforts to reduce greenhouse gas emissions. This effort is particularly important for perishable products as they heavily have carbon-emitting impacts. To mitigate these impacts, more environmentally friendly practices such as demand substitution are being increasingly used in the food sector. In this regard, production and inventory planning for perishables plays a pivotal role in achieving carbon emission goals. Traditionally, these planning models have only focused on minimizing setup, production and inventory costs.

However, shifting toward more sustainable practices necessitates more advanced modeling in decision-making. In particular, carbon-aware planning models with perishable inventory address the need to minimize operational costs while simultaneously restricting carbon emissions through setup, production and inventory.

In production planning, lot-sizing is a fundamental problem that determines the optimal setups and production amounts while meeting demand with minimum operational costs. With the increasing need for sustainable policies, integrating carbon-awareness into the lot-sizing has become more relevant to the academic community as well as the industry. Early research on lot-sizing with carbon emission constraints [Absi et al., 2013, 2016, Benjaafar et al., 2013] introduced various carbon policies and emission constraints into production planning models and highlighted their impacts on operational decisions. These policies and constraints are important for high-emission sectors such as the food industry, where the production and inventory of perishables contribute heavily to carbon footprints. In particular, carbon-aware lot-sizing helps decision-makers obtain cost-optimal policies aligned with their carbon emission goals.

In inventory management, perishable products present challenges due to their fixed shelf-lives and stock-outs. In the literature, some studies [Deniz, 2007, Karaesmen et al., 2011, Nahmias, 1982] have already emphasized the importance of managing age distributions in stock. On the other hand, demand substitution has also been shown to be an effective strategy for minimizing stock-outs [Hsu et al., 2005, Lang and Domschke, 2010]. To address these real-world challenges, the stochastic lot-sizing problem can be extended to include age-differentiated demand and stockout-driven substitution. As perishable products contribute significantly to carbon emissions through production and inventory, carbon-aware lot-sizing remains as an important problem in obtaining cost-optimal and sustainable policies in the food industry.

Stochastic programming approaches are efficient in obtaining solutions for the stochastic lot-sizing under practical applications, see, e.g., [Thevenin et al., 2021, 2022]. However, the carbon-aware lot-sizing under demand substitution remains underexplored in production and inventory planning for perishable products as they significantly contribute to carbon emissions whereas substitution among these products offers reduced waste and increased service levels. We aim to contribute to the literature by proposing a stochastic programming framework that combines carbon-awareness and substitution for perishable products, moving toward more sustainable and efficient production and inventory systems.

2 Problem Formulation

We consider a vertically integrated food producer operating a shared capacitated production system for multiple products with substitutable demands. The producer has to make setup, production, and substitution decisions for each product. Moreover, the producer policy dictates that the demand cannot be negative and backlogged due to food regulations, and the unmet demand results in lost sales. We denote the set of products by $\mathcal{I} = \mathcal{I}_r \cup \mathcal{I}_g$, where $\mathcal{I}_r$ and $\mathcal{I}_g$ are the sets of regular and green products, respectively. Green products emit less carbon dioxide during production compared to regular products. However, they have shorter shelf lives and produce more emissions during storage. Due to environmental regulations, the producer has to ensure that the periodic carbon emission during the setup, production, and storage does not exceed the unitary environmental impact. For each product $i \in \mathcal{I}$, we define the set of age-classes as $\mathcal{K}_i = \{1, \ldots, a_i\}$, where $a_i > 1$ is the shelf-life of product i. We denote the set of periods in the planning horizon by $\mathcal{T}$.

Let us now denote the setup decision by x_{it} that takes the value of one if product i enters production in period t, and zero otherwise. The production and the shortage quantity of product i in period t are denoted by y_{it} and s_{it}, respectively. The setup, production and shortage cost of product i in period t are denoted by c_{it}^{setup}, c_{it}^{prod} and c_{it}^{lost}, respectively. Let v_{ikt} represent the inventory amount of product i in age-class k at the end of period t. The company follows a first-in-first-out (FIFO) issuing policy for each product. For each product i in age-class k, the inventory cost is given by c_{ikt}^{hold}. The inventory amount of product $i \in \mathcal{I}$ that exceeds its shelf-life a_i within the planning horizon cannot be carried over to the next period and is disposed of with a waste cost denoted by c_{it}^{waste}. For brevity, we also define the decision variable z_{it} as the total inventory amount of product i at the end of period t. The demand of product i follows a stochastic process denoted by $\mathcal{D}_i = \{D_{i1}, \ldots, D_{i|\mathcal{T}|}\}$. For each product i, we define the set of products that can be used to satisfy the demand of product i in the planning horizon as $\mathcal{I}_i^- \subseteq \mathcal{I}$ whereas the set of products that product i can be used as a substitute for is denoted by $\mathcal{I}_i^+ \subseteq \mathcal{I}$. Let q_{ijt} be the amount of product i that is used to substitute the demand of product j in period t. For each product i and product $j \in \mathcal{I}_i^-$, there is a substitution cost denoted by c_{ijt}^{sub} in period t. Note that no substitution cost is incurred for any product that is used for its own demand i.e., $c_{iit} = 0$.

In the remainder of our paper, we will simplify our notation by omitting the index of interest when referring to the vector of decision variables or parameters. For example, x_t denotes the vector in the form of $(x_{1t}, \ldots, x_{|\mathcal{I}|t})$. We let $d_{it} = D_{it}(\omega_{it})$ be the realization of the demand process of product i in period t. We also denote by $D_{[1,t]} = (D_1, \ldots, D_t)$ the history of demand up to period t whereas the history of demand realizations up to time t is given by $d_{[1,t]} = (d_1, \ldots, d_t)$.

For each product i, the unit carbon emissions that result from setup operations, production, and inventory in period t are denoted by e_{it}^{setup}, e_{it}^{prod} and e_{it}^{inv}, respectively. Due to environmental regulations, the company wants to ensure that the average carbon emission during setup operations, production, and inventory in period t does not exceed the maximum carbon emission E_t^{setup}, E_t^{prod} and E_t^{inv}, respectively. For brevity, let us define the unit carbon emission vector as $e_{it} := (e_{it}^{setup}, e_{it}^{prod}, e_{it}^{inv})$ and maximum carbon emission vector as $E_t := (E_t^{setup}, E_t^{prod}, E_t^{inv})$. Finally, we define the following set denoted by $\mathcal{M}_t$ as the feasible emission set that restricts the total periodic emission resulting from setup, production, and inventory in period t.

Now we are ready to present the stochastic dynamic programming formulation for our problem:

$$
\begin{aligned}
\mathcal{Q}_t(v_{t-1}, d_{[1,t]}) = \min \quad & \sum_{i \in \mathcal{I}} (c_{it}^{setup} x_{it} + c_{it}^{prod} y_{it} + c_{it}^{lost} s_{it} + c_{it}^{waste} v_{i,a_i,t}) \\
& + \sum_{i \in \mathcal{I}} (\sum_{k \in \mathcal{K}_i'} c_{it}^{hold} v_{ikt} + \sum_{j \in \mathcal{I}_i^-} c_{ijt}^{sub} q_{ijt}) \\
& + \mathbb{E}_{\boldsymbol{D}_{t+1}} \{\mathcal{Q}_{t+1}(\boldsymbol{v}_t, \boldsymbol{D}_{[t+1]}) | D_{[1,t]} = d_{[1,t]}\} && \text{(1a)} \\
\text{s.t. For each } i \in \mathcal{I}: & && \text{(1b)} \\
& d_{it} + s_{it} = \sum_{j \in \mathcal{I}_i^+} q_{jit} && \text{(1c)} \\
& v_{i,a_i,t} - u_{i,a_i-1,t} = v_{i,a_i-1,t-1} - \sum_{j \in \mathcal{I}_i^-} q_{ijt} && \text{(1d)} \\
& v_{i,k+1,t} - u_{ikt} = v_{ik,t-1} - u_{i,k+1,t}, \quad k \in \mathcal{K}_i'' && \text{(1e)} \\
& v_{i1t} = y_{it} - u_{i1t} && \text{(1f)} \\
& u_{ikt} \leq M_{it} r_{ikt}, \ v_{i,k+1,t} \leq M_{it}(1 - r_{ikt}), \quad k \in \mathcal{K}_i' && \text{(1g)} \\
& z_{it} = \sum_{i \in \mathcal{K}_i} v_{ikt} && \text{(1h)} \\
& (x_t, y_t, z_t) \in \mathcal{M}_t(e_t, E_t) && \text{(1i)} \\
& y_{it} \leq C_{it} x_{it} && \text{(1j)} \\
& s_{it} \leq d_{it} && \text{(1k)} \\
& r_t \in \{0,1\}^{|\mathcal{I}| \times |\mathcal{K}_i|}, \ u_t, v_t \in \mathbb{R}_+^{|\mathcal{I}| \times |\mathcal{K}_i|}, \ q_t \in \mathbb{R}_+^{|\mathcal{I}_i^+| \times |\mathcal{I}_i^-|} && \text{(1l)} \\
& x_t \in \{0,1\}^{|\mathcal{I}|}, \ y_t, s_t \in \mathbb{R}_+^{|\mathcal{I}|}. && \text{(1m)}
\end{aligned}
$$

3 First Numerical Results

The multi-stage stochastic programming formulation in (1) is computationally intractable due to its high-dimensional state space and exponential growth in scenario trees. To approximate production and inventory policies, we utilize a two-stage stochastic programming framework. Scenarios are generated through uniform sampling. We first adopt a static-static uncertainty approach, where setup and production decisions are made in the first stage, while second-stage variables depend on the specific scenario. The resulting deterministic equivalent problem is solved using the state-of-the-art commercial solver Gurobi. Due to pending industry data, for our initial experiments, we use synthetic data, focusing on only two products: a "green" product and a "regular" product. Demand is assumed to follow a lognormal distribution, and results are analyzed across different scenario sizes, varying setup costs, and demand levels. We also compare our stochastic programming framework with the deterministic approximation in which the mean of the corresponding distribution is used for each product.

Our experiments yield several insights. When varying setup costs, we observe that low or no setup costs reduce total inventory levels and lead to higher production and substitution decisions. However, the impact of setup costs diminishes when total demand exceeds capacity. On the other hand, demand variations highlight the value of substitution: it reduces waste and lost sales, particularly under low-demand scenarios. In congested systems with high demand, substitution becomes less effective as the system's flexibility decreases. Also, emissions from production and inventory are influenced by setup costs and demand levels. Low setup costs result in higher inventory-related emissions, whereas high setup costs reduce emissions through lower production levels. Additionally, emissions are lowest in low-demand scenarios, with substitution offering significant reductions in both waste and lost sales, particularly under moderate congestion.

References

N. Absi, S. Dauzère-Pérès, S. Kedad-Sidhoum, B. Penz, and C. Rapine. Lot sizing with carbon emission constraints. *European Journal of Operational Research*, 227 (1):55–61, 2013. ISSN 0377-2217.

N. Absi, S. Dauzère-Pérès, S. Kedad-Sidhoum, B. Penz, and C. Rapine. The single-item green lot-sizing problem with fixed carbon emissions. *European Journal of Operational Research*, 248(3):849–855, 2016. ISSN 0377-2217.

S. Benjaafar, Y. Li, and M. Daskin. Carbon footprint and the management of supply chains: Insights from simple models. *IEEE Transactions on Automation Science and Engineering*, 10(1):99–116, 2013.

B. Deniz. *Essays on Perishable Inventory Management*. Doctoral dissertation, Tepper School of Business, Carnegie Mellon University, Pittsburgh, 2007.

V. N. Hsu, L. Chung-Lun, and X. Wen-Qiang. Dynamic lot size problems with one-way product substitution. *IIE Transactions*, 37(3):201–215, 2005.

I. Z. Karaesmen, A. Scheller-Wolf, and B. Deniz. *Managing Perishable and Aging Inventories: Review and Future Research Directions*, pages 393–436. Springer US, New York, NY, 2011. ISBN 978-1-4419-6485-4.

J. C. Lang and W. Domschke. Efficient reformulations for dynamic lot-sizing problems with product substitution. *OR Spectrum*, 32(2):263–291, 2010. ISSN 1436-6304.

S. Nahmias. Perishable inventory theory: A review. *Operations Research*, 30(4): 680–708, 1982.

S. Thevenin, Y. Adulyasak, and J.-F. Cordeau. Material requirements planning under demand uncertainty using stochastic optimization. *Production and Operations Management*, 30(2):475–493, 2021.

S. Thevenin, Y. Adulyasak, and J.-F. Cordeau. Stochastic dual dynamic programming for multiechelon lot sizing with component substitution. *INFORMS Journal on Computing*, 34(6):3151–3169, 2022.

Dynamic Re-optimization of Timed Routes in Multistep Production Planning

Camil Zarrouk
Mines Saint-Etienne, Univ. Clermont Auvergne, CNRS, UMR 6158 LIMOS, CMP, Department of Manufacturing Sciences and Logistics, Gardanne, France, and STMicroelectronics, Crolles, France
camil.zarrouk@st.com

Nabil Absi
Mines Saint-Etienne, Univ. Clermont Auvergne, CNRS, UMR 6158 LIMOS, CMP, Department of Manufacturing Sciences and Logistics, Gardanne, France
absi@emse.fr

Quentin Christ
STMicroelectronics, Crolles, France
quentin.christ@st.com

Stéphane Dauzère-Pérès
Mines Saint-Etienne, Univ. Clermont Auvergne, CNRS, UMR 6158 LIMOS, CMP, Department of Manufacturing Sciences and Logistics, Gardanne, France
dauzere-peres@emse.fr

Renaud Roussel
STMicroelectronics, Crolles, France
renaud.roussel@st.com

1 Introduction

Semiconductor manufacturing includes probably the most complex production processes, with each product requiring several hundred steps to be processed on heterogeneous machines. Additionally, re-entrant flows and congestion must be considered, which results in a production lead time ranging from 8 to 20 weeks from a blank silicon wafer to a wafer with finished chips. This complexity is further exacerbated by the large volume of wafers processed in a single factory with hundreds of heterogeneous machines, as well as the high number of different products in high-mix factories.

In this work, we consider an operational production planning problem encountered in semiconductor manufacturing. This is a multi-product, multi-step, multi-machine production planning problem, with the goal of minimizing the work-in-process, inventory and backlog costs, while respecting machine capacity. Given the scale of operations, where tens of thousands of wafers are manufactured every day, this is a highly complex optimization problem.

To address this problem, a novel and efficient column generation approach was proposed in [1]. This approach was later extended in [2] to consider various industrial constraints. However, to improve the convergence of the column generation approach for industrial instances, it is crucial to start with an appropriate initial set of columns. Since production plans are executed over a rolling horizon and span several months, we propose, in this work, to integrate and adapt previously generated timed routes to enhance the stability of the production plan and to significantly reduce computational times.

2 Concept of Timed Route

The approach in [1] relies on the notion of "timed routes", where each production step is assigned to a specific time period in the manufacturing route of a product. Thus, a timed route provides a complete temporal representation of the production flow of a given product, specifying precisely in which period the production capacity is required. The problem consists of optimizing the production quantities along these timed routes. However, the number of possible timed routes grows exponentially with the number of production steps, making exhaustive enumeration impractical in industrial instances. This justifies the use of a column generation approach, where only the most relevant timed routes are generated. The notion of timed route was extended in [2] to include machine assignment, leading to the notion of "machine timed route", where each production step is also assigned to a specific machine.

By solving the problem with continuous variables, it is possible to find optimal solutions even for industrial instances, typically with computational times of approximately 24 hours. However, this is still far too long for operational deployment in industrial settings. Therefore, techniques to accelerate the resolution process are essential.

3 Warm Start approach

In semiconductor manufacturing, it is rare to compute a production plan from scratch. Typically, plans are determined in a rolling horizon approach, where a new plan is calculated periodically (for example, every week), but due to long cycle times, many products from the previous plan remain in production. Plans may also be computed in what-if scenarios, where only limited changes are introduced. In this case, a large number of products will not be impacted by the changes. In both cases, the plan computed in a previous calculation can be reused as a starting point to determine the new plan, which can greatly improve the quality of the initial solution and accelerates convergence, thus reducing the computational time to find the optimal solution. This strategy is known as a "*Warm Start*".

Timed routes are particularly well suited to apply a *Warm Start* approach, because generating a new production plan can reduce to adjusting the quantities associated with the existing timed routes. It is also possible to reuse previously generated timed routes to generate new ones.

However, reusing previously generated timed routes is not trivial, and several strategies can be considered:

- *Limiting the number of reused timed routes.* More reused timed routes may lead to a better initial solution, but may also slow down subsequent iterations.
- *Adapting timed routes to fit the new context.* Adjustments may be needed when demand or machine availability has changed.
- *Focusing the re-optimization on impacted products.* For what-if analyses, we can decide to focus only on a subset of products.

4 Numerical results

To assess the effectiveness of the *Warm Start* strategy (i.e., reusing the optimal solution from a previous resolution), we solve the same instance twice: Without (*From Scratch*) and with warm start. In our first computational experiments, we used a reduced data set with a limited number of products, ensuring computational times of a few minutes.

In this work, we consider industrial instances provided by the most advanced site of STMicroelectronics in Crolles, France. The instances include 30 products and 700 machines, and the plan is determined over 60 periods (days). The different instances correspond to different moments in time. In Tables 1 and 2, the first line represents to the different instances and the second line, respectively the third line, corresponds to the computational time to find an optimal solution *From Scratch*, respectively with *Warm Start*. Finally, the last line shows the time saving in percentage obtained by using *Warm Start*.

As shown in Table 1, the warm start strategy yields time savings of 90% in what-if scenarios. In rolling horizon planning (Table 2), results are more variable and instance-dependent, but generally lead to significant gains. We expect these time savings to be even more significant on full datasets, as the ratio of timed routes not impacted by the changes should be much larger. Preliminary experiments on full instances are showing promising results. For example, in a what-if scenario, finding

	CPU time (seconds)								
Instance	#1	#2	#3	#4	#5	#6	#7	#8	#9
From Scratch	405	329	917	603	356	420	512	610	344
Warm Start	31	8	31	84	20	18	49	36	24
Saved (%)	92%	98%	97%	86%	94%	95%	91%	94%	93%

Table 1: Comparison of CPU times between *Warm Start* and *From Scratch*: What-if analysis.

	CPU time (seconds)								
Instance	#1	#2	#3	#4	#5	#6	#7	#8	#9
From Scratch	329	382	483	451	596	629	696	777	770
Warm Start	267	341	223	217	571	480	164	879	563
Saved (%)	19%	11%	54%	52%	4%	24%	76%	-13%	27%

Table 2: Comparison of CPU times between *Warm Start* and *From Scratch*: Rolling horizon.

an optimal solution *From Scratch* took approximately 33 hours, while *Warm Start* converged in just 1.5 hours. Moreover, *Warm Start* allows a good solution to be reached very quickly: A few seconds are enough to obtain a 99%-optimal solution, whereas it took 20 hours with *From Scratch*.

References

[1] Sébastien Beraudy, Timed route approaches for large multi-product multi-step capacitated production planning problems, *European Journal of Operational Research* (2022), 300(2), 602-614.

[2] Benjamin Anthouard, Timed Route Approaches for Production Planning with Time Constraints, *2023 International Workshop on Lot Sizing (2023).*

Decision Strategies on a Two-level Multi-Stage Stochastic Lot Sizing Problem considering Blending

Maurício Rocha Gonçalves
UNESP-IBILCE (Brazil)
mauricio.rocha@unesp.br

Raf Jans
HEC Montréal/GERAD (Canada)
raf.jans@hec.ca

Silvio Alexandre de Araujo
IBILCE-UNESP/GERAD (Brazil)
silvio.araujo@unesp.br

Abstract

Blending problems occur when different components need to be mixed in order to form an end product. The recipe typically has some flexibility as long as some specific quality conditions are met. We consider this as a multi-period, multi-level lot sizing problem in which we need to determine the purchase quantities of the components and the production quantities for the end products (and their respective setups). Lead times for the availability of the components and the end products are assumed. Demand uncertainty is represented by a scenario tree in which the decision-making unfolds period by period as demand is stochastically revealed. The objective is to minimize the expected total cost, including inventory, production, purchasing, setup, and lost sales costs. We present different decision-making framework combinations for deciding on the setup and the production quantities of the end products and the components in the multi-stage stochastic setting. We additionally develop a heuristic method that employs a sequence of two-stage stochastic approximations for solving the resulting scenario-tree based models. Computational experiments are conducted to evaluate the performance of the methods to approach the stochastic problem with respect to the associated costs and solution difficulty.

1 Problem statement

We consider a production planning problem involving $1, \ldots, |\mathcal{P}|$ distinct end products. These products are obtained through a blending process that mixes $1, \ldots, |\mathcal{C}|$ different components. The bill of materials of the products also associates the specifications that must be satisfied over $1, \ldots, |\mathcal{Q}|$ relevant quality characteristics. The demand d^E for the end products must be met over $1, \ldots, |\mathcal{T}|$ discrete time periods.

An integrated plan is required for both the purchasing of the components and the production of the end products [1]. Due to the presence of lead times, the end products produced and the components purchased in a given time period may only become available in some time periods later.

The plan involves deciding on the purchase and the production quantities, denoted by $x^C \geq 0$ for the components and $x^E \geq 0$ for the products, which are only possible when there are respective setups $y^C \in \{0,1\}$ and $y^E \in \{0,1\}$. The blending decision, represented by $p^E \geq 0$, determines the quantity of component used in the production of end products. Inventory levels $s^C \geq 0$ for the components and $s^E \geq 0$ for the products are carried between periods. Any unmet product demand is considered a lost sale, denoted by $l^E \geq 0$.

The decision-making process is characterized by demand uncertainty in a multi-stage stochastic setting. The demand is realized between the time periods $1, \ldots, |\mathcal{T}|$ by an underlying probability distribution. The demand realization at each stage has a finite number of possible outcomes. We use a scenario tree to represent the stochastic structure of the problem [2]. In the scenario tree, each node $n \in \mathcal{N}$ represents a point in the sequential unfolding of demand realizations $d^E(n)$, associated with a specific probability $\mathbb{P}(n)$ and the corresponding decision-making at a given stage $t(n) \in \mathcal{T}$.

The resulting multi-stage stochastic problem based on the scenario tree formulation is (1)-(12).

$$\begin{aligned} \min \sum_{n \in \mathcal{N}} \sum_{i \in \mathcal{C}} \mathbb{P}(n) \left(sc^C_{it(n)} y^C_i(n) + vc^C_{it(n)} x^C_i(n) + hc^C_{it(n)} s^C_i(n) \right) \\ + \sum_{n \in \mathcal{N}} \sum_{j \in \mathcal{P}} \mathbb{P}(n) \left(sc^E_{jt(n)} y^E_j(n) + vc^E_{jt(n)} x^E_j(n) + hc^E_{jt(n)} s^E_j(n) + pc^E_{jt(n)} l^E_j(n) \right) \end{aligned} \tag{1}$$

Subject to:

$$x_j^E(n) \le M_{jt(n)}^E y_j^E(n) \qquad \forall j \in \mathcal{P},\ n \in \mathcal{N} \tag{2}$$

$$x_i^C(n) \le M_{it(n)}^C y_i^C(n) \qquad \forall i \in \mathcal{C},\ n \in \mathcal{N} \tag{3}$$

$$s_j^E(a(n)) + l_j^E(n) + \sum_{t=1}^{t(n)} \mathbf{1}_{\{t+lt_{jt}^E=t(n)\}} x_j^E(a_t(n)) = s_j^E(n) + d_j^E(n) \qquad \forall j \in \mathcal{P},\ n \in \mathcal{N} \tag{4}$$

$$s_i^C(a(n)) + \sum_{t=1}^{t(n)} \mathbf{1}_{\{t+lt_{it}^C=t(n)\}} x_i^C(a_t(n)) = s_i^C(n) + \sum_{j\in\mathcal{P}_i} p_{ij}^E(n) \qquad \forall i \in \mathcal{C},\ n \in \mathcal{N} \tag{5}$$

$$l_j^E(n) \le d_j^E(n) \qquad \forall j \in \mathcal{P},\ n \in \mathcal{N} \tag{6}$$

$$x_j^E(n) = \sum_{i\in C_j} p_{ij}^E(n) \qquad \forall j \in \mathcal{P},\ n \in \mathcal{N} \tag{7}$$

$$qu_{qj} x_j^E(n) \ge \sum_{i\in C_j} qa_{qi} p_{ij}^E(n) \qquad \forall j \in \mathcal{P},\ n \in \mathcal{N} \tag{8}$$

$$ql_{qj} x_j^E(n) \le \sum_{i\in C_j} qa_{qi} p_{ij}^E(n) \qquad \forall j \in \mathcal{P},\ n \in \mathcal{N} \tag{9}$$

$$x_i^C(n),\ s_i^C(n) \ge 0, \quad y_i^C(n) \in \{0,1\} \qquad \forall i \in \mathcal{C},\ n \in \mathcal{N} \tag{10}$$

$$x_j^E(n),\ s_j^E(n),\ l_j^E(n) \ge 0, \quad y_j^E(n) \in \{0,1\} \qquad \forall j \in \mathcal{P},\ n \in \mathcal{N} \tag{11}$$

$$p_{ij}^E(n) \ge 0 \qquad \forall i \in \mathcal{C},\ \forall j \in \mathcal{P},\ n \in \mathcal{N} \tag{12}$$

The objective function in (1) minimizes the expected total cost, which includes the setup, purchasing, and inventory costs for the components, plus the setup, production, inventory, and lost sales costs for the end products. Constraints (2) guarantee that the production of products occurs only if there is setup. Constraints(3) ensure that there is setup for the purchase of components. Constraints (4) and (5) describe the inventory balance for products and components, respectively. In these constraints, $a(n)$ denotes the parent node of node n, which links nodes across subsequent stages through inventory carrying. The auxiliary parameter $\mathbf{1}_{\{.\}}$ is used to indicate stocking under the lead times for purchased and produced the quantities, which takes the value 1 whenever the subscribed condition is verified, and 0 otherwise. Constraints (6) impose that the amount of demand that is lost is at most the total realized demand. The blending of components into end products is described by constraints (7), where constraints (8) and (9) impose the final composition to meet the quality conditions. The domain for the decisions associated with the components is determined by constraints (10). Finally, the domain for the decisions associated with the products is determined by (11), (12).

2 Decision-making settings

The multi-stage stochastic formulation (1)–(12) assumes that decisions are made in response to realized demand. In contrast to this fully dynamic decision-making setting, early planning is often necessary in practice [3]. The purchasing phase is typically determined in advance of the demand realization to ensure component availability, as it directly constrains the subsequent production phase. The production phase, in turn, may follow different decision timing depending on its anticipation relative to the demand realization. Table 1 summarizes the combinations of different anticipatory strategies considered for the purchasing and production planning, which lead to alternative settings for the multi-stage stochastic problem.

Table 1: Decision timing for purchase and production under different settings.

Problem Setting		**Purchasing Decisions**		**Production Decisions**	
Purchasing	Production	Setup y^C	Quantity x^C	Setup y^E	Quantity x^E
Static	Static	Stage Start	Stage Start	Stage Start	Stage Start
Static	Static-Dynamic	Stage Start	Stage Start	Stage Start	After Demand Realization
Static	Dynamic	Stage Start	Stage Start	After Demand Realization	After Demand Realization

3 Sequential heuristic approximation

We propose a heuristic approach in which the multi-stage problem is decomposed into a chain of two-stage stochastic subproblems [4]. The chain of subproblems has an initial subproblem defined at a dummy stage, followed by one subproblem for each node in the scenario tree. This chain of subproblems is solved sequentially across stages, iteratively using previously determined solutions to guide the construction of decisions in subsequent iterations. The goal is to progressively generate a pool of candidate variable values for the setup decisions for the components and the products, and then apply a rule to determine their final values. The final setup variable values form a partial solution for the multi-stage stochastic problem. Thus, the remaining decision variables for the multi-stage stochastic problem are obtained by optimizing the multi-stage stochastic model with these setup decision variables fixed.

4 Computational results

We carried out computational experiments using test instances under different parameter settings within the problem to evaluate the different multi-stage stochastic modeling approaches. Table 2 reports the computational results across 48 instances using the different decision-making settings to address the multi-stage stochastic problem. The reported values include the optimal solution count, the average solution time (in seconds), and the average expected cost. The header 'Exact' corresponds to the results regarding multi-stage stochastic models directly solved using the commercial solver, and 'STS' corresponds to the results using the setup variable values from the sequential two-stage stochastic heuristic. The average running time of the heuristic is 14.7 seconds.

Table 2: Summarized results for the Expected Cost and Solution Time across 48 test instances.

Problem Setting		**Opt. Count**		**Sol. Time**		**Exp. Cost**	
Purchasing	**Production**	Exact	*STS*	Exact	*STS*	Exact	*STS*
Dynamic	Dynamic	44	48	433.4	0.8	25,876,081	30,241,516
Static	Dynamic	46	48	223.6	0.8	25,939,077	30,306,594
Static	Static-Dynamic	46	48	231.5	0.8	25,945,894	30,306,594
Static	Static	48	48	14.8	0.8	25,973,742	30,348,875

5 Conclusions

The computational results highlight the trade-offs between the expected cost and the computational effort associated with different anticipatory strategies. Notably, the problem settings with lower degrees of anticipativity in the decision-making tend to result in higher objective costs, whereas those with greater anticipativity yield lower costs. The decision frameworks with more anticipation also result in a simplified mathematical model structure when compared to the fully dynamic multi-stage stochastic problem, which results in faster solution times.

The computational results also demonstrate the performance of the sequential two-stage stochastic approximation as an alternative solution approach. The sequential heuristic strategy yields implementable solutions while significantly reducing the final computational time of the mathematical models. The outcome of this approximation is heuristic values for the setup decision variables that lead to increased objective costs in the multi-stage stochastic problem.

References

[1] Fiorotto, D. J., Jans, R., and de Araujo, S. A., Integrated lot sizing and blending problems, *Computers & Operations Research*, **131**, 105255 (2021).

[2] Pochet, Y., and Wolsey, L. A., *Production Planning by Mixed Integer Programming*, Springer, Volume 149, Number 2 (2006).

[3] Bookbinder, J. H., and Tan, J.-Y., Strategies for the probabilistic lot-sizing problem with service-level constraints, *Management Science*, **34**(9), 1096–1108 (1988).

[4] Rocha Gonçalves, M., Jans, R., and de Araujo, S. A., Integrated lot sizing and blending problems under demand uncertainty, *Les Cahiers du GERAD*, Technical Report G-2025-02, Groupe d'études et de recherche en analyse des décisions (GERAD), Montréal QC, Canada (2025). Available at: https://www.gerad.ca/en/papers/G-2025-02

Stochastic Lot-Sizing and Scheduling for Quality-Differentiated Disassembly

Sajjad Hedayati
Department of Industrial Systems Engineering and Product Design, Ghent University
Industrial Systems Engineering Core lab, FlandersMake@UGent, Belgium
sajjad.hedayati@ugent.be

Stijn De Vuyst
Department of Industrial Systems Engineering and Product Design, Ghent University
stijn.devuyst@ugent.be

Abstract

We address a multi-period stochastic disassembly scheduling problem using a two-stage program where first-stage scheduling is coupled with a complex recourse stage governed by a shared procurement cost. To solve this NP-hard problem, we propose a decomposition-inspired Simulated Annealing (SA) heuristic that uses an embedded MIP solver to evaluate the recourse. We benchmark the SA against a full extensive-form MIP and a deterministic model to analyze the trade-offs between heuristic performance, solution exactness, and the value of stochastic modeling.

1 Introduction

The transition to a circular economy hinges on the ability to recover valuable modules and materials from end-of-life products efficiently. A pr imary op erational challenge in this domain is disassembly scheduling, which involves assigning and sequencing returned product cores on available machines that recover a number of still usable modules from them, with the goal of meeting a certain module demand. This process is fundamentally complicated by uncertainty: the quantity and quality of modules recovered from a core are often known only in a probabilistic sense. Effective planning, therefore, requires methods that can produce robust schedules in the face of these random module yields [1, 3]. The challenge, which we address in this work, extends beyond simple scheduling. Optimal planning must integrate short-term, operational scheduling decisions with long-term, tactical inventory policy over a multi-period horizon. The schedule chosen in the current period directly impacts the stochastic supply of modules available to meet future demands. This creates a difficult trade-off: a schedule that seems efficient locally may lead to costly future stockouts or overages. A successful model must navigate this interplay between proactive scheduling and reactive, dynamic inventory management.

To tackle this problem, we make the following contributions. First, we formulate a two-stage stochastic program that captures this dynamic link between scheduling and multi-period inventory decisions under yield uncertainty. Second, recognizing the NP-hard nature of the resulting model, we develop a Simulated Annealing heuristic—drawing on Logic-Based Benders Decomposition—to efficiently so lve th e two-stage stochastic scheduling and inventory model under yield uncertainty, and benchmark its performance against both the full extensive-form MIP and a deterministic expected-value model.

2 Problem Definition

We address a multi-period disassembly scheduling problem over a finite horizon of T periods. In period $t \in \mathcal{T} = \{1, \ldots, T\}$, a set $\mathcal{J}_t$ of returned product cores arrives which must be scheduled for disassembly on one machine from a set $\mathcal{M}$ of heterogeneous parallel machines. Disassembling a core can potentially yield modules of type $r \in \mathcal{R}$, but the exact quantity of each recovered module type is uncertain. We assume however that the yield ξ_{jkrt}, i.e. the number of usable modules of type r obtained from disassembling core $j \in \mathcal{J}_t$ on machine k, is a random variable with known probability distribution. These probabilities can follow from historical analysis, from an expert's assessment or from other knowledge in the domain. In period t, there is a demand of D_{rt} modules of type r and we try to cover that demand by optimally scheduling the cores on the machines. A schedule is both an assignment x_{jkt} (1 if core $j \in \mathcal{J}_t$ is disassembled on machine $k \in \mathcal{M}$, and 0 otherwise) and a sequence y_{ijkt} (1 if core i is disassembled immediately before core j on machine k, 0 otherwise, with $i, j \in \mathcal{J}_t$).

We formulate the problem as a two-stage stochastic program. In the first stage, we determine a proactive schedule, whose performance is then evaluated under each future scenario $\omega \in \Omega$, where Ω is a finite set of s cenarios. The realized yields from this schedule then become an input to the second stage, driving the reactive decisions on procurement, salvage, and inventory management. Specifically, second-stage recourse variables determining the inventory policy are $v_{rt}(\omega)$ (1 if type-r modules are procured in period t, 0 otherwise). In period t of scenario ω, $q^+_{rt}(\omega)$ and $q^-_{rt}(\omega)$ respectively are the number of procured and salvaged modules of type r and $I_{rt}(\omega)$ is their inventory level at the end of the period.

On machine k, disassembly of job j requires a processing time of p_{jk}, and the setup time between jobs i and j is sequence-dependent, denoted by s_{ijk} for all $i, j \in \mathcal{J}_t$.

Moreover, $H_{\max}$ represents the maximum capacity in hand on each machine during period t. There is a cost c_{jk} for assigning core j to machine k, a cost c^{time} per unit makespan time, an inventory holding cost h_r per period and per module of type r, and a salvage price c^-_{rt} per type-r module in period t. Procurement incurs a fixed cost F_t for placing an order in period t, plus a variable cost c^+_{rt} per type-r module. Denote a schedule as $\mathbf{x} = (x_{jkt}, y_{ijkt}; i, j \in \mathcal{J}_t, k \in \mathcal{M}, t \in \mathcal{T})$, then the first-stage formulation becomes

$$
\begin{aligned}
\min_{\mathbf{x}} \quad & \sum_{t\in\mathcal{T}} \Big(c^{\text{time}} C_{\max,t} + \sum_{j\in\mathcal{J}_t}\sum_{k\in\mathcal{M}} c_{jk} x_{jkt} \Big) + \mathbb{E}_\omega[Q(\mathbf{x},\omega)] && \text{(MP)}\\
\text{s.t.} \quad & \sum_{k\in\mathcal{M}} x_{jkt} = 1 && \forall t \in \mathcal{T}, j \in \mathcal{J}_t\,,\\
& \sum_{i\in\mathcal{J}_{0t}, i\neq j} y_{ijkt} = x_{jkt} && \forall t \in \mathcal{T}, k \in \mathcal{M}, j \in \mathcal{J}_t\,,\\
& \sum_{j\in\mathcal{J}_{0t}, j\neq i} y_{ijkt} = x_{ikt} && \forall t \in \mathcal{T}, k \in \mathcal{M}, i \in \mathcal{J}_t\,,\\
& \sum_{j\in\mathcal{J}_t} y_{0jkt} \leqslant 1 && \forall t \in \mathcal{T}, k \in \mathcal{M}\,,\\
& C_{jt} \geqslant C_{it} + s_{ijk} + p_{jk} - M(1 - y_{ijkt}) && \forall t, k, i \in \mathcal{J}_{0t}, j \in \mathcal{J}_t, i \neq j\,,\\
& C_{\max,t} \geqslant C_{jt} && \forall t \in \mathcal{T}, j \in \mathcal{J}_t\,,\\
& C_{\max,t} \leqslant H_{\max} && \forall t \in \mathcal{T}\,,\\
& x_{jkt}, y_{ijkt} \in \{0,1\},\ C_{jt}, C_{\max,t} \geq 0 && \forall t \in \mathcal{T}, k \in \mathcal{M}, i, j \in \mathcal{J}_t\,,
\end{aligned}
$$

where $\mathcal{J}_{0t} = \mathcal{J}_t \cup \{0\}$ and M is a big enough positive number. For a fixed schedule $\mathbf{x}$ and scenario ω, the second-stage recourse $Q(\mathbf{x},\omega)$ is

$$
\begin{aligned}
Q(\mathbf{x},\omega) = \min & \sum_{t\in\mathcal{T}} \Big(F_t z_t(\omega) + \sum_{r\in\mathcal{R}} \big(c^+_{rt} q^+_{rt}(\omega) - c^-_{rt} q^-_{rt}(\omega) + h_r I_{rt}(\omega) \big) \Big)\\
\text{s.t.} \quad & \sum_{j,k} \xi_{jkrt}(\omega) x_{jkt} + I_{r,t-1}(\omega) + q^+_{rt}(\omega) - q^-_{rt}(\omega) && \forall r \in \mathcal{N}, t \in \mathcal{T}\,,\\
& \qquad = D_{rt} + I_{rt}(\omega)\\
& q^+_{rt}(\omega) \leqslant U_{rt} v_{rt}(\omega) && \forall r \in \mathcal{N}, t \in \mathcal{T}\,,\\
& v_{rt}(\omega) \leqslant z_t(\omega) && \forall r \in \mathcal{N}, t \in \mathcal{T}\,, \qquad (1)\\
& z_t(\omega), v_{rt}(\omega) \in \{0,1\},\ q^\pm_{rt}(\omega) \in \mathbb{N},\ I_{rt}(\omega) \in \mathbb{N} && \forall r \in \mathcal{N}, t \in \mathcal{T}\,,
\end{aligned}
$$

with initial inventory $I_{r0} = 0$ for all types, and where U_{rt} is a sufficiently large upper bound on procurement of type-r in t, e.g. $U_{rt} = \sum_{t'=t}^{T} D_{rt'}$. The linking constraint (1) enforces the fixed cost F_t if any type-r module is procured in period t.

3 Solution Methodologies

Our primary approach is a Simulated Annealing (SA) heuristic inspired by the structure of Logic-Based Benders Decomposition (LBBD) [2]. In this framework, the SA serves as a heuristic master problem to explore scheduling decisions. For each candidate schedule, we solve the second-stage recourse subproblem to optimality using an MIP solver. We also directly solve the extensive form of the formulation above, in which the expectation in the MP objective is replaced by an average over $|\Omega|$ randomly generated yield scenarios. The resulting cost is then used to guide the SA search and update its solution.

4 Computational Study and Conclusion

In our experiments, we consider a planning horizon of $T = 3$ with $|\mathcal{J}_t| = 6$, $|\mathcal{M}| = 2$, $|\mathcal{R}| = 3$, and $|\Omega| = 10$, under a CPU time limit of 60 s. The extensive MIP formulation timed out, yielding a best incumbent solution of 1294.28 and a lower bound of 1222.27. Our SA heuristic reached a solution cost of 1303.40, whereas the deterministic expected-value model (EVM) schedule cost 1312.12. This corresponds to a value of the stochastic solution (VSS) of 17.83, thereby demonstrating the benefit of the stochastic model.

We formulated a two-stage stochastic program for disassembly and inventory under yield uncertainty, and showed that an SA heuristic delivers near-optimal solutions within a practical time. Future work will develop multi-stage formulations to relax the implicit perfect-foresight in our second stage and incorporate adaptive updates of yield distributions.

References

[1] Suzanne, E., Absi, N., and Borodin, V., Towards circular economy in production planning: challenges and opportunities, *European Journal of Operational Research*, 287(1), 168-190 (2020)

[2] Hooker, J. N., Planning and scheduling by Logic-Based Benders Decomposition, *Operations Research*, 55(3), 588-602 (2007)

[3] Liu, K. and Zhang, Z.-H., Capacitated disassembly scheduling under stochastic yield and demand, *European Journal of Operational Research*, 269(1), 244-257 (2018)

Lot-sizing under decision-dependent uncertainty: A probing-enhanced stochastic programming approach

Franco Quezada
University of Santiago of Chile, Industrial Engineering Department, Chile
franco.quezada@usach

Céline Gicquel
Université Paris-Saclay, LRI, F-91190 Gif-sur-Yvette, France
gicquel@lisn.fr

Safia Kedad-Sidhoum
Conservatoire national des arts et métiers, CEDRIC, F-75003 Paris, France
kedad_sidhoum@cnam.fr

Bernardo Pagnoncelli
SKEMA Business School, Université Côte d'Azur, Lille, France
bernardo.pagnoncelli@skema.edu

Abstract

We address the multi-item capacitated lot-sizing problem under decision-dependent uncertainty via a probing-enhanced stochastic programming framework. Demand is correlated with another random vector, and the decision-maker can acquire partial information by probing components of this vector, conditioning decisions on observed covariates. This generalizes classical models by embedding information acquisition into a three-stage framework. We propose a compact reformulation that removes non-anticipativity constraints, yielding stronger relaxations and better tractability. We extend classical inequalities and introduce value-function cuts that capture the link between probing and recourse costs. These are embedded in a branch-and-cut algorithm with a primal heuristic. Results show our method outperforms off-the-shelf solver, reducing optimality gaps by up to 85%, and achieving gaps below 1.5% on average. Results highlight the importance of structured reformulation, valid inequalities, and heuristics in solving decision-dependent stochastic programs.

1 Introduction

Lot-sizing is a core problem in production and inventory management, focused on determining optimal production or ordering quantities over time to meet demand at minimal cost. While traditional models often assume known or exogenously uncertain demand, stochastic programming approaches have been developed to address uncertainty more realistically by incorporating probabilistic demand distributions [1, 2]. However, these models typically treat uncertainty as fixed and uncontrollable, overlooking the possibility of actively reducing it through information acquisition.

We consider a classical Multi-Item Capacitated Lot-Sizing Problem (MCLSP) under demand uncertainty, where a single production resource is used to manufacture multiple items over a finite planning horizon. Each item faces uncertain demand in each period, modeled as a random variable with a known probability distribution. The planner must decide in advance which items to produce and in what quantities, while respecting production capacity limits and incurring setup, production, holding, and lost sales costs. The objective is to minimize the total expected cost, accounting for both early-stage decisions and later adjustments once demand is realized.

In many real-world settings, firms can influence the accuracy of demand forecasts by strategically observing information such as consumer behavior, demand trends, pricing, or promotional data at a cost. Such practical considerations motivate decision-dependent uncertainty models, particularly the probing-enhanced stochastic programming framework [3], where early decisions determine what information should be acquired. We apply this framework to the stochastic MCLSP, allowing planners to selectively probe external signals (covariates) that are statistically correlated with future demand. The resulting model follows a three-stage structure: first, the planner chooses which covariates to probe; second, setup and production decisions are made based on the observed information; and third, after demand is fully revealed, the planner adjusts through inventory or lost sales decisions. This approach is especially valuable in environments with high uncertainty and limited resources for information gathering, enabling more informed and cost-effective production planning.

2 Branch-and-cut-based decomposition approach

We develop a scalable branch-and-cut-based decomposition algorithm for solving the probing-enhanced MCLSP, built upon three key ingredients. First, we introduce a novel reformulation (ALF) that improves numerical stability and yields stronger relaxations by avoiding classical non-anticipativity constraints in the representation of decision-dependent uncertainty (NAF). Second, we extend the classical (k, U) inequalities to incorporate probing decisions, linking setup and production variables with information acquisition actions. These inequalities strengthen the formulation and improve the convergence of the algorithm. Third, we propose a new class of value-function-based inequalities derived from item-level lower bounds, which capture the cost impact of probing decisions on recourse actions. These components are integrated into a decomposition framework that exploits the structure of single-item subproblems. The formulation serves as the master problem, while the subproblems are used to generate both primal bounds and valid inequalities. As there are an exponential number of the extended (k, U) inequalities, we added them dynamically to the formulation to strengthen its linear relaxation. The value-function-based inequalities are added a priori to the formulation based on the solution of single-item subproblems. Together, these elements enable an efficient and flexible branch-and-cut algorithm capable of solving medium-scale instances with high-dimensional uncertainty.

3 Results

Table 1 summarizes the numerical results. Columns $|\Omega|$ and $|\Gamma|$ indicate the number of scenarios used for the random variables, while $|\mathcal{I}|$ denotes the number of items. Column `Gap` shows the relative optimality gap, computed as `Gap` $= |UB - LB|/UB$, where UB and LB are the upper and lower bounds, respectively. Column `Time (s)` reports the average total computation time in seconds. Columns #`VI` and #`Nodes` display the total number of valid inequalities generated and the number of branch-and-bound nodes explored.

Instance				NAF				ALF			
$\|\Omega\|$	$\|\Gamma\|$	$\|\mathcal{I}\|$	Method	Gap (%)	Time (s)	#VI	#Nodes	Gap (%)	Time (s)	#VI	# Nodes
10	5	5	CPX	0.92	648.12	0	3590	0.70	613.31	0	4895
			B&C	1.61	824.91	2937	3235	1.69	805.35	2801	2619
		10	CPX	9.77	844.43	0	8032	8.03	868.18	0	7329
			B&C	1.69	958.47	4231	1310	1.68	958.00	3076	967
	10	5	CPX	11.40	766.76	0	4447	5.59	804.18	0	11968
			B&C	1.86	957.10	4629	1961	2.00	956.85	2946	521
		10	CPX	30.40	895.72	0	3814	14.06	902.76	0	3818
			B&C	4.90	866.69	4433	29	3.39	859.65	3254	0

Table 1: Performance of Branch-and-Cut algorithm .

The results in Table 1 show that the proposed ALF consistently outperforms the standard NAF, both when solved directly with CPLEX and within the branch-and-cut framework. Across all instance sizes, ALF achieves lower optimality gaps, particularly in large-scale or high-scenario settings. For example, in the largest configuration ($|\Gamma| = 10$, $|\mathcal{I}| = 10$), ALF reduces the gap from 30.40% (NAF) to 14.06% under CPLEX, and further to just 3.39% with the branch-and-cut algorithm. The branch-and-cut algorithm also provides significant improvements over CPLEX alone, especially for NAF, where it reduces the gap from 11.40% to 1.86% for medium-sized

$\sigma = 30$						$\sigma = 60$					
$\varrho = -0.5$		$\varrho = -0.7$		$\varrho = -0.9$		$\varrho = -0.5$		$\varrho = -0.7$		$\varrho = -0.9$	
$\texttt{VPS}_{gap}$	$\texttt{EVPPI}_{gap}$	$\texttt{VPS}_{gap}$	$\texttt{EVPPI}_{gap}$	$\texttt{VPS}_{gap}$	$\texttt{EVPPI}_{gap}$	$\texttt{VPS}_{gap}$	$\texttt{EVPPI}_{gap}$	$\texttt{VPS}_{gap}$	$\texttt{EVPPI}_{gap}$	$\texttt{VPS}_{gap}$	$\texttt{EVPPI}_{gap}$
3.95	0.51	6.43	0.98	10.15	1.30	8.89	0.84	15.12	1.39	22.29	2.18

Table 2: Value of probing solution for different correlation coefficient ϱ and standard deviation σ.

instances ($|\Gamma| = 10$, $|\mathcal{I}| = 5$). ALF benefits similarly, achieving gaps below 2% in most configurations when combined with branch-and-cut.

To evaluate the benefits of information acquisition, we define two metrics: the Value of the Probing Solution (VPS) and the Expected Value of Perfect Probing Information (EVPPI), analogous to VSS and EVPI in stochastic programming. VPS captures the gain from integrating costly probing, while EVPPI reflects the additional benefit if all information were freely available. Results show that higher demand-covariate correlation significantly increases VPS, with relative improvements exceeding 45% under high penalty and variability. Likewise, higher lost-sales penalties boost the value of probing, confirming its importance in high-risk settings. Across all settings, EVPPI remains modest (typically $< 5\%$), indicating that most value is already captured through selective, costly probing. This confirms the efficiency and practicality of the proposed approach.

Acknowledgment. This research benefited from ANID, Beca Postdoctorado en el extranjero, Becas Chile, Folio 74230025 and the support of the FMJH Program Gaspard Monge for optimization and operations research and their interactions with data science, and the Energy4Climate Interdisciplinary Center (E4C) of IP Paris and Ecole des Ponts ParisTech, which is in part supported by 3rd Programme d'Investissements d'Avenir [ANR-18-EUR-0006-02]. The authors also acknowledge the financial support from the Programa de Cooperación Científica ECOS-ANID project number ECOS230013 and the Programa Regional STIC-AMSUD project number AMSUD240008.

References

[1] J.H. Bookbinder and J.-Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34(9):1096–1108, 1988.

[2] Horst Tempelmeier. Stochastic lot sizing problems. In *Handbook of Stochastic Models and Analysis of Manufacturing System Operations*, pages 313–344. Springer, 2013.

[3] Zhichao Ma, Youngdae Kim, Jeff Linderoth, James R Luedtke, and Logan R Matthews. Probing-enhanced stochastic programming. *arXiv preprint arXiv:2407.10669*, 2024.

The Production Routing Problem with Stochastic Demand and Service Levels

Ali Kermani
HEC Montréal, CIRRELT, and GERAD, H3T 2A7, Canada
ali.kermani@hec.ca

Jean-François Cordeau
HEC Montréal, CIRRELT, and GERAD, H3T 2A7, Canada
jean-francois.cordeau@hec.ca

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

Abstract

The Production Routing Problem is an integrated problem that simultaneously considers lot sizing, inventory and routing decisions. We consider the context in which demand is uncertain while certain service levels have to be satisfied. The problem is modeled as a two-stage stochastic problem using a scenario approach in which the production setup decisions are made in the first stage, while the other decisions, such as the production and delivery quantities and the routing, are made in the second stage. This leads to a highly complex problem for which we develop a matheuristic. We iteratively solve three distinct subproblems. The first subproblem aims to find a good production setup plan. The second subproblem determines production and delivery quantities, while the third subproblem determines the routing decisions. The computational experiments show the effectiveness of this heuristic to find high-quality solutions when considering various types of service levels.

1 Introduction

The Production Routing Problem (PRP) integrates production, inventory, and routing decisions to optimize supply chain operations over a finite horizon. While the classical deterministic PRP assumes known demand, real-world settings often involve uncertainty. Early stochastic models addressed this by fixing routes in advance, but this limits flexibility [1]. To overcome this, we introduced the Stochastic PRP with

Adaptive Routing (SPRP-AR), where routing decisions are made after demand realization, improving cost efficiency by up to 6.5% [3].

Another important aspect under uncertainty is ensuring service levels. While various service level metrics have been studied in the lot-sizing literature, they primarily address production and inventory planning [2, 4]. In contrast, the PRP introduces the added complexity of routing decisions. Despite its practical importance, the integration of service level constraints into the PRP framework remains largely unexplored. This study is the first to incorporate four service level measures: α (probability of no stockout), β (fill rate), γ (expected backlog to average demand ratio), and δ (expected backlog to maximum backlog ratio), into the SPRP-AR framework. Each is applied at different levels of granularity (per customer or aggregated, per period or over the whole horizon), offering flexibility in aligning with diverse operational needs. This strategy is particularly valuable in settings where outsourcing is not feasible.

The study proposes a novel two-stage stochastic model with adaptive routing and service level constraints, following a static-dynamic strategy where setup decisions are made in the first stage while adapting other decisions based on realized demand. An iterative matheuristic (IMH) algorithm is developed to solve the model in three phases, gradually refining setup, inventory, and routing decisions. Computational results on benchmark instances demonstrate the value of service level integration, adaptive routing, and the impact of the service level granularity, offering practical insights for supply chain decision-makers facing uncertainty.

2 Problem Definition

We propose a two-stage formulation for the SPRP-AR, where a single production facility serves multiple customers over a finite planning horizon under uncertain demand. A fleet of homogeneous vehicles delivers products to customers, and both production and transportation decisions must satisfy capacity and inventory constraints. In this setting, the first-stage decisions involve determining production setups, while second-stage decisions, including production quantities, deliveries, and routing, are made after the actual demand is realized. To model demand uncertainty, we use a scenario-based approach where each scenario represents a possible realization of customer demand. Products can be stored at either the plant or customer locations, and any unmet demand may be backlogged and fulfilled later.

To ensure customer satisfaction under uncertainty, we incorporate four distinct service level constraints into the model. These include: (1) the α service level, which limits the probability of stockouts; (2) the β fill rate, which controls the proportion of demand met immediately; (3) the γ level, which constrains the average backlog to average demand ratio; and (4) the δ level, which limits the expected backlog relative to its maximum. These service levels are applied with varying levels of granularity, either per customer or aggregated across all customers, and either per period or over the entire planning horizon. Their inclusion directly influences how inventory and delivery decisions are structured, enabling the model to reflect diverse service expectations and operational policies.

3 Iterative Matheuristic Algorithm

To solve the SPRP-AR efficiently, we propose an IMH that decomposes the problem into three subproblems. The first phase solves a simplified two-level lot-sizing problem with a single dummy customer to quickly generate feasible production setup plans. This abstraction helps reduce the complexity of the original problem and serves as a foundation for more detailed planning in later phases. In the second phase, given the fixed setup plan, we solve a restricted PRP with a single vehicle and aggregated capacity to determine production quantities and delivery amounts. At this stage, routing decisions are not yet considered, and the focus is on meeting demand while ensuring service level feasibility.

In the third phase, we incorporate routing by solving a deterministic multi-vehicle PRP separately for each scenario. This step refines delivery routes while respecting service level constraints and previously established decisions. The algorithm alternates between two iterative loops: an outer loop for diversifying setup plans, and an inner loop for intensifying solutions by refining production and routing based on updated visit costs. Once feasibility is established, the second phase is updated to include approximate routing costs, improving solution quality. This structured, scenario-based approach enables efficient handling of uncertainty, adaptive routing, and service level measures.

4 Results

We compared the average objective function values for different service level types across various granularity levels. Figure 1 provides the results for the β service level as an example. The results show that the plant level-global (β^{plant}) configuration yields the lowest costs, due to its high flexibility. As service levels are imposed from the plant's perspective over the entire horizon, demand imbalances across customers and periods are absorbed more easily. The customer level-global ($\beta^{customer}$) case follows, with higher costs resulting from stricter service enforcement for each customer across the entire horizon. Interestingly, this configuration still has lower costs than the plant level-single period (β_c^{plant}) case, where service levels are applied per period but aggregated across customers.

In general, imposing service levels for each period leads to higher costs compared to global-level enforcement, as it requires more consistent demand satisfaction over

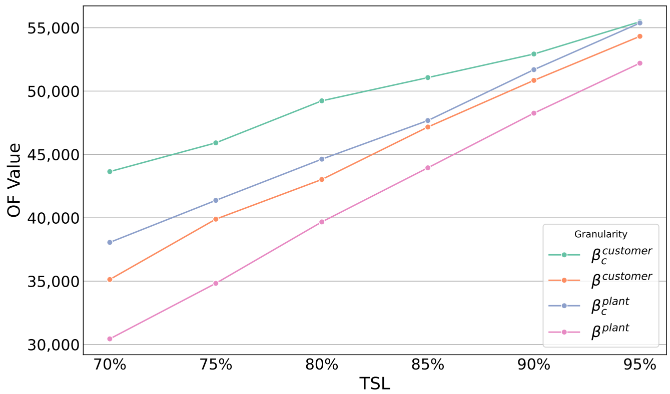

Figure 1: Objective function values for β service level for different granularity levels

time. Among per-period configurations, customer-level ($\beta_c^{customer}$) setting results in the highest costs, compared to service levels aggregated over all customers. However, the cost gap between plant and customer granularity levels narrows as the target service level (TSL) increases. Overall, enforcing service levels at the customer level increases costs due to stricter fulfillment requirements, while plant-level case offers greater flexibility. These findings are consistent across multiple settings with varying numbers of scenarios and uncertainty levels, highlighting the importance of selecting an appropriate granularity based on operational priorities.

References

[1] Adulyasak, Yossiri and Cordeau, Jean-François and Jans, Raf, Benders Decomposition for Production Routing Under Demand Uncertainty, Operations Research, 63, 4, 851-867 (2015)

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

[3] Kermani, Ali and Cordeau, Jean-François and Jans, Raf, A progressive hedging-based matheuristic for the stochastic production routing problem with adaptive routing, Computers & Operations Research, 169, 106745 (2024)

[4] Tempelmeier, Horst, Stochastic Lot Sizing Problems, 192, 313–344, Springer, (2013)

A simple heuristic for computing non-stationary inventory policies

Onur A. Kilic
University of Groningen, Groningen, The Netherlands
o.a.kilic@rug.nl

S. Armagan Tarim
Hacettepe University, Ankara, Turkey
armagan.tarim@hacettepe.edu.tr

Abstract

We consider a finite-horizon periodic-review inventory system with fixed replenishment costs that faces non-stationary demands. The structure of the optimal control policy for this system has long been known. However, finding optimal policy parameters requires solving a large-scale stochastic dynamic program. To circumvent this, we devise a recursion-free approximation for the cost function of the problem. This translates into an efficient and effective heuristic to compute policy parameters that significantly outperforms earlier heuristics. Our approach is easy-to-understand and easy-to-use as it follows by elementary methods of shortest paths and convex minimization.

1 Introduction

Today, industries are experiencing non-stationary (stochastic and time-varying) demand more frequently as product life cycles are getting increasingly shorter in response to fast technological progress and rapid changes in consumer preferences (Simchi-Levi et al., 2003; Chopra and Meindl, 2007). When a product life cycle spans a short period of time, the magnitude of non-stationarity becomes drastic because the demand rate changes rapidly as a product moves from one phase of the life cycle to another. Also, in most environments, demand is often heavily seasonal and has a significant trend. Therefore, the demand rate also changes within the phases of the product life cycle. The conventional methods of inventory control that are tailored for stationary demands have very limited applicability in such environments. Hence, firms must employ alternative methods to effectively match their supply to non-stationary demand (Kurawarwala and Matsuo, 1996; Graves and Willems, 2008).

Managing inventories is challenging when demand is non-stationary. That is because fluctuations in demand must be reflected in replenishments. That is non-stationary demand compels non-stationary inventory control. The challenge further intensifies when replenishments require fixed costs, as is often the case for companies with distant suppliers. In such systems, decision-makers are exposed to an inventory control problem where the non-stationarity of demand affects not only the size but also the timing of replenishments. This also manifests itself in the mathematical models of such problems where the associated cost functions (even under the mildest of assumptions) are time-dependent and non-convex. It is indeed widely accepted that mathematical models of inventory problems with non-stationary demands are complicated in terms of computational needs and user understanding (Silver, 1981, 2008).

This paper aims to address the aforementioned challenges in the context of a finite-horizon periodic-review inventory system with fixed replenishment costs that faces non-stationary demands. The relevance of the system is evident as it appears in many retail, wholesale, and industrial environments. The associated inventory control problem is well-addressed in the literature. The optimal control policy is known to be a non-stationary (s, S) policy (Scarf, 1960). However, finding optimal policy parameters is demanding as it requires constructing the optimal cost function recursively for each and every period in the planning horizon by solving a stochastic dynamic program. This is viable when in case of discrete demands (see e.g. Bollapragada and Morton, 1999; Lulli and Sen, 2004), but, even then, it is computationally expensive as one needs to consider all possible demand realizations and inventory levels in each period—both of which could be arbitrarily large in number. To that end, we propose an approximation for the non-convex optimal cost function. Our approximate cost function has two important properties. First, it is recursion-free. That is, it can be obtained without resorting to a stochastic dynamic program. Second, it is defined as the pointwise minimum of only a few convex functions. Hence, while it has multiple local minima like the optimal cost function, its minimizer and level sets can easily be obtained. These properties enable us to translate the approximate cost function into an efficient and effective heuristic to compute policy parameters. The heuristic proceeds with a simple computational procedure wherein the policy parameters for all periods in the planning horizon can be obtained by solving a single deterministic shortest path problem and a series of root-finding problems. We further accelerate this computational procedure by means of several algorithmic refinements. These make it possible to handle problem instances with hundreds of periods in a matter of seconds. We numerically compare our heuristic against the heuristics by Askin (1981) and Bollapragada and Morton (1999) as well as the optimal stochastic dynamic program. The results show that it significantly outperforms the earlier heuristics and yields almost-optimal results for a variety of demand patterns and cost parameters. Our approach is very accessible as it builds on readily available methods of shortest paths and convex minimization. Thus it is appealing for both practical and educational purposes.

2 Relevant literature

The characterization of optimal control policies for inventory problems has always been a center of interest in the inventory management literature. In this context, Scarf's (1960) proof of the optimality of (s, S) policies is of particular importance. Scarf considered the finite-horizon dynamic inventory problem with stationary demands and showed that the optimal control decision for each period n can be characterized by two critical values (s_n, S_n). That is, a replenishment order should be issued only if the inventory level is below s_n and if so the order quantity should increase the inventory level to S_n. This seminal result paved the way for a large number of studies. Iglehart (1963) showed that the optimal policy converges to a stationary (s, S) policy if the planning horizon is sufficiently long. This is followed by a variety of studies on exact and approximate methods to compute the optimal stationary (s, S) policies (see e.g. Veinott and Wagner, 1965; Wagner et al., 1965; Johnson, 1968; Sivazlian, 1971; Archibald and Silver, 1978; Ehrhardt, 1979; Porteus, 1979; Federgruen and Zipkin, 1984; Zheng and Federgruen, 1991; Feng and Xiao, 2000).

Scarf's (1960) proof immediately applies to systems with non-stationary demands. However, while the structure of the optimal policy is known, finding optimal non-stationary policy parameters is challenging as it requires constructing an optimal cost function for each period recursively by solving a stochastic dynamic program. This is viable when demands and hence inventory levels are integer-valued (see e.g. Bollapragada and Morton, 1999; Lulli and Sen, 2004). However, the difficulty prevails as it requires evaluating all possible inventory levels and demand realizations in each period, which could be arbitrarily large in number. Indeed, such a numerical procedure is complex and computationally expensive to be implemented in practice (see e.g. Bollapragada and Morton, 1999; Silver, 1981, 2008; Neale and Willems, 2009).

Notwithstanding, not much work has been done on the non-stationary inventory control problem as compared to its stationary counterpart. Silver (1978) and Askin (1981) developed heuristics that determine policy parameters by using the least expected cost per period criterion. Bollapragada and Morton (1999) proposed a conceptually different approach where the non-stationary problem is approximated by a series of stationary problems. These stationary problems are constructed by averaging demands over a number of consecutive periods based on an expected cycle time, and the policy parameters of each stationary problem are computed using stationary analysis. Xiang et al. (2018) presented an approach where policy parameters are approximated by iteratively solving a series of mixed integer non-linear programs for each period of the planning horizon. While effective, this method is computationally too expensive for practical use. In the current study, we follow the aforementioned line of research and develop a new approach to compute policy parameters heuristically, based on approximating the non-convex cost function by a sequence of convex functions.

References

Archibald, B. C. and Silver, E. A. (1978). (s, S) policies under continuous review and discrete compound poisson demand. *Management Science*, 24(9):899–909.

Askin, R. (1981). A procedure for production lot sizing with probabilistic dynamic demand. *IIE Transactions*, 13(2):132–137.

Bollapragada, S. and Morton, T. (1999). A simple heuristic for computing non-stationary (s, S) policies. *Operations Research*, 47(4):576–584.

Chopra, S. and Meindl, P. (2007). *Supply Chain Management. Strategy, Planning & Operation*. Prentice Hall, 3rd edition.

Ehrhardt, R. (1979). Power approximation for computing (s, S) inventory policies. *Management Science*, 25(8):777–786.

Federgruen, A. and Zipkin, P. (1984). An efficient algorithm for computing optimal (s, S) policies. *Operations Research*, 32(6):1268–1285.

Feng, Y. and Xiao, B. (2000). A new algorithm for computing optimal (s, S) policies in a stochastic single item/location inventory system. *IIE Transactions*, 32(11):1081–1090.

Graves, S. C. and Willems, S. P. (2008). Strategic inventory placement in supply chains: Nonstationary demand. *Manufacturing & Service Operations Management*, 10(2):278–287.

Iglehart, D. L. (1963). Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Science*, 9(2):259–267.

Johnson, E. L. (1968). On (s, S) policies. *Management Science*, 15(1):80–101.

Kurawarwala, A. A. and Matsuo, H. (1996). Forecasting and inventory management of short life-cycle products. *Operations Research*, 44(1):131–150.

Lulli, G. and Sen, S. (2004). A branch-and-price algorithm for multistage stochastic integer programming with application to stochastic batch-sizing problems. *Management Science*, 50(6):786.

Neale, J. J. and Willems, S. P. (2009). Managing inventory in supply chains with nonstationary demand. *Interfaces*, 39(5):388–399.

Porteus, E. L. (1979). An adjustment to the Norman-White approach to approximating dynamic programs. *Operations Research*, 27(6):1203–1208.

Scarf, H. E. (1960). Optimality of (S, s) policies in the dynamic inventory problem. In Arrow, K. J., Karlin, S., and Suppes, P., editors, *Mathematical Methods in the Social Sciences*, pages 196–202. Stanford University Press, Stanford, CA.

Silver, E. A. (1978). Inventory control under a probabilistic time-varying, demand pattern. *IIE Transactions*, 10(4):371–379.

Silver, E. A. (1981). Operations research in inventory management: A review and critique. *Operations Research*, 29(4):628–645.

Silver, E. A. (2008). Inventory management: An overview, Canadian publications, practical applications and suggestions for future research. *INFOR*, 46:15 – 28.

Simchi-Levi, D., Kaminsky, P., and Simchi-Levi, E. (2003). *Designing and Managing the Supply Chain–Concepts, Strategies and Case Studies*. McGraw Hill, 2nd edition.

Sivazlian, B. (1971). Dimensional and computational analysis in stationary (s, S) inventory problems with gamma distributed demand. *Management Science*, 17(6):307–311.

Veinott, A. F. and Wagner, H. M. (1965). Computing optimal (s, S) inventory policies. *Management Science*, 11(5):525–552.

Wagner, H. M., O'Hagan, M., and Lundh, B. (1965). An empirical study of exactly and approximately optimal inventory policies. *Management Science*, 11(7):690–723.

Xiang, M., Rossi, R., Martin-Barragan, B., and Tarim, S. A. (2018). Computing non-stationary (s, S) policies using mixed integer linear programming. *European Journal of Operational Research*, 271(2):490–500.

Zheng, Y. S. and Federgruen, A. (1991). Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Operations Research*, 39(4):654–665.