14th International Conference on Stochastic Models of Manufacturing and Service Operations

June 9 – 14, 2024, Naples, Italy

SMMSO 2024

Conference Proceedings

Contents

Committees

Organising Committee

- Giuseppe Converso (University of Naples Federico II, Italy)
- Mosé Gallo (University of Naples Federico II, Italy)
- Elisa Gebennini (Universitas Mercatorum, Italy)
- Andrea Grassi (University of Naples Federico II, Italy)
- Guido Guizzi (University of Naples Federico II, Italy)
- Teresa Murino (University of Naples Federico II, Italy)
- Valentina Popolo (Università Telematica Pegaso, Italy)
- Liberatina C. Santillo (University of Naples Federico II, Italy)
- Silvestro Vespoli (University of Naples Federico II, Italy)

Scientific Committee

- J. Buzacott (York University, Canada)
- M. Di Mascolo (Laboratoire G-SCOP, France)
- K. Furmans (KIT, Germany)
- S. B. Gershwin (MIT, USA)
- S. Helber (Leibniz University Hannover, Germany)
- Y. Jang (KAIST, South Korea)
- F. Karaesmen (Koç University, Turkey)
- G. Liberopoulos (University of Thessaly, Greece)
- J. MacGregor Smith (University of Massachusetts, Amherst, USA)
- A. Matta (Politecnico di Milano, Italy)
- S. Meerkov (University of Michigan, Ann Arbor, USA)
- C. Papadopoulos (Aristotle University of Thessaloniki, Greece)
- R. Stolletz (University of Mannheim, Germany)
- B. Tan (Özyeğin University, Turkey)
- H. Tempelmeier (University of Cologne, Germany)
- T. Tolio (Politecnico di Milano, Italy)

Welcome message

Dear Participants,

it is with great pleasure that we welcome you to the 14th International Conference on Stochastic Models of Manufacturing and Service Operations, SMMSO 2024, organized by the University of Naples Federico II.

This conference marks the fourteenth in a successful series of biennial conferences that began in 1997 across Europe. Initially held in Greece (1997, 1999, 2001, 2003, and 2005), the conference has since been hosted by various other countries: the Netherlands in 2007, Italy in 2009, Turkey in 2011, Germany in 2013, before returning to Greece in 2015, Italy in 2017, Germany in 2019, and France in 2022. This year, we are honored to welcome you back to Italy, where we will gather along the beautiful Neapolitan coast, close to the stunning sea.

We have prepared a comprehensive program that bridges research and practice in manufacturing and service operations. The conference will feature one keynote and 25 presentations, organized into 9 sessions. We are delighted to have Prof. Stefan Helber from Leibniz Universität Hannover as our keynote speaker and Eng. Giovanni Abete of A.Abete srl as our industrial speaker. We extend our heartfelt thanks to them and all the authors for their invaluable contributions to the success of this conference.

We would like to thank our sponsors: the University of Naples Federico II; the Department of Chemical, Materials and Industrial Production Engineering; PSB srl; Nemesi Project as supporting project. Special thanks go to the members of the Scientific Committee for their support in organizing this conference.

Our deepest gratitude goes to Christos Papadopoulos, who was able to create this series of conferences, and to the organizers of previous editions who ensured its continuation in time. We also remember James MacGregor Smith, a dedicated supporter of this conference from its inception, who sadly passed away recently.

Lastly, we extend our sincere thanks to all participants and their accompanying persons for attending SMMSO 2024. We hope you enjoy your stay in Naples and the activities we have arranged for you, and that you leave with enriching scientific discussions and friendly exchanges with other members of the SMMSO community.

Sincerely,

The SMMSO 2024 Organizing Committee

Proceedings

Keynote

Prof. Stefan Helber, Leibniz Universität Hannover

Title: Some thoughts on machine learning, performance evaluation and optimization

Abstract: Machine learning, as a sub-area of artificial intelligence, can be seen as the approach to distill and statistically describe underlying structures behind large data sets of well-understood observations with the overarching intention to support some type of decision-making. Since digital and connected stochastic manufacturing and service systems nowadays create such large data sets, the question arises how and to which extent machine learning can or cannot easily be used to better understand the underlying ”mechanics” of those systems and to facilitate their design and operation. This requires some sort of digital twin of the considered system or process. However, while the general public may see machine learning (or artificial intelligence at large) as the final way to solve all kinds of formally untractable problems, at least my own experience is, to date, less positive and clear-cut. In the talk, I will try to shed some light on the opportunities and pitfalls of using machine learning in the context of performance evaluation and optimization, in particular regarding stochastic manufacturing and service systems.

A Digital Twin Based Decision Support System for the Management of an Operating Room

Leah RIFI
Industrial Engineering Center, IMT Mines Albi, Albi, Tarn, France, email: leah.rifi@mines-albi.fr

Canan PEHLIVAN
Industrial Engineering Center, IMT Mines Albi, Albi, Tarn, France, email: canan.pehlivan@mines-albi.fr

Cléa MARTINEZ
Industrial Engineering Center, IMT Mines Albi, Albi, Tarn, France, email: clea.martinez@mines-albi.fr

Maria DI MASCOLO
Univ. Grenoble Alpes, CNRS, Grenoble INP, G-SCOP, 38000 Grenoble, France, email: maria.di-mascolo@grenoble-inp.fr

Franck FONTANILI
Industrial Engineering Center, IMT Mines Albi, Albi, Tarn, France, email: franck.fontanili@mines-albi.fr

With healthcare demand rising worldwide, hospital services are increasingly needed. Hospitals' performance is tightly linked to their surgical suite performance, which makes it necessary for surgical suites to be efficient. In this paper, we focus on the operating room schedule execution and related decision-making. We thus propose a digital twin-based decision support system for the prospective and retrospective simulation and analysis of the operating room schedule execution. We describe the developed prototype (inputs/output/parameters/modeling) and its functionalities and show its application to an operating room inspired by a real case.

Keywords: Operating Room, Digital Twin, OR Management, Decision Support System, Modeling and Simulation, Uncertainties

1 Context and problematic

As healthcare demand escalates globally, the necessity for medical services grows, and the efficiency of hospitals' surgical suites becomes pivotal. These suites not only represent a significant portion of hospital budgets, with 40%-50% allocated to them, but they also constitute around 30% of overall healthcare costs (Kaye et al., 2020; Macario et al., 1997). Thus, ensuring the effectiveness and profitability of operating rooms (OR) is imperative. However, the management of a surgical suite is complex due to several factors: the **variety of patient pathways** (outpatient and inpatient, elective and non-elective patients), the **diversity of professionals** working within the suite (surgeons, anesthesiologists, nurses, etc.), the **interactions with internal and external units** (wards, diagnostic facilities, external service providers…), and the **uncertainties inherent in medical practice** (variability in procedure durations and unexpected events such as cancellations and emergency arrivals…) (Zhu et al., 2019).

Against this backdrop, the objective of this research is to advocate for a performant, resilient, robust organization of surgical suites. Performance is evaluated through key metrics such as staff workload (overtime), patient satisfaction (waiting time), and operational efficiency (operation room utilization). A robust organization can maintain performance levels despite inherent variability, while a resilient one can withstand significant disruptions without compromising performance. Achieving this objective not only enhances efficiency and profitability but also contributes to delivering safe and high-quality medical care.

To attain this objective, the study focuses on optimizing the operating schedule, known as the OR planning and scheduling problem. This involves decision-making across four hierarchical levels: strategic, tactical, offline operational, and online operational. While existing literature primarily addresses strategic, tactical, and offline operational planning, there's a notable gap in addressing disruptions at the online operational level, crucial for real-time decision-making (Guerriero & Guido, 2011).

Within this context, we investigate the following general research question: "How can we promote and maintain the performance of a surgical suite's organization under uncertainties?" To answer this question, we propose: (1) A methodology for a prospective analysis of the provisional schedule, and (2) A methodology for a retrospective analysis of the performed schedule. Leveraging insights from a digital twin-based decision support system (DT-DSS), the study seeks to improve decision-making processes for OR managers, facilitating both proactive schedule preparation and post-operative analysis. A case study based on the Hôpital Privé de La Baie (GIE Vivalto Santé, France) serves as a proof of concept, demonstrating the applicability of the proposed methodologies and their potential to enhance operational efficiency across various surgical suites.

The remainder of this article is divided into two parts: our solution proposal and an illustrative case.

2 Solution Proposal

2.1 Prospective and Retrospective Analyses Description

We propose a prospective approach to analyzing the provisional schedule, that is divided into five key steps. The objectives are to (1) anticipate and address potential disruptions before the actual execution of the schedule, and (2) to evaluate the effectiveness of predictive disruption management, and (3) improve the surgery scheduling and sequencing. First, we assess the practicality and feasibility of the provisional schedule (step #1). Then, we evaluate the provisional schedule's performance (step #2). Following this, in the robustness assessment phase, we analyze the impact of stochastic durations on schedule performance to gauge its resilience against uncertainties (step #3). Subsequently, in the resilience assessment step, we introduce additional non-elective arrivals and examine the schedule's performance under different scheduling strategies (step #4). Finally, we simulate the execution of the provisional schedule within a stochastic environment incorporating both variable durations and non-elective arrivals in order to assess the performance, the robustness and the resilience of the provisional schedule simultaneously (step #5).

On the other hand, we propose a retrospective analysis of the performed schedule. Here, the aim is to evaluate the quality of reactive disruption management and to identify whether performance shortcomings originate from offline or online decisions. To begin with, we replay the performed schedule in a deterministic environment to assess compliance with resource constraints (step #1). Then, we evaluate the performed schedule's performance (step #2). Finally, we analyze any disparities between the performed and provisional schedules to determine whether they stem from offline or online operational decisions (step #3).

2.2 Prototype Description

In this section, we present the DT-DSS developed prototype, outlining its inputs, outputs, parameters, and modeling approach. The input data comprises a relational database consisting of three tables: the master surgery schedule, the attributes of scheduled cases, and the durations of scheduled cases. The output is represented through a dashboard, showcasing key performance indicators such as staff overtime, patient waiting time, and operating room utilization. Additionally, the dashboard includes a Gantt Chart Diagram illustrating the status of each OR, including idle time, resource allocation, setup, and procedure duration.

The parameters of the prototype encompass various aspects, including the initial schedule type, the process type, the constraints on resources, the duration type, whether to keep or not the non-elective cases of the initial schedule, whether to add or not non-elective arrivals to the Initial schedule, and the number of replications.

The Digital Twin-based Decision Support System (DT-DSS) is constructed utilizing a modeling and simulation tool, Flexsim Healthcare®, which facilitates the modeling of material resources, human resources, and processes within the surgical suite. This tool enables the simulation of schedule execution, providing insights into the system's performance under different scenarios.

3 Application to an Operating Room Inspired by A Real-Case

In this section, we present a synthesis of both Step #1 and Step #2 of the retrospective analysis, focusing on a one-day use case inspired by our hospital partner.

In Figure 1, we display the OR state Gantt Chart Diagram of the performed schedule for different scheduling strategies (real-life, first fit, worst fit, and best fit). Each row corresponds to an OR and each color corresponds to an OR state: setup with anesthesiologist (yellow), setup without anesthesiologist (blue), surgical procedure (green), reversal (purple), idle time (light gray), and off schedule (dark grey). We indicate non-elective cases with a red or an orange line. For better readability, we display all the ORs for the real-life scenario, but only the ORs with non-elective cases for the other scenarios.

During the real performed schedule execution (step #1), the OR manager dealt with two urgent cases and scheduled them in OR#1 and #2. Both cases were entered at the earliest time possible and were the last case in their respective OR. In step #2, we simulate the performed schedule execution as it happened in real-life, but we change the urgent case scheduling solution. BF and WF strategies give the same scheduling solution because there are no shift end durations available for any of the urgent cases. We note that the FF strategy leads to scheduling the case in an OR#2 that has been closed for several hours, and that the BF strategy leads to scheduling both cases one after the other.

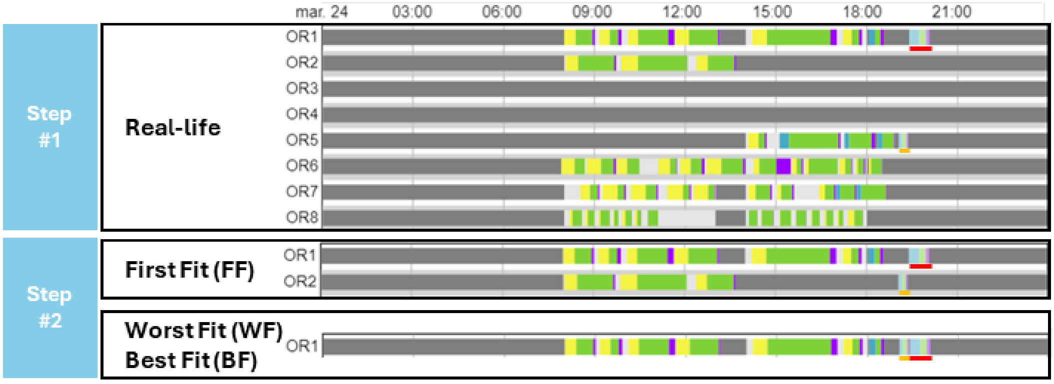

Figure 1 – Real-life scheduling of the non-elective (step #1) VS other scheduling solutions (step #2)

In Table 1, we display the value of two of our KPI for each scenario. The OR utilization and the staff overtime are the same for all scenarios because the urgent cases had arrival times and preoperative durations such that they could only be scheduled in overtime, after the last OR shift end. We note that both KPI do not reach their performance target. The prospective analysis results also offer insight into whether the prospective analysis was relevant or not.

Table 1 - KPI value for each scenario

KPI	Real-life	FF	BF / WF	Rate Target
OR Utilization (percentage of OR shift)	77.3%	77.3%	77.3%	~85%
Staff Overtime (percentage of OR shift)	9.6%	9.6%	9.6%	≤ 5%

The information gathered during these two steps can help support organizational decisions such as opening or closing OR (punctually or regularly) or serve to either encourage or discourage OR managers to implement specific non-elective scheduling strategies.

References

Guerriero, F., & Guido, R. (2011). Operational research in the management of the operating theatre: A survey. *Health Care Management Science*, *14*(1), 89–114. https://doi.org/10.1007/s10729-010-9143-6

Kaye, D. R., Luckenbaugh, A. N., Oerline, M., Hollenbeck, B. K., Herrel, L. A., Dimick, J. B., & Hollingsworth, J. M. (2020). Understanding the Costs Associated With Surgical Care Delivery in the Medicare Population. *Annals of Surgery*, *271*(1), 23–28. https://doi.org/10.1097/sla.0000000000003165

Macario, A., Vitez, T. S., Dunn, B., McDonald, T., & Brown, B. (1997). Hospital Costs and Severity of Illness in Three Types of Elective Surgery. *Anesthesiology: The Journal of the American Society of Anesthesiologists*, *86*(1), 92–100.

Zhu, S., Fan, W., Yang, S., Pei, J., & Pardalos, P. M. (2019). Operating room planning and surgical case scheduling: A review of literature. *Journal of Combinatorial Optimization*, *37*(3), 757–805. https://doi.org/10.1007/s10878-018-0322-6

A Preliminary Investigation on Bernoulli Models for Battery Energy Storage Systems with State-Dependent Transition Probabilities

Mosè Gallo, Andrea Grassi, Silvestro Vespoli
Dipartimento di Ingegneria Chimica dei Materiali e della Produzione Industriale (DICMAPI), Università degli Studi di Napoli Federico II, Napoli, NA, Italy mose.gallo@unina.it, andrea.grassi@unina.it, silvestro.vespoli@unina.it

Elisa Gebennini
Facoltà di Scienze Tecnologiche e dell'Innovazione, Universitas Mercatorum, Roma, RM, Italy elisa.gebennini@unimercatorum.it

This paper proposes the application of a stochastic model for single-buffer systems with restart level and state-dependent transition probabilities to a Battery Energy Storage System (BESS) in a smart grid. The model represents the charging and discharging processes of a BESS, taking into account the uncertainty in renewable generation and load demand. BESS is treated as a buffer between the upstream renewable source and the downstream load, with state-dependent charging and discharging efficiencies and a specific restart level which controls the input energy flow and, consequently, the system performance. The model is solved analytically using an isolation technique, allowing the derivation of a closed-form solution for the probability distribution of the system states.

Key words: Stochastic modeling; State-dependent transition probabilities; Markov processes; Battery

1. Introduction

The increasing penetration of renewable energy sources in modern power systems has led to new challenges in terms of grid stability and energy management (Zhao et al. 2012). The intermittent and unpredictable nature of solar and wind generation requires the adoption of flexible resources, such as energy storage systems, to balance supply and demand in real time (Fernandez-Blanco et al. 2017). In this context, the management of Battery Energy Storage Systems (BESS) has become a crucial task, complicated by the inherent uncertainty in renewable generation and load demand, making precise scheduling of charging and discharging cycles difficult (Ghiassi-Farrokhfal et al. 2016). Several approaches have been proposed in the literature to address this problem, ranging from deterministic optimization methods (Wu et al. 2014) to stochastic dynamic programming (Zhang et al. 2013). However, most of these techniques rely on simplified battery models and do not fully capture the complex dynamics of the storage process, such as the dependence of charging and discharging efficiencies on the state of charge (SOC) (Rao et al. 2005).

In this paper, we propose a stochastic model for the BESS that accounts for uncertainty in renewable generation and load demand. The proposed model represents the charging and discharging processes of the BESS as a buffer between the upstream renewable source and the downstream load, with state-dependent charging and discharging efficiencies. A key feature of the model is the inclusion of a restart level, which serves to control the input energy flow. By setting a minimum state of charge threshold for the battery to start charging again, the model aims to reduce the probability of charging the BESS with low efficiency values, thus increasing the overall system efficiency and battery shelf life. In fact, a BESS system based on lithium technology is known to work well when charged between 20% and 80% of the charge level Zhang and Zhang (2021). This detailed description of the system is based on a discrete-time Markov process with state-dependent transition probabilities which is solved analytically using an isolation technique.

2. Modeling approach

This section aims to model a stochastic system consisting of a BESS, i.e. the buffer, which receives discrete unit amounts of energy from an upstream charging process and releases discrete unit amounts of energy to a downstream discharging process. Specifically,

- the energy flow is modeled as a discrete flow of unit amounts of energy (e.g., 1 unit amount of energy can be equal to 1 Ah);
- the input probability that one unit amount of energy enters the BESS (from the upstream process) depends on the charge level, i.e., the input probability is state-dependent;
- the output probability that one unit amount of energy leaves the BESS (to the downstream process) depends on the charge level, i.e., the output probability is state-dependent;
- if the BESS is empty no output flow is allowed (the discharging process is starved);
- if the BESS is fully charged no input flow is allowed (the charging process is blocked);
- once the BESS gets fully charged, only the discharging flow is allowed until a certain threshold level (called "restart level" in the sequel) is reached: the charge level can only decrease or remain constant in this phase;
- once the restart level is reached, the charging process is allowed to restart: the charge level can increase/decrease/remain constant according to the stochastic dynamics of the system.

Since charging flow is not always allowed, it is possible to identify two different system behaviors that we call *standard operation* and *energy drainage*. In the *standard operation* behavior both the charging and the discharging processes are allowed: the energy storage receives one unit amount of energy in a time unit with probability $p_1^{(i)}$, being i the current charge level; it releases one unit amount of energy in a time unit with probability $p_2^{(i)}$, being i the current charge level. In the *energy drainage* behavior only the discharging process is allowed while the charging process is interrupted: the charge level i decreases with probability $p_2^{(i)}$ or remains constant with probability (1-$p_2^{(i)}$). The switch from the *standard operation* behavior to the *energy drainage* behavior occurs when the BESS is fully charged (the charge level i is equal to the energy storage capacity N); the switch back from the *energy drainage* behavior to the *standard operation* behavior occurs when the charge level reaches the restart level, denoted as L in the sequel. Hence, for charge levels below the restart level $(i < L)$ both flows are allowed (the only feasible behavior is the *standard operation* behavior); for charge levels above the restart level $(i \geq L)$, we have to distinguish which of the two behaviors the system is following.

We assume here that the system states coincide with the charge levels, by distinguishing between the two behaviors when we are above the restart level L. Hence, it is convenient to partition the state space into three partitions:

- *Standard operation A* partition: it consists of L states $(0, 1, \ldots, L-1)$ and represents the system below the restart level, where the only feasible behavior is the *standard operation* behavior;
- *Standard operation B* partition: it consists of $N-L+1$ states $(L, L+1, \ldots, N-1, N)$ and represents the system at the restart level or above the threshold level, given that the system is following the *standard operation* behavior (both flows are allowed);

- *Energy drainage* partition: it consists of $N - L - 1$ states $(L + 1, \dots, N - 1)$ and represents the system above the restart level given that the system is following the *energy drainage* behavior (the charge level is decreasing after having reached the full state).

Given the three aforementioned state partitions, the so-called *isolation technique* can be applied in order to simplify the mathematical treatment of the problem by treating each partition independently, i.e. in "isolation". This can be done by exploiting the proposition that, in steady state, the probability of switching from partition P_j to partition P_k equals the probability of switching from partition P_k to P_j. In other words, the probability of entering a partition must be equal to the probability of leaving that partition. Once each partition is solved in "isolation", the original system can be rebuilt by introducing the concept of "partition probability", i.e., the probability for the system being in one and only one of the three partitions at a given time instant. This approach leads to the exact closed-form solution for the stationary probability distribution of the system states. Thus, by applying the isolation technique, a set of analytical formulas can be derived to express the probability that the BESS is at a given charge level.

3. Conclusion

This paper presented a stochastic model for BESS that captures the uncertainty in renewable generation and load demand. By solving the model analytically, we obtained a closed-form solution for the stationary probability distribution of the state of the system. This detailed description of the behavior of the energy storage system could serve as a foundation for optimization or as a benchmark for evaluating heuristic management strategies. Future research directions include extending the model to multiple energy storage devices, incorporating battery degradation mechanisms, and integrating the model with real-time energy management optimization frameworks.

Acknowledgments

This study was carried out within the MICS (Made in Italy – Circular and Sustainable) Extended Partnership and received funding from the European Union Next-GeneraFonEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1551.11-10-2022, PE00000004). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

References

Fernandez-Blanco, Ricardo, Yury Dvorkin, Bolun Xu, Yishen Wang, Daniel S Kirschen. 2017. Optimal energy storage siting and sizing: A wecc case study. *IEEE Transactions on Sustainable Energy* **8**(2) 733–743.

Ghiassi-Farrokhfal, Yashar, Catherine Rosenberg, Srinivasan Keshav, Marie-Benedicte Adjaho. 2016. Joint optimal design and operation of hybrid energy storage systems. *IEEE Journal on Selected Areas in Communications* **34**(3) 639–650.

Rao, V., G. Singhal, A. Kumar, N. Navet. 2005. Battery model for embedded systems. *18th International Conference on VLSI Design held jointly with 4th International Conference on Embedded Systems Design*. 105–110. doi: 10.1109/ICVD.2005.61.

Wu, Yuan, Vincent K. N. Lau, Danny H. K. Tsang, Li Ping Qian, Limin Meng. 2014. Optimal energy scheduling for residential smart grid with centralized renewable energy source. *IEEE Systems Journal* **8**(2) 562–576. doi: 10.1109/JSYST.2013.2261001.

Zhang, Shuzhi, Xiongwen Zhang. 2021. A multi time-scale framework for state-of-charge and capacity estimation of lithium-ion battery under optimal operating temperature range. *Journal of Energy Storage* **35** 102325. doi:https://doi.org/10.1016/j.est.2021.102325.

Zhang, Yan, Nikolaos Gatsis, Georgios B Giannakis. 2013. Robust energy management for microgrids with high-penetration renewables. *IEEE Transactions on Sustainable Energy* **4**(4) 944–953.

Zhao, Junhua, Fushuan Wen, Zhao Yang Dong, Yusheng Xue, Kit Po Wong. 2012. Optimal dispatch of electric vehicles and wind power using enhanced particle swarm optimization. *IEEE Transactions on Industrial Informatics* **8**(4) 889–899.

A sample path-based method for the digital twin prediction update synchronization problem of unreliable production lines

Yi-Xing Gou
Department of Industrial Engineering, School of Mechanical Engineering, Northwestern Polytechnical University, Xi'an, Shaanxi 710072, P. R. China, pacpos.yxgou@gmail.com

Andrea Matta
Department of Mechanical Engineering, Politecnico di Milano, Milano 20156, Italy, andrea.matta@polimi.it

Barış Tan
Faculty of Engineering & Faculty of Business, Özyeğin University, Çekmeköy Campus Nişantepe District, Orman Street, Çekmeköy, Istanbul 34794, Türkiye, baris.tan@ozyegin.edu.tr

The digital twin prediction update synchronization problem determines whether or not to update the performance prediction from the digital twin at each observation period depending on the observed state of the physical system. Existing approaches provide solutions for the prediction update problem, but they can be applied only to simple systems. In this study, we propose a sample path-based method to solve the prediction update problem for unreliable production lines composed of multiple machines and finite buffers. The method estimates the performance measure for each synchronization decision by partially observing the state of the system. An optimal state-dependent synchronization policy is determined based on the observed state to balance the prediction bias and the synchronization cost. The results show that partially observing the state of the bottleneck machine, rather than fully observing the state of the system, is efficient for solving the problem without requiring a very long sample path.

Key words: Digital twin synchronization; Prediction update; Sample path; Unreliable production lines

1. Introduction

Digital twins are considered important parts of smart manufacturing to monitor and control physical systems more accurately and optimize manufacturing processes (Tao et al. 2019). The development of IoT technologies makes data collection and communication more efficient, which enables digital twins to capture the behavior of physical systems in increasing detail. A digital twin of a production system involves a lot of variables that reflect the dynamics of the physical system, which increases the difficulty of aligning the digital twin with its physical system. Synchronizing all variables of the digital twin with the data from the physical system may take much time and resources (Modoni et al. 2019, Sargent 2013). Not synchronizing the digital twin with the actual state of the physical system results in significant prediction errors that may cause production loss and management challenges (Zipper and Diedrich 2019, Zipper 2021). Consequently, it is necessary to determine when and how to synchronize the digital twin with its physical system in the most efficient way.

The optimal digital twin synchronization problem for production systems is defined as a stochastic control problem, and the prediction update synchronization problem of an unreliable machine is solved analytically in Tan and Matta (2024). The prediction update problem aims to decide whether or not to update the prediction from the digital twin depending on the observed state of the physical system. A sample path-based method is introduced in Tan and Matta (2024) that can be used to solve the problem. The method requires estimating the system performance for each synchronization decision by observing the state of the system from the sample path. A state-dependent synchronization policy is determined based on the observed state of the system. In the case of a single unreliable machine, the state of the system is fully observable and is either up or down. For an unreliable production line, the state of the system is composed of states of all machines and buffers. If we fully observe the state of the system, the available data in a given sample path that can be used to estimate the performance prediction is limited. This will result in an inaccurate prediction when the number of system states is large. In this study, partial observations of the production lines are used to solve the problem. The main contribution of this work is a sample path-based method that estimates the performance prediction and determines the optimal state-dependent policy based on the partially observed state of the system.

2. Prediction update problem of unreliable production lines

A production line consisting of M unreliable machines and $M-1$ finite buffers is considered in this study. Machine $m_i, i = 1, 2, \ldots, M$ follows the geometric reliability model. Buffer $b_i, i = 1, 2, \ldots, M-1$ has a finite capacity denoted by B_i. The state of the system at time t is a set of the states of all machines and buffers, denoted by $y_t = (\alpha_t^1, \ldots, \alpha_t^M, \beta_t^1, \ldots, \beta_t^{M-1})$. α_t^i denotes the state of machine $m_i, i = 1, 2, \ldots, M$ at time t, which is 1 if machine m_i is up and 0 if it is down. β_t^i denotes the buffer level of buffer b_i at time t, $\beta_t^i \in \{0, 1, \ldots, B_i\}$. If not all machines and buffers are observable or not all observations are used to determine the synchronization policy, y_t includes only the states of the machines and the buffers observed or used.

A digital twin, which uses a discrete event simulation model of the physical system, is adopted to predict the throughput of the production line based on the available history. The throughput is evaluated each Δ cycles at times $\Delta, 2\Delta, \ldots, N\Delta$ until the end of the planning observation period. The throughput within a given time interval $[t, t+\tau)$ is the expected number of parts produced by the production line during this interval denoted by $\mathbb{E}[TH_{t,t+\tau}]$.

Synchronizing the digital twin with the current observations allows obtaining a better prediction of the throughput. This comes at a cost that includes retrieving the data from the production line, executing a simulation experiment, and changing the production resources based on the updated prediction. The synchronization decision at time $n\Delta$ is $\mathbf{u}_n = \mathrm{H}_n$, in which H_n is equal to 1 if the digital twin is synchronized at time $n\Delta$ and 0 otherwise. The synchronization cost of decision $\mathbf{u}_n = \mathrm{H}_n$ at time $n\Delta$ is calculated with $C_{\mathrm{DT}}(\mathbf{u}_n) = c_{\mathrm{H}}\mathrm{H}_n$, where c_{H} is the cost for each synchronization. We consider a state-dependent policy that the synchronization decision is taken depending on both the evaluation period n and the observed state of the system y_n. Not synchronizing the digital twin results in inaccurate prediction, which incurs a cost of prediction bias. The bias cost at time $n\Delta$ denoted with $C_{\mathrm{B}}(n, \mathbf{S}_n, \mathbf{u}_n)$ depends on the difference between the best estimation $R_n^*(y_n)$ with the most recent observations y_n and the estimation from the digital twin based on synchronization decision at time $n\Delta$, $R_n^{\mathrm{DT}}(\mathbf{S}_n, \mathrm{H}_n)$, i.e., $C_{\mathrm{B}}(n, \mathbf{S}_n, \mathbf{u}_n) = c_{\mathrm{B}}(R_n^{\mathrm{DT}}(\mathbf{S}_n, \mathrm{H}_n) - R_n^*(y_n))^2$. The prediction update synchronization problem is to find a policy that determines when to update the throughput prediction from the digital twin to minimize the expected total cost of prediction bias and of synchronizations over N observation periods:

$$\min_{\{\mathrm{H}_n\}} \quad \sum_{n=1}^{N} (c_{\mathrm{H}}\mathrm{H}_n + c_{\mathrm{B}}\mathbb{E}[(R_n^{\mathrm{DT}}(\mathbf{S}_n, \mathrm{H}_n) - R_n^*(y_n))^2]). \tag{1}$$

A sample path-based method is proposed to determine the optimal state-dependent policy for the model described in Equation (1). The method estimates the throughput based on the observed state of the last synchronization at time $k\Delta$, $r_{n,k}^{\mathrm{DT}}(y_k) = \mathbb{E}[TH_{n\Delta,(n+1)\Delta}|y_k]$. A sample path with L periods consists of the states of all machines and buffers during $L\Delta$ cycle times. The state of the system at time t in the sample path is denoted with o_t. If the state of the system at time $l\Delta$ in the sample path is equal to the state of the digital twin at the last synchronization at time $k\Delta$, i.e., $o_{l\Delta} = y_k$, the throughput during the next period after $n-k$ periods at time $l\Delta$ in the sample path can be used to estimate the throughput at time $n\Delta$ based on state y_k. Then $r^{\mathrm{DT}}_{n,k}(y_k)$ can be obtained by using the average of the observations from the sample path. Accordingly, the estimate from the digital twin based on the synchronization decision H_n at time $n\Delta$, $R^{\mathrm{DT}}_n(\mathbf{S}_n, \mathrm{H}_n)$ is evaluated as

$$R^{\mathrm{DT}}_n(\mathbf{S}_n, \mathrm{H}_n) = \begin{cases} r^{\mathrm{DT}}_{n,n}(y_n), & \mathrm{H}_n = 1 \\ r^{\mathrm{DT}}_{n,k}(y_k), & \mathrm{H}_n = 0 \end{cases}. \tag{2}$$

The best estimate of the throughput at time $n\Delta$ is obtained by taking synchronization action with $\mathrm{H}_n = 1$, i.e., $R^*_n(y_n) = r^{\mathrm{DT}}_{n,n}(y_n)$. Then, an optimal state-dependent policy is determined to solve the problem formulated in Equation (1).

3. Numerical results and conclusions

To validate the performance of the sample path-based method, a simulation-based method is also proposed in this work. The method differs from the sample path-based method in that $R^*_n(y_n)$ is evaluated using simulation based on the current full observations of the system at time $n\Delta$. The performance of both methods is compared in two test cases. Some results are shown in Figure 1.

The numerical results show that observing the bottleneck machine is more critical to determining an optimal state-dependent policy than observing other machines in unreliable production lines. As the length of the sample path increases, the average cost and the average number of synchronizations obtained using the sample path-based method almost stay stable. Longer sample paths are not required for the cases analyzed. The simulation-based method yields policies with lower average costs than the sample path-based method, but it takes much more time as the sample path becomes longer. For more complex systems, the sample path-based method is more efficient to apply in solving the prediction update synchronization problem. This will be studied in the future.

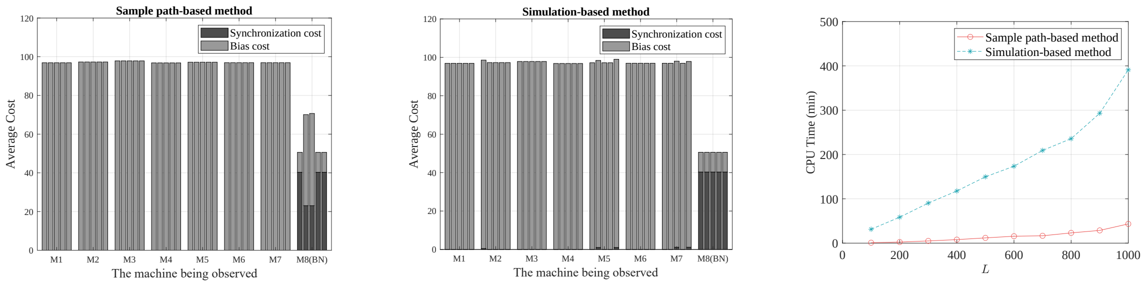

Figure 1. Effect of the observed machine on the average cost obtained using both methods over five experiments ($c_{\mathrm{H}} = 20, L = 1000$), and CPU time conducting one experiment using both methods for different values of L

References

Modoni, G. E., E. G. Caldarola, M. Sacco, W. Terkaj. 2019. Synchronizing physical and digital factory: benefits and technical challenges. *Procedia CIRP* **79** 472–477.

Sargent, R. G. 2013. Verification and validation of simulation models. *Journal of Simulation* **7**(1) 12–24.

Tan, B., A. Matta. 2024. The digital twin synchronization problem: Framework, formulations, and analysis. *IISE Transactions* **56**(6) 652–665.

Tao, F., Q. Qi, L. Wang, A.Y.C. Nee. 2019. Digital twins and cyber-physical systems toward smart manufacturing and industry 4.0: Correlation and comparison. *Engineering* **5**(4) 653–661.

Zipper, H. 2021. Real-time-capable synchronization of digital twins. *IFAC-PapersOnLine* **54**(4) 147–152.

Zipper, H., C. Diedrich. 2019. Synchronization of industrial plant and digital twin. *2019 24th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*. IEEE, 1678–1681.

A state-dependent decomposition method for discrete-time open tandem queues

Christoph Jacobi
Institute for Material Handling and Logistics, Karlsruhe Institute of Technology, Karlsruhe, Germany, jacobi@kit.edu

Decomposition often is the only feasible and computationally efficient approach to compute steady-state performance measures for queueing networks. However, performance results may be subject to severe approximation errors since decomposition methods usually assume that the connecting stream can be approximated by renewal processes. To overcome the renewal assumption, we present the semi-Markov arrivals decomposition approach (SMAD). SMAD is a refined decomposition approach, where the connecting stream between the upstream and the downstream station is described by a semi-Markov process. Using this modelling approach, state-dependent inter-departure times from the upstream queue are preserved for downstream queuing analysis. Numerical results demonstrate that the approach produces accurate results, compared to simulation.

Key words: Decomposition; Point process; semi-Markov process

---

1. Introduction

We consider a discrete-time open tandem queue, where the upstream queue is of type $M/G/1$, and the downstream queue is of type $G/G/1$. Decomposition is often the only feasible and computationally efficient approach to compute steady-state performance measures in this type of queuing network. This approach partitions the network into individual queuing systems and analyses them in isolation. It is based on the assumption that the output stream of the upstream $M/G/1$-queue – which is fed into the downstream $GI/G/1$-queue – can be approximated by a renewal process. However, it is well known that the departure process is a point process that is generally difficult to deploy for queueing system analysis (Whitt 1981, 1982). Recently, it has been shown that performance results may be subject to severe approximation errors when applying the renewal decomposition method in this tandem queue (Jacobi and Furmans 2022). Thus, we outline a novel approach to overcome the renewal assumption. We use a semi-Markov arrival process to model the connecting stream between the upstream and the downstream queue.

2. Literature review

Decomposition approaches for open queueing networks generally rely on two basic assumptions (Govil and Fu 1999): First, it is assumed that the individual queueing systems in the network can be treated as being statistically independent. Second, it is assumed that the input to each queueing system is a renewal process. In the continuous-time domain, this approach was first by applied by Kuehn (1979) with modifications presented by Shanthikumar and Buzacott (1981), Whitt (1983), and Reiman (1990). In the discrete-time domain, Haßlinger and Rieger (1996) and Furmans (2004) proposed refinements of these so-called parametric decomposition methods which allows for the computation of the entire probability distributions of performance measures.

A crucial problem for decomposition methods in the continuous-time domain is the computation of the variability measures for the internal flows, and for the departure stream. The Queueing Network Analyzer (QNA) (Whitt 1983) employs two procedures (Whitt 1982) to approximate the point departure processes by renewal processes, the stationary interval method and the asymptotic method. Since neither of both methods yields promising results for a wide range of variability parameters, Whitt (1983) introduces a hybrid procedure based on the work by Albin (1984a,b). Finally, the variability of departure stream of the $GI/G/1$-queue is observed as an approximation of the stationary interval method (Whitt 1984).

Haßlinger and Rieger (1996) present a refinement of the parametric decomposition approach for the analysis of open queueing networks in the discrete-time domain. The discrete distribution of superpositions of renewal processes is reversibly obtained by the distribution of the minimum of the residual times of all superposed flows. A recursive method and a faster approach based on the z-transform for the computation of the stochastic split of a renewal process are presented. Despite discussing the renewal assumption and its implications, Haßlinger and Rieger (1996) state that "further study is needed to construct [...] representations of non-renewal processes" in the discrete-time domain that enable the computation of exact results.

3. Modelling approach

To overcome the renewal assumption for the analysis of tandem queues with Poisson arrivals and general service times, we introduce the semi-Markov decomposition approach (SMAD). The novelty of this decomposition method is that a semi-Markov process (SMP) is used to model the connecting stream between the upstream $M/G/1$- and the downstream $G/G/1$-queue. Let the stochastic process $\mathcal{Z} = \{(N_k, D_k), k = 1, 2, ...\}$ denote a semi-Markov process where $N \in \mathbb{N}_0$ is the number of customers in the upstream $M/G/1$-queue immediately after the departure instance of customer k, and $D_k \in \mathbb{N}$ is the inter-departure time between customers k and $k+1$. Let the probability function

$$f(t \,|\, i) = P(D = t \,|\, N = i) \tag{1}$$

denote the conditional probability that the inter-departure time is equal to t, given that the embedded Markov chain of the semi-Markov process Z_k is in state $N_k = i$. The probability function $f(t \,|\, i)$ is equal to the service time, if the system is not empty immediately after the departure instance (that is, $i > 0$), and equal to the sum of the remaining inter-arrival time and the service time, if the system is starving after departure instance k (that is, $i = 0$). For downstream queueing analysis, we deploy the discrete-time $SM/G/1$-queue, which has been introduced by Rieger and Haßlinger (1994).

4. Numerical results

We consider a tandem queue where the service time distributions are equal at the upstream and the downstream station, $P(B = 15) = P(B = 16) = 0.5$, and the arrival stream is defined by $\lambda = 0.0613$. The utilisation of the tandem queue is $\rho = 0.950$. We compute the probability distribution of waiting time at the downstream queue, and compare the results to the waiting time distributions obtained with the renewal decomposition approach and simulation. While the renewal decomposition method computes an expected waiting time $E(W) = 1.74$, SMAD computes an expected waiting time $E(W) = 2.11$, and simulation yields an expected waiting time $E(W) = 2.10$. We performed a Chi-Square Goodness-of-Fit Test and found a significant relationship between the waiting time distributions computed with SMAD, and simulation $\left(\chi^2(11; 612{,}682) = 5.69, p = .893\right)$.

5. Conclusion

Decomposition approaches for open queueing networks approximate the interconnecting streams as renewal processes. While this assumption allows for computationally efficient models, performance results obtained at downstream queues might be prone to considerable approximation errors. The novelty of SMAD is that a SMP is used to model the connecting stream between the upstream $M/G/1$- and the downstream $G/G/1$-queue. Thus, SMAD captures the state-dependent inter-departure times in the upstream $M/G/1$-queue departure process. While SMAD computes performance results with great accuracy, state space explosion of the embedded Markov chain in the SMP remains a concern. Thus, introducing a state space limit to increase the computationally efficiency is a natural extension to the method.

Acknowledgments
The author wishes to thank J. George Shanthikumar for his guidance and support in this work. The author received funding from the Karlsruhe House of Young Scientists (KHYS), which is gratefully acknowledged.

References

Albin, S.L. 1984a. Approximating a point process by a renewal process, II: Superposition arrival processes to queues. *Operations Research* **32**(5). doi:10.1287/opre.32.5.1133.

Albin, S.L. 1984b. Approximating queues with superposition arrival processes. Ph.D. thesis, Department of Industrial Engineering and Operations Research, Columbia University, New York, NY, United States.

Furmans, K. 2004. A framework of stochastic finite elements for models of material handling systems. *8th International Material Handling Research Colloquium*. Graz, Austria.

Govil, M.K., M.C. Fu. 1999. Queueing theory in manufacturing: A survey. *Journal of Manufacturing Systems* **18**(3). doi:10.1016/S0278-6125(99)80033-8.

Haßlinger, G., E.S. Rieger. 1996. Analysis of open discrete time queueing networks: A refined decomposition approach. *Journal of the Operational Research Society* **47**. doi:10.1057/jors.1996.76.

Jacobi, C., K. Furmans. 2022. Point and interval estimation of decomposition error in discrete-time open tandem queues. *Operations Research Letters* **50**(5). doi:10.1016/j.orl.2022.07.009.

Kuehn, P. 1979. Approximate analysis of general queuing networks by decomposition. *IEEE Transactions on Communications* **27**(1). doi:10.1109/TCOM.1979.1094270.

Reiman, M. 1990. Asymptotically exact decomposition approximations for open queueing networks. *Operations Research Letters* **9**(6). doi:10.1016/0167-6377(90)90055-A.

Rieger, E.S., G. Haßlinger. 1994. An analytical solution for the discrete time single server system with semi-markovian arrivals. *Queueing Systems* **18**. doi:10.1007/BF01158775.

Shanthikumar, J.G., J.A. Buzacott. 1981. Open queueing network models of dynamic job shops. *International Journal of Production Research* **19**(3). doi:10.1080/00207548108956652.

Whitt, W. 1981. Approximating a point process by a renewal process: The view through a queue, an indirect approach. *Management Science* **27**(6). doi:10.1287/mnsc.27.6.619.

Whitt, W. 1982. Approximating a point process by a renewal process, I: Two basic methods. *Operations Research* **30**(1). doi:10.1287/opre.30.1.125.

Whitt, W. 1983. The queueing network analyzer. *The Bell System Technical Journal* **62**(9). doi:10.1002/j.1538-7305.1983.tb03204.x.

Whitt, W. 1984. Approximations for departure processes and queues in series. *Naval Research Logistics Quarterly* **31**(4). doi:10.1002/nav.3800310402.

An approximate analytical method for the performance evaluation of semiconductor front-end fabrication integrating photolithography inspection strategies

Maria Chiara MAGNANINI
Department of Mechanical Engineering, Politecnico di Milano, Milano, Italy, email: mariachiara.magnanini@polimi.it

Riccardo POMI
Department of Mechanical Engineering, Politecnico di Milano, Milan, Italy, riccardo.pomi@polimi.it

Dragan DJURDJANOVIC
Cockrell School of Engineering, University of Texas at Austin, Austin TX, USA, dragan.djurdjanovic@utaustin.edu

Tullio TOLIO
Department of Mechanical Engineering, Politecnico di Milano, Milano, Italy, email: tullio.tolio@polimi.it

In semiconductor manufacturing, lithography represents the core process of frontend fabrication as the quality outcome in terms of overlay errors depends entirely on it. Hence, particular attention is devoted to the inspection of each wafer layer, having 100% measurements of markers distributed across a wafer with subsequent long inspection times. At the same time, process control is based on each layer's overall measurements, discouraging companies from improving productivity by reducing inspection time. As a consequence, in this context product, process and system are extremely inter-related. Recent developments in joint product-process modelling show that robust model-based control coupled with optimal down-selection of measurement markers enables improved process control without increasing the defects. However, when considering the system level, new dynamics should be accounted for in order to take decisions about production system configuration and operations. This paper proposes a novel analytical model for the evaluation of quality and productivity performance in manufacturing systems characterized by propagation of quality errors, process adaptation and alternative inspection policies. The proposed model is general but particularly useful for the semiconductor sector. Application of this method to an industrial-scale semiconductor manufacturing system shows that when product-process-system are considered together, global optimal solutions can be achieved.

Keywords: Manufacturing systems; Decomposition; Quality control; Semiconductors.

1 Motivation and objectives

Performance evaluation is an essential step towards improvement of manufacturing systems and involves estimating the effect of various factors on performance indicators such as throughput, work in process level, efficiency, among others. Quality is considered one of the most critical factors that can influence system behavior and performance. Therefore, efforts for the integration in performance evaluation models start with including production resources that may produce parts with imperfect quality. Contributions from the community include the integration of SPC methods in analytical models, the evaluation and improvement of systems including quality-quantity coupled operations, quality deterioration, caused by machine stoppages and long waiting times, scrapping policies in response to machine failures, considering buffer-less systems with both presence and absence of deterioration memory, or scrapping policies resulting from stoppages of production systems, and differentiating between long lasting failures, causing deterioration, and rapid failures. The semiconductor manufacturing system is recognized as a highly intricate production process. The initial phase, known as wafer manufacturing or the front-end, incurs significant costs. During this phase, circuits are methodically layered onto the wafer using a series of sequential procedures. Numerous processing steps are involved in this phase. Consequently, the dynamics, performance, and characteristics of both the process and the end product are determined by an extensive range of factors. The propagation of multi-stage dynamics has a clear impact on the responsiveness of the quality strategy. At product level, each stage operates a transformation on the product and may add product deviations, in form of overlays. Particular attention is devoted to the inspection of each wafer layer, having 100% measurements of markers distributed across a wafer with subsequent long inspection times. Despite technological advancements and research & development efforts in this direction, inspection time could not be reduced yet, and still represents a limit for the system productivity (Chien 2020). To ensure the required quality, process control based on models such as stream of variation (SoV) was introduced and studied (Graff 2023). Errors propagation is analytically described and incorporated in process control models robust to inaccuracies between process parameters and error generation. As a consequence, here product, process and system are extremely inter-related, and a clear understanding of the interaction between quality and production is envied to improve the capacity planning of fabs (Ghasemi 2020). This work integrates state-of-the-art quality models for inspection optimization in lithography combined with model-based process control into an approximate analytical model. Given the unique characteristics, this novel approach is particularly suitable for semiconductor fabrication, but it has general validity. The novel contributions of this work can be summarized as follows:

1. State-based Markovian synthetic representation of semiconductor fabrication stages in lithography threads embedding selective inspection strategies and model-based process control;
2. System-level stochastic model with close-to-reality assumptions, as split and scrap, integrating propagation of product quality and inspection errors along stages;
3. Joint optimization of productivity and quality at system-level in semiconductor fabrication for the downsizing of measurement selection.

2 Methodology and results

The system includes K stages, where each stage S_k is composed of one photolithography machine and one inspection station. Stages are decoupled by buffers B_k characterized by buffer capacity N_k. Both photolithography machines and inspection stations are fully reliable, and no failures occur in either stage. If the downstream carrier has enough capacity, no blocking phenomenon affects the lithography machines. Similarly, starvation may occur if the lithography thread is not adequately balanced. Whenever an inspection station performs a conformity check on the wafer and finds that the patterned layer is defective, a parameter tuning of the photolithography machine is immediately performed, as usual in run-to-run control. At the same time, the defective wafer is promptly unloaded from the inspection machine and rejected from the line. This preserves downstream station capacity from being wasted on processing defective wafers. Given the high automation level and the type of process, both photolithography and inspection have constant cycle time. The target performance measures include the average total throughput th_K at the end of the lithography thread, the average total throughput of good wafers th_K^G and the average throughput of expected defective wafer th_K^{NG}. If the inspection station inspects 100% of the markers on the patterned layer of the wafer, the throughput of expected defective wafers th_K^{NG} includes only the wafers for which the last patterned layer is out of specification. On the other hand, if the inspection station inspects a reduced number of markers on the patterned layer, the throughput of expected defective wafers th_K^{NG} includes also those wafers with non-compliant layers that have not been detected by inspection stations $k = 1, \dots, K-1$. From the perspective of the production flow, there is no conservation of flow, as the production flow decreases along the lithography thread, because at each stage in-process scrap may occur, when non-compliant patterned layers are detected. The product and process model utilizes a robust run-to-run control that considers not only overlay errors but also stack-up overlay errors. These errors are described by the summation of the overlay of non-adjacent layers, using Zernike polynomial-based models (Zhang 2022). The product-process model selects the best combination of a given percentage of available measurement points $m, m \in [0,1]$, known as markers, for the robust control of lithography errors. This ensures that the measure of overlay errors at all the candidate measurement points is minimized. However, it is important to note that the down-selection of markers for the wafer inspection may not detect bad wafers, allowing them to continue along the manufacturing line. The definition of the system-level Markovian model extends the work introduced by Magnanini (2023). Hence, the state space for the stage S_k at system level is $\Omega_k = \{G, BD, BND, S_G, S_{BD}, S_{BND}, B_G, B_{BND}, NQ_1, NQ_j, \dots, NQ_{k-1}\}$, where:

- Local states: state G represents the production condition of good wafers at the lithography machine that afterwards are correctly identified as compliant by the inspection station; state BD represents the production condition of bad wafers at the lithography machine that afterwards are correctly identified as not-compliant by the inspection station; state BND represents the production condition of bad wafers at the lithography machine that afterwards are not detected as not-compliant by the inspection station, due to the down-selection of markers m.

- Remote states $\{S_G, S_{BD}, S_{BND}\}$ in which the stage is upstream limited from the perspective of production flow, i.e. the considered stage is starved or slowed down.
- Remote states $\{B_G, B_{BND}\}$ in which the stage is downstream limited from the perspective of production flow, i.e. the considered stage is blocked or slowed down.
- Remote states $\{NQ_1, NQ_j, \ldots, NQ_{k-1}\}$ in which the stage is processing defective layers from previous stages that were not detected by the down-selection of markers in upstream inspection stations.

The decomposition method obtains the probability density function of each Building Block, then the two-level decomposition approach obtains the transition rate matrix for each Integrated Machine by lumping of the Building Block state-space into the Integrated Machine state-space. In this way, the number of states do not explode. Results in Fig. 1 shows that reducing the number of inspected markers clearly improves the productivity with respect to the baseline case (100%). On the other hand, when the yield is analyzed, it is possible to notice that using a good uncertainty model of the propagation of deviations coupled with the process control guarantees the identification of the optimal set of measurement markers, not only to the optimal number of measurement markers to be inspected. In particular, the highest productivity in terms of good wafers is obtained when the measurement reduction is set according to the proposed novel product-process-system model, as it exploits the knowledge from the product and process level, together with the knowledge about the dynamics at system level.

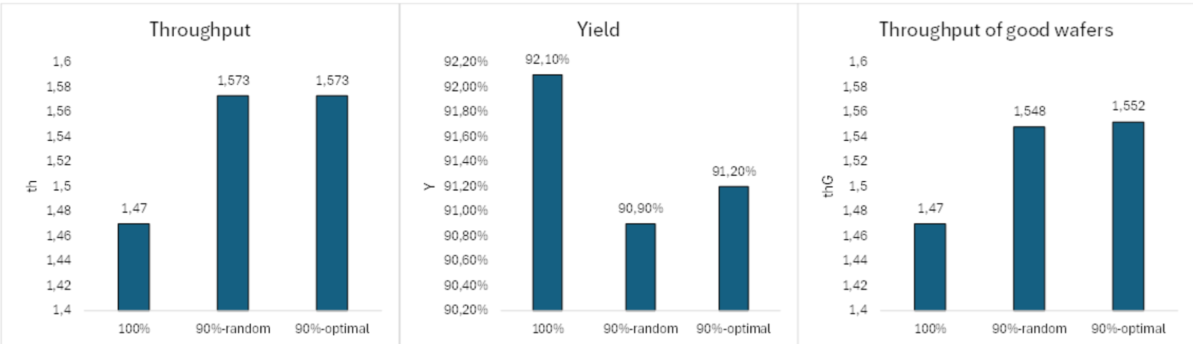

Figure 1. Comparison of performance measures with alternative quality control strategies.

References

Arinez, J., Meerkov, S. M., Zhang, L. 2010. Quality/Quantity Improvement in an Automotive Paint Shop: A Case Study. IEEE Transactions on Automation Science and Engineering 7(4), 755-761.

Chien, C. F., Chen, Y. H., Lo, M. F. 2020. Advanced quality control (AQC) of silicon wafer specifications for yield enhancement for smart manufacturing. IEEE Transactions on Sem. Manuf,, 33(4), 569-577.

Colledani, M., Tolio, T. 2011. Integrated analysis of quality and production logistics performance in manufacturing lines. International Journal of Production Research 49(2), 485-518.

Ghasemi, A., Azzouz, R., Laipple, G., Kabak, K. E., Heavey, C. 2020. Optimizing capacity allocation in semiconductor manufacturing photolithography area–Case study: Robert Bosch. Journal of Manufacturing Systems, 54, 123-137.

Graff, N., Hanasusanto, G. A., Djurdjanović, D. 2023. Robust control of maximum photolithography overlay error in a pattern layer. CIRP Annals, 72(1), 429-432.

Kim, J., Gershwin S. B. 2005. Integrated quality and quantity modeling of a production line. OR Spectrum 27, 287-314.

Liberopoulos, G., Kozanidis, G., Tsarouhas, P. 2007. Performance Evaluation of an Automatic Transfer Line with WIP Scrapping During Long Failures. Manufacturing & Service Operations Management 9(1), 62-83.

Magnanini, M. C., Tolio, T. 2023. A Markovian model of asynchronous multi-stage manufacturing lines fabricating discrete parts. Journal of Manufacturing Systems, 68, 325-337.

Tempelmeier, H., Bürger, M. 2001. Performance evaluation of unbalanced flow lines with general distributed processing times, failures and imperfect production. IIE Transactions 33(4), 293-302.

Zhang, H., Feng, T., Djurdjanovic, D. 2022. Dynamic down-selection of measurement markers for optimized robust control of overlay errors in photolithography processes. IEEE Transactions on Semic. Manuf., 35(2), 241-255.

Auction-Based Mechanism for Job Allocation in Hybrid Manufacturing Planning and Control Architectures

Silvestro Vespoli, Andrea Grassi, Guido Guizzi
Dipartimento di Ingegneria Chimica dei Materiali e della Produzione Industriale (DICMAPI), Università degli Studi di Napoli Federico II, Napoli, NA, Italy silvestro.vespoli@unina.it, andrea.grassi@unina.it, guido.guizzi@unina.it

Valentina Popolo
Facoltà di Ingegneria, Università Telematica Pegaso, Napoli, NA, Italy valentina.popolo@unipegaso.it

In modern manufacturing environments, effective allocation of production resources is crucial for addressing new market challenges. This work proposes an auction-based mechanism for horizontal bargaining within a semi-heterarchical architecture. The mechanism is built around three critical functions: activation, evaluation, and bidding. Particularly noteworthy are the two distinct evaluation approaches proposed to ascertain the value of jobs for auction. The first approach, inspired by a M/G/n queue model, aims to estimate the potential impact of a job on production line throughput. The second approach is strategically focused on strengthening the decision-making capabilities of the lower-level controller. It evaluates how the inclusion of the auctioned job could broaden the range of operational choices at the lower level, potentially optimizing system performance. Through a preliminary simulation study, we examined the effectiveness of this auction-based mechanism across a spectrum of scenarios. The results show that the proposed approach effectively aligns production processes, especially in settings where job distribution and processing time variability are critical.

Key words: Auction-based mechanism; Semi-heterarchical architecture; Manufacturing planning and control; Stochastic modeling

1. Introduction

The challenges posed by Industry 4.0 in manufacturing planning and control (MPC) require agile and responsive production systems that can efficiently allocate resources (Borangiu et al. 2020, Ivanov et al. 2018). Among various MPC architectures, semi-heterarchical architectures have emerged as a promising approach, combining the advantages of hierarchical coordination and decentralised reactivity (Antons and Arlinghaus 2022). In the architecture proposed by Grassi et al. (Grassi et al. 2020), High-Level Controllers (HLCs) coordinate production processes at a higher level of abstraction, while lower-level controllers focus on executing tasks and managing local resources. A key component of this architecture is the Job Ready Queue (JRQ), a virtual job queue positioned before the production system that contains jobs awaiting insertion by the dispatcher.

This work introduces an auction-based mechanism to facilitate horizontal bargaining between HLCs, enhancing job allocation and resource utilisation. The mechanism aims to improve decision-making and operational flexibility within semi-heterarchical structures by enabling HLCs to negotiate and exchange jobs based on their suitability and potential impact on production performance.

2. Proposed Auction-Based Mechanism

The proposed auction-based mechanism consists of three key functions: activation, evaluation, and bidding. A sealed bid second-price auction, also known as a Vickrey auction, was chosen for this mechanism because of its ability to ensure the highest degree of cooperation among participants.

2.1. Activation Function

The activation function is responsible for identifying the suitable jobs for auction. In the proposed case, the HLC monitors the JRQ and triggers the auction process when a job has remained in the queue for a pre-defined duration, indicating that it may benefit from being allocated to a different HLC.

2.2. Evaluation Function

The evaluation function plays a crucial role in determining the value of an auctioned job. In this work, we propose two distinct approaches to job evaluation, each grounded in stochastic modelling principles. The main conceptual difference between the two approaches lies in their focus: the first approach aims to estimate the impact of the auctioned job on the system's throughput, while the second approach focusses on enhancing the decision-making capabilities of the lower-level controller by considering the potential impact on the range of operational choices.

2.2.1. M/G/1 Queue Model Approach The first evaluation approach is inspired by the M/G/1 queue model, which assumes exponentially distributed interarrival times and generally distributed service times. The idea is to use the M/G/1 model to evaluate the cycle time before and after the addition of the auctioned job, in order to assess the impact on the line's throughput once the job has been accepted. The relative increase in throughput is the key metric in this approach.

- Calculates the cycle time of the queue (CT_{q_i}) for each machine i as follows:

$$CT_{q_i} = \left(\frac{1+c_i^2}{2}\right)\left(\frac{u_i}{1-u_i}\right)(t_i) \tag{1}$$

where u is the machine utilisation and t_i and c_i are the mean and the coefficient variation of processing time of jobs for machine i.

- Evaluate TH with Little Law with the auctioned job j (TH_j) and without (TH).
- Then the evaluation of the acution job j, P_j is obtained:

$$P_j = \frac{TH_j - TH}{TH} \tag{2}$$

2.2.2. Symmetric Point Analysis Approach The second approach is based on symmetric point analysis. For each HLC, a matrix containing the processing times of the jobs in the JRQ for each machine is considered. The evaluation function considers two points in space: I, which represents the point of imbalance of the job in the JRQ without the auctioned job, and I_j, which represents the auctioned job j. The evaluation function follows these steps:

- Calculate the average processing times (t_i) of the jobs in the JRQ for each machine i. This results in a $[1] \times [m]$ vector T containing the average processing times for each machine:

$$T = [t_1, \ldots, t_m] \tag{3}$$

- Define the desired Imbalance point I as:

$$I = -1 \cdot T \tag{4}$$

- Evaluate the distance P_l between I and I_j:

$$P_j = \sqrt{(i_{j,1} - t_1)^2 + (i_{j,2} - t_2)^2 + \cdots + (i_{j,m} - t_m)^2} \tag{5}$$

where $i_{j,k}$ represents the component of I for machine i and t_i represents the average processing time for each machine.

2.3. Bidding Function

The bidding function determines the value that each HLC proposes to an auctioned job based on the evaluation results and the HLC's current production goals and constraints. The bidding function considers three components: the evaluation result, the relative wealth of the HLC (R), and the ratio of the target cycle time to the actual cycle time (VAR_{CT}). The bid value is determined as follows:

$$Bid = P_j \cdot R \cdot VAR_{CT} \tag{6}$$

3. Conclusion

This work presents an auction-based mechanism for job allocation in hybrid manufacturing planning and control architectures. The proposed approach leverages the principles of auction theory and stochastic modelling to enhance decision-making and operational flexibility within semi-heterarchical structures. A preliminary simulation study evaluated the effectiveness of the auction-based mechanism across various scenarios with different distributions of job processing times and production line configurations. The results demonstrated that the auction-based mechanism outperformed traditional job allocation approaches, particularly in scenarios with high variability in job processing times. The mechanism's ability to dynamically allocate jobs based on their potential impact on production performance led to improved throughput and resource utilisation. The two evaluation approaches proposed in this work proved to be effective in capturing the stochastic nature of job processing times and providing valuable insight for informed decision making.

Acknowledgments

This study was carried out within the MICS (Made in Italy – Circular and Sustainable) Extended Partnership and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1551.11-10-2022, PE00000004). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

References

Antons, Oliver, Julia C. Arlinghaus. 2022. Distributing decision-making authority in manufacturing–review and roadmap for the factory of the future. *International Journal of Production Research* **60**(13) 4342 – 4360. doi:10.1080/00207543.2022.2057255.

Borangiu, Theodor, Damien Trentesaux, Andr'e Thomas, Paulo Leit ao, Jos'e Barata. 2020. Digital transformation of manufacturing through cloud services and resource virtualization. *Computers in Industry* **108** 150–162.

Grassi, Andrea, Guido Guizzi, Liberatina Carmela Santillo, Silvestro Vespoli. 2020. A semi-heterarchical production control architecture for industry 4.0-based manufacturing systems. *Manufacturing Letters* **24** 43–46. doi:https://doi.org/10.1016/j.mfglet.2020.03.007.

Ivanov, Dmitry, Alexandre Dolgui, Boris Sokolov. 2018. A digital supply chain twin for managing the disruption risks and resilience in the era of industry 4.0. *Production Planning & Control* **29**(9) 775–788.

Design of unreliable flow lines: How to jointly allocate buffer space and spare parts

Florian E. SACHS
University of Cologne, Faculty of Management, Economics and Social Sciences, Albertus-Magnus-Platz, 50923 Cologne, Germany
Technical University of Munich, TUM School of Management, Bildungscampus 9, 74076 Heilbronn, Germany,
florian.sachs@uni-koeln.de

Gudrun P. KIESMÜLLER
Technical University of Munich, TUM School of Management, Bildungscampus 9, 74076 Heilbronn, Germany,
gudrun.kiesmueller@tum.de

Stefan HELBER
Leibniz University Hannover, Faculty of Economics and Management, Königsworther Platz 1, 30167 Hannover, Germany,
stefan.helber@prod.uni-hannover.de

The design of flow lines is an extensive area of research whereby most publications focus on buffer allocation to cope with random influences like breakdowns or random processing times. In addition, spare parts can increase the machines' availability directly. However, only a few papers consider the simultaneous optimization of buffer space and the possible number of spare parts in stock. The literature hardly even considers the evaluation of such systems. We are the first to introduce the joint optimization of buffer space and spare parts for flow lines of arbitrary length. First, we aim to allocate buffer capacities and spare parts efficiently. Since the buffer allocation problem is NP-hard, we can expect only to find near-optimal solutions. Second, we demonstrate the algorithmic behavior of different greedy and metaheuristics on this design problem. We illustrate how to exploit the problem structure to solve it almost optimally. Third, we generate managerial insights into allocating spare parts in manufacturing systems with buffers. We show that spare parts tend to be more effective when arranged at or near the center of a flow line, as it is already known for buffers. Moreover, we provide details on the combined buffer and spare part allocations.

Key words: Manufacturing; spare parts; stochastic flow lines; buffer allocation problem; heuristics

1. Research problem

A principle in the analysis of manufacturing systems is the installation of buffers between sequential machines to decouple their production processes coping with machine downtimes or differences in processing times. However, installing buffers could require high investments (see, e.g., Liberopoulos 2018), although buffers do not directly tackle machine downtimes. Nevertheless, it is well-known that downtimes cause high costs, too. Recent publications (see, e.g., Sachs et al. 2022) analyzed the impact of spare part provisioning on manufacturing systems to address this problem. There is, however, no publication on the problem of how buffers and spare parts should be allocated jointly in long flow lines.

This project contributes to the research area of manufacturing systems in three ways: First, we set up a new, general optimization problem that is highly relevant for the praxis but has never been studied before. Second, we introduce solution methods and how they need to be adjusted to solve this specific problem. Third, we demonstrate how the simultaneous optimization of buffer capacities and spare part stocks under cost considerations influences the design of flow lines.

2. Problem assumptions

Our analysis is based on an I-machines flow line as depicted in Figure 1 with $J := I - 1$ intermediate buffers B_j each with an capacity of $C_j, j = 1, \dots, J$. The manufacturing system processes discrete, non-perishable parts and produces a single product. We assume a saturated system, i.e., an infinite supply of raw material and infinite storage capacity for end products. Usually, machines for different production steps are close by so that we can neglect transportation times. We assume that all random times are mutually stochastically independent.

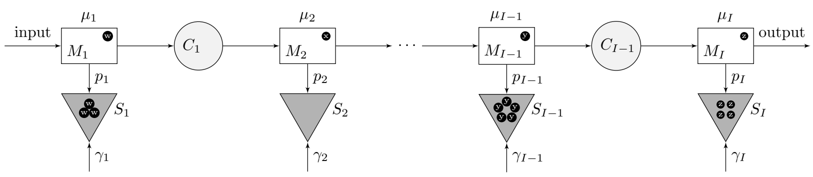

Figure 1. Model of the production system

Each of the machines contains one unit of a machine-specific, failure-prone critical component, i.e., a failure of such a component leads to a downtime of the specific machine. Each build-in component can only fail while the machine is producing (operation-dependent failures). We assume exponentially distributed times to failure with a failure rate of $p_i, i = 1, \dots, I$. The exponential distribution is well-suited for components not suffering from wear-out effects. If a spare part is available, we assume a repair-by-replacement strategy. Otherwise, an order is placed, and the machine is nonfunctional until the spare arrives.

The stock-keeping of spare parts is organized via a one-for-one replenishment policy, which is reasonable as fixed ordering costs are neglectable compared to holding costs. Hence, we have the basestock level S_i per component type and at maximum $Q_i = S_i + 1$ units of the component of machine i. Ordered units have an exponentially distributed procurement time with replenishment rate $\gamma_i, i = 1, \dots, I$. This assumption does not hamper the applicability of our results since only the mean of the distribution matters (Palm-Khintchine theorem).

We model the processing time exponentially distributed with processing rate $\mu_i, i = 1, \dots, I$. We apply blocking after service, i.e., a machine is blocked after the production process finishes.

3. Key methodology

We formulate the primal buffer and spare part allocation problem (BSAP-P) based on the primal buffer allocation problem (BAP-P) according to the notation of Weiss et al. (2019) as

$$\min_{\substack{C_1,\dots,C_J \\ Q_1,\dots,Q_I}} \sum_{j=1}^{J} c^j_{\text{buffer}} \cdot C_j + \sum_{i=1}^{I} c^i_{\text{spare}} \cdot Q_i \tag{1}$$

$$\text{s.t.} \quad TP(\mathbf{C}, \mathbf{Q}) \geq TP^T, \tag{2}$$

$$0 \leq C_{\min} \leq C_j \leq C_{\max}, \qquad C_j \in \mathbb{N}_0 \qquad \forall j = 1, \dots, J \tag{3}$$

$$1 \leq Q_{\min} \leq Q_j \leq Q_{\max}, \qquad Q_i \in \mathbb{N} \qquad \forall i = 1, \dots, I \tag{4}$$

where $TP(\mathbf{C}, \mathbf{Q})$ is the expected throughput, TP^T is the target throughput, $C_{\min}$ and $Q_{\min}$ ($C_{\max}$ and $Q_{\max}$) are the lower (upper) bounds for buffer capacities and number of units. $\mathbf{C}$ and $\mathbf{Q}$ are vector representations of $C_j, j = 1, \ldots, J$, and $Q_i, i = 1, \ldots, I$.

The evaluation of a specific flow line uses a recently published decomposition approach. Due to the included stochasticity, we apply two pseudo-gradient-based approaches for optimization. The first one starts with a minimal system where $C_j = 1, \forall j = 1, \ldots, J$, (due to numeric stability) and $Q_i = 1, \forall i = 1, \ldots, I$, and increases one component per iteration step. Accordingly, $\mathbf{C}^{\text{new}}$ and $\mathbf{Q}^{\text{new}}$ originate from $\mathbf{C}$ and $\mathbf{Q}$ by increasing one single component by one. We choose to increment the buffer $C_j, j = 1, \ldots, J$, or the spare part stock $Q_i, i = 1, \ldots, I$, by maximizing

$$\max \left\{ \frac{TP(\mathbf{C}^{\text{new}}, \mathbf{Q}^{\text{new}}) - TP(\mathbf{C}, \mathbf{Q})}{\sum_{j=1,\ldots,J} c^j_{\text{buffer}} \cdot (C^{\text{new}}_j - C_j) + \sum_{i=1,\ldots,I} c^i_{\text{spare}} \cdot (Q^{\text{new}}_i - Q_i)} \right\}, \tag{5}$$

thus using the maximum throughput-per-cost increase. In addition, we apply a decrasing greedy heuristic, simulated annealing, and a genetic algorithm to check whether these algorithms find better solutions (due to the integer optimization problem).

4. Major results

We can validate that the pseudo-gradient approaches perform excellent compared to a brute-force optimization (enumerating the whole decision space) and find optimal results in the tested instances with three-machine systems. Furthermore, we can show for larger systems that our pseudo-gradient results cannot be improved via the analyzed meta-heuristics.

The simultaneous optimization realizes significant cost benefits. Often, only spare parts cannot achieve target throughputs since they do not cope with variability in processing times. Hence, buffers are essential for high-throughput flow lines.

In balanced settings (identical machine parameters, identical costs), the spare part allocation shows a bowl pattern like buffers are known to do. In unbalanced settings, we observe spare parts to be allocated at or around bottlenecks even if the costs are relatively high. The specific allocation is highly cost-dependent.

All in all, we found that well-known buffer-allocation characteristics also hold for manufacturing systems with spare parts. Furthermore, we are the first to observe and explain spare part allocation in flow lines. In specific, we find that some results from the spare part literature still hold, whereas rules of thumb cannot explain the complex interaction between buffer places and spare parts.

References

Liberopoulos, George. 2018. Performance evaluation of a production line operated under an echelon buffer policy. *IISE Transactions* **50**(3) 161–177. doi:10.1080/24725854.2017.1390800.

Sachs, Florian E, Stefan Helber, Gudrun P Kiesmüller. 2022. Evaluation of unreliable flow lines with limited buffer capacities and spare part provisioning. *European Journal of Operational Research* **302**(2) 544–559. doi:10.1016/j.ejor.2022.01.005.

Weiss, Sophie, Justus A Schwarz, Raik Stolletz. 2019. The buffer allocation problem in production lines: Formulations, solution methods, and instances. *IISE Transactions* **51**(5) 456–485. doi:10.1080/24725854.2018.1442031.

Digital Twin for Virtual Learning - Industry Case

Young Jae Jang and Taejoon Kim
Affiliation 1[Industrial and System Engineering, KAIST], Daejeon, South Korea, yjang@kaist.ac.kr

Sangpyo Hong
Affiliation 2[Research and Development, DAIM Research], Seoul, South Korea, sphong@daimresearch.com

This paper explores the implementation of a Digital Twin for Virtual Learning (DT-VL) system, developed by DAIM Research and KAIST, in a semiconductor manufacturing facility. Utilizing a reinforcement learning algorithm within a Massive Fleet Robot Agent (MFRA) framework, the system optimizes robot routing and allocation in Automated Material Handling Systems (AMHS), particularly addressing Overhead Hoist Transportation (OHT) challenges. The deployment demonstrated significant improvements, including a 32% enhancement in delivery times and a 20% increase in operational capacity, enabling the facility to operate efficiently with fewer robots and achieve substantial cost savings, thereby highlighting the efficacy of integrating advanced digital twins and machine learning technologies in industrial operations.

1. Massive Fleet Robot Agent (MFRA) – AMHS

Massive Fleet Robot Agent (MFRA) refers to extensive collaboration among autonomous robots, aiming to achieve a shared objective. The term ”massive” emphasizes the fleet’s size, with each ”robot agent” independently performing tasks. The Overhead Hoist Transportation (OHT) system, used in semiconductor fabrication facilities or FABs, is a prime example of a Massive Fleet Robot Agent (MFRA) based Automated Material Handling System (AMHS). These systems utilize numerous vehicles to lift and move batches between different machines. Modern FABs often employ over a thousand OHT vehicles, highlighting the vast scale and complexity of these operations (Hong et al. (2022)).

2. DIGITAL TWIN FOR VIRTUAL LEARNING

We introduce the application of a reinforcement learning algorithm within the MFRA-AMHS framework, illustrating how digital twin technology enhances the effectiveness of the reinforcement learning process. The algorithms and approaches employed are based on previous research including Hong et al. (2022) and Hwang and Jang (2020).

Q-learning, a reinforcement learning method, trains machine learning models through trial and error interactions with its environment, adjusting strategies based on outcomes Sutton and Barto (2018).

In high-tech manufacturing, especially with MFRA-AMHS, robots follow predefined routes to navigate efficiently and safely within spatial constraints. This setup, similar to urban traffic control, ensures both efficiency and safety, preventing accidents on the factory floor.

For modeling, MFRA-AMHS systems are seen as directional networks of nodes and edges, representing various points robots travel through. Control is managed node-by-node, where decisions on routing and task allocation are made based on the optimal path calculations using the Q-function, which estimates travel times between nodes. This process ensures smooth flow and effective task assignment within the manufacturing system.

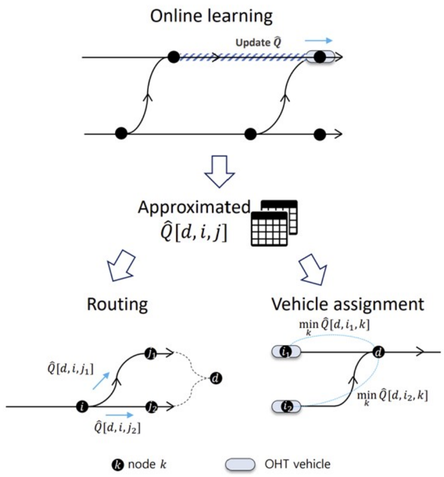

Figure 1. Overall Q-function value-based robot routing and control scheme

Routing decisions are guided by Q-function values to minimize travel time and efficiently allocate tasks, demonstrating an organized method for managing large fleets of robots within complex manufacturing setups. For example, consider a scenario where a robot, destined for node d, approaches node i. The robot must choose between edge (i, j) and edge (i, k). The decision entity at node i compares $Q_i(d, j)$ and $Q_i(d, k)$ to determine which edge represents the shortest estimated travel time. The edge with the lower Q-function value is chosen, guiding the robot along the optimal route.

Robot allocation also leverages the Q-function values. Suppose a new task at node d requires a robot for transportation. Node entities at nodes $i = 1, \ldots, n$, where idle robots are stationed, initially calculate $c_i = \min_k Q_i(d, k)$ to identify which robot can reach the load in the shortest time. The optimal robot is determined by $i^* = \arg\min_i c_i$, and it is then assigned to transport the new load at node d. Readers interested in a more comprehensive understanding of the methodologies and empirical results are encouraged to consult the following articles: Hong et al. (2022) and Hwang and Jang (2020). Factories are dynamic entities that must adapt quickly to shorter product cycles and new technologies. As machines are upgraded and new products launched, both the factory layout and the flow of parts change, requiring robots' paths to be adjusted. This adjustment is streamlined through the Digital Twin for Virtual Learning (DT-VL), which mirrors changes in the factory setup and updates the robots' Q-tables accordingly.

The DT-VL serves as a virtual model of the MFRA system, promptly reflecting changes like new machines or products that impact robot movement and operations. It simulates these updates in advance, ensuring that robots function seamlessly when production changes occur. This approach, known as Zero-shot Learning, emphasizes fast simulations to update learning parameters virtually before applying them in the real system, significantly reducing the costs and time associated with physical trials Haarnoja et al. (2023).

3. Industry Case

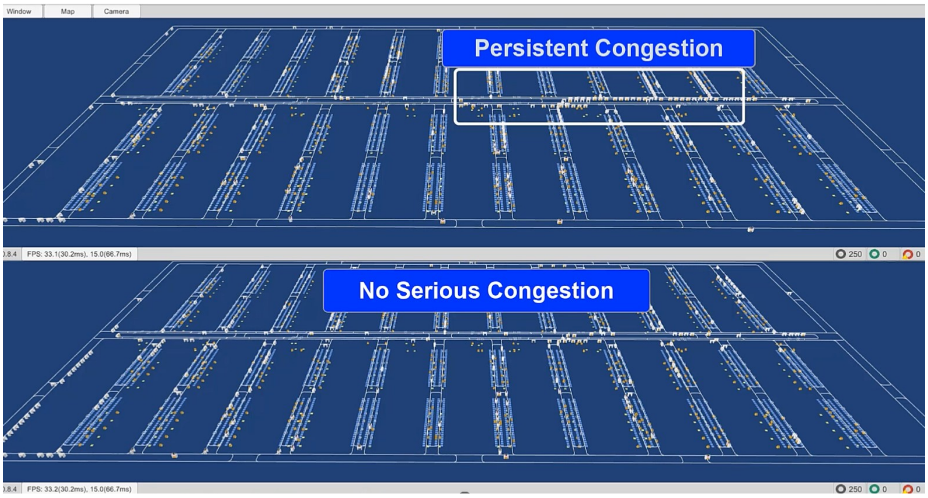

Figure 2. Movie clip of the system operating 1,000 OHT robots

DAIM Research and KAIST developed an Overhead Hoist Transportation (OHT) system for a Korean memory chip manufacturer, an application of the Massive Fleet Robot Agent (MFRA) system critical in semiconductor production. A movie clip 2 featured with this article shows the deployment of 1,000 robots and the traffic congestion challenges inherent in such dense robotic operations, where traditional manual interventions often failed.

To address these issues, the chip manufacturer collaborated with DAIM Research and KAIST, implementing a Reinforcement Learning (RL) algorithm coupled with a Digital Twin for Virtual Learning (DT-VL). This solution effectively managed robot assignments and alleviated congestion, as demonstrated in the movie clip (Screen shot depicted in Figure 2).

The DT-VL system's value is highlighted by the need for rapid adaptation in semiconductor factories, where updates are frequent. The implementation of the RL and DT-VL solution markedly improved robot efficiency with delivery times and capacity improving by 32% and 20%, respectively. This efficiency allowed the factory to reduce its robot fleet to 800, saving USD 16 million based on the cost of each robot.

References

Haarnoja, Tuomas, Ben Moran, Guy Lever, Sandy H Huang, Dhruva Tirumala, Jan Humplik, Markus Wulfmeier, Saran Tunyasuvunakool, Noah Y Siegel, Roland Hafner, et al. 2023. Learning agile soccer skills for a bipedal robot with deep reinforcement learning. *arXiv preprint arXiv:2304.13653* .

Hong, Sangpyo, Illhoe Hwang, Young Jae Jang. 2022. Practical q-learning-based route-guidance and vehicle assignment for oht systems in semiconductor fabs. *IEEE Transactions on Semiconductor Manufacturing* **35**(3) 385–396.

Hwang, Illhoe, Young Jae Jang. 2020. Q (λ) learning-based dynamic route guidance algorithm for overhead hoist transport systems in semiconductor fabs. *International Journal of Production Research* **58**(4) 1199–1221.

Sutton, Richard S, Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.

Energy-Efficiency Control of Manufacturing Systems via Active Inference

Yavar Taheri Yeganeh
Department of Mechanical Engineering, Politecnico di Milano, Milan, Lombardy, Italy, yavar.taheri@polimi.it

Mohsen A. Jafari
Department of Industrial and Systems Engineering, Rutgers University, Piscataway, New Jersey, United States, jafari@soe.rutgers.edu

Andrea Matta
Department of Mechanical Engineering, Politecnico di Milano, Milan, Lombardy, Italy, andrea.matta@polimi.it

We are exploring active inference theory to develop intelligent decision-making models for optimizing energy-efficient control in manufacturing systems. Inspired by insights from neuroscience, active inference provides a unified and probabilistic framework that integrates perception, learning, and action. It represents an emerging domain in artificial intelligence, bridging generative models with decision-making processes. Utilizing a deep active inference agent, we investigate potential control strategies in parallel and identical machine workstations, with a specific focus on promoting sustainability in manufacturing systems. We initially concentrate on stationary manufacturing environments, and subsequently extend our analysis to non-stationary cases. We leverage advancements in previously developed active inference agents, along with existing reinforcement solutions, for the systems under study. Our study compares the performance of the active inference-based method with reinforcement learning to evaluate the advancement of the proposed methodology.

Key words: Active Inference; Reinforcement Learning; Energy-Efficient Control; Manufacturing Systems

---

1. Introduction

Energy-efficient control (EEC) is vital in the manufacturing sector due to its substantial impact on global energy consumption. EEC strategies focus on minimizing energy usage by optimizing machine states, particularly during idle periods (Loffredo et al. 2023). Traditional EEC methods often require complete system knowledge, which is impractical in dynamic, real-world environments. Reinforcement learning has shown promising performance in optimizing manufacturing processes without prior system knowledge (Loffredo et al. 2023) but struggles with rapid adaptation to changing conditions. Active inference, based on the free energy principle (FEP), offers an alternative by unifying perception, learning, and decision-making under uncertainty through a Bayesian framework (Friston 2010). It has been applied successfully in complex decision-making tasks in fields such as robotics, enabling agents to navigate uncertain and dynamic scenarios (Pezzato et al. 2023). This research aims to build upon advancements in active inference-based decision-making and apply it to EEC in manufacturing systems to demonstrate its potential.

2. Active Inference Agent

The active inference framework posits that organisms actively interact with their environment by updating beliefs and actions based on sensory inputs to reduce *surprise*. Active inference agents have an internal generative model parameterized by θ that interacts with the world similarly to a

Partially Observable Markov Decision Process (POMDP). The core concept is FEP, which leads to optimizing the model by minimizing Variational Free Energy (VFE) to reduce *surprise*, quantified by $-\log P_\theta(o_t)$, as follows (Fountas et al. 2020):

$$\theta^* = \arg\min_\theta \left(\mathbb{E}_{Q_\phi(s_t,a_t)} \left[\log Q_\phi(s_t, a_t) - \log P_\theta(o_t, s_t, a_t) \right] \right). \tag{1}$$

The agent's actions aim to minimize Expected Free Energy (EFE or G), which is calculated during planning by simulating future trajectories, π, up to a horizon $\tau \geq t$ (Fountas et al. 2020):

$$G(\pi, \tau) = \mathbb{E}_{P(o_\tau|s_\tau,\theta)} \mathbb{E}_{Q_\phi(s_\tau,\theta|\pi)} \left[\log Q_\phi(s_\tau, \theta|\pi) - \log P(o_\tau, s_\tau, \theta|\pi) \right]. \tag{2}$$

EFE comprises three terms: Reward-like expected *surprise*, which pertains to how close the future predictions are to the preference, state uncertainty, and model parameter uncertainty. The framework (as depicted in Fig. 1) includes first optimizing model parameters θ to fit observations based on VFE in Eq. 1 and then making decisions based on negative accumulated EFE in Eq. 2, incorporating a softmax function. The active inference agent architecture includes encoder, transition,

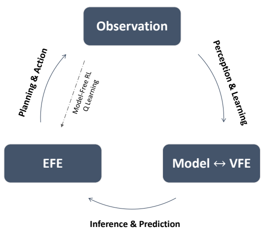

Figure 1. Active Inference Framework

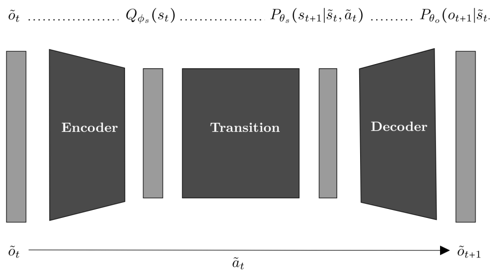

Figure 2. Agent Architecture

and decoder modules (as depicted in Fig. 2), implemented with neural networks. This architecture is probabilistic and similar to a variational autoencoder. Each module generates parameters of pre-selected distributions given a sample. Calculating EFE for all possible trajectories is infeasible, so Fountas et al. (2020) proposed using Monte Carlo Tree Search (MCTS) coupled with an inference action module (i.e., $Q_{\phi_a}(a_t)$). This module approximates the posterior distribution over actions using the prior obtained from the MCTS. Building on the proposed agent by Fountas et al. (2020), we explored various perspectives to design an agent for the EEC task in a manufacturing system. However, there are peculiar features that challenge the original agent's effectiveness. The system is highly stochastic with delays in responding to agent policy, reflected in reward functions that measure average performance (e.g., throughput) over a long horizon. This complicates learning dynamics and planning strategy. Notably, the agent architecture (Fig. 2) predicts one step ahead, often similar to the previous observation due to stochasticity. Therefore, we introduced modifications to tailor the architecture to our problem. To address the limitations of finite horizon EFE, we propose a hybrid architecture that incorporates longer horizons via deep Q-learning. This approach balances short-term EFE and long-term considerations using a hyperparameter, γ. Instead of using the previous habitual structure within Monte Carlo Tree Search (MCTS) as described by Fountas (2020), we trained $Q_{\phi_a}(a_t)$ based on deep Q-learning. Additionally, we modified the transition module to allow multiple steps, controlled by a hyperparameter (e.g., s = 90), enabling multi-step predictions. Due to the computational expense of MCTS, we replaced it with repeated actions in the transition and calculated EFE. This method assesses the impact of actions over a short period, using repeated action simulations at every decision step.

3. Results and Conclusion

Our experiments focused on controlling a real industrial workstation consisting of six parallel-identical machines with finite upstream capacity buffer, as described in (Loffredo et al. 2023). For the reward, similar to (Loffredo et al. 2023), we balanced ($\phi = 0.97$) the ratio of throughput and energy consumption against the *ALL ON* policy over the past 8 hours. We considered both 1-step transitions and multi-step transitions, taking repeated actions during planning for each of the possible policies (i.e., determining how many machines to keep *ON*) to then calculate their EFE. We tested the performance of our agents 50 times during different training iterations, each on independent systems initialized with a random agent after warm-up. We evaluated the performance of our agents using metrics such as test reward, throughput loss, total energy savings, and energy savings per part percentage compared to the *ALL ON* policy over an 8-hour window. Fig. 3 presents the comparison for a single set of hyperparameters except s and γ, while Table 1 shows the performance and average learned policy distribution for specific γ values. These results demonstrate the efficacy of our introduced modifications. Hyperparameters significantly influence agent performance, highlighting the importance of proper tuning for notable improvements and effective control. This underscores the potential of our proposed methodology for EEC applications. Importantly, the framework and formalism of active inference agents exhibit notable promise for non-stationary scenarios, where model-free agents may struggle to adapt swiftly. Future work will concentrate on extending experiments and tailoring the methodology for such non-stationary scenarios.

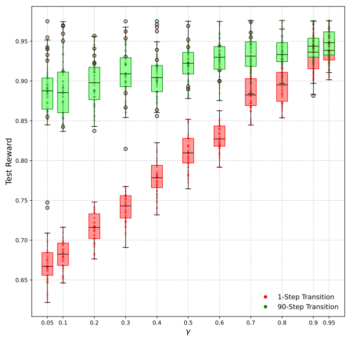

Figure 3. Results

Performance Metrics (%)				
γ	Reward	TH Loss	EN Saving	EN/Part Saving
0.05	89.05±0.86	8.88±0.45	10.78±0.50	2.13±0.10
0.3	91.11±0.80	6.71±0.40	8.83±0.47	2.30±0.12
0.6	92.88±0.75	4.84±0.35	7.94±0.42	3.29±0.15

Policy Probabilities (A_n: Keeping n Machines *ON*)							
γ	$P(A_0)$	$P(A_1)$	$P(A_2)$	$P(A_3)$	$P(A_4)$	$P(A_5)$	$P(A_6)$
0.05	0.00±0.00	0.00±0.00	0.07±0.00	0.18±0.01	0.22±0.01	0.27±0.02	0.26±0.01
0.3	0.00±0.00	0.00±0.00	0.05±0.00	0.13±0.01	0.17±0.01	0.38±0.02	0.27±0.01
0.6	0.00±0.00	0.00±0.00	0.03±0.00	0.07±0.01	0.10±0.01	0.57±0.02	0.23±0.01

Table 1. Performance and the Learned Policy for 90-Step Transition.

References

Fountas, Z., Noor Sajid, Pedro A. M. Mediano, Karl J. Friston. 2020. Deep active inference agents using monte-carlo methods. *ArXiv* **abs/2006.04176**.

Friston, Karl. 2010. The free-energy principle: a unified brain theory? *Nature reviews neuroscience* **11**(2) 127–138.

Loffredo, Alberto, Marvin Carl May, Louis Schäfer, Andrea Matta, Gisela Lanza. 2023. Reinforcement learning for energy-efficient control of parallel and identical machines. *CIRP Journal of Manufacturing Science and Technology* **44** 91–103.

Pezzato, Corrado, Carlos Hernández Corbato, Stefan Bonhof, Martijn Wisse. 2023. Active inference and behavior trees for reactive action planning and execution in robotics. *IEEE Transactions on Robotics* **39**(2) 1050–1069.

Energy-Efficient Resource Scheduling of a Single Resource with Production Requirements: A Joint Simulation and Scenario-Based Mathematical Programming Approach

Siamak Khayyati
HEC Liege, Management School of the University of Liege, Liege, Belgium
s.khayyati@uliege.be

Oktay Karabağ
Department of Industrial Engineering, İzmir University of Economics, Sakarya Caddesi No:156, Balçova, İzmir, Turkey
oktay.karabag@ieu.edu.tr

Berk Berkman, Şebnem Demirtaş
College of Engineering, Koç University, Rumeli Feneri Yolu, İstanbul, Turkey
bberkman20@ku.edu.tr, sdemirtas20@ku.edu.tr

Barış Tan
Faculty of Engineering & Faculty of Business, Özyeğin University, Çekmeköy Campus Nişantepe District, Orman Street, Çekmeköy, Istanbul, 34794, Türkiye

Improving energy efficiency in manufacturing yields significant cost and environmental benefits. Advanced data collection and execution technologies allow the implementation of data-driven dynamic control policies that turn on and off resources depending on the real-time data to save energy. Motivated by an automotive producer's paint oven on-off scheduling, we consider a single resource that operates in on, off and warmup modes. When the resource is on, it can be turned off immediately. However, when the turned-off resource is turned on, there is a delay for warm-up. The energy consumption in the off mode is the lowest. We present a scenario-based mathematical programming formulation to determine the optimal on-off schedule of a single resource that minimizes the average energy and early/late production costs while meeting the production requirements. We employ a cutting plane algorithm to solve this problem. We discuss using the scenario-based formulation with the random arrival times generated by a detailed digital twin of the production system feeding the resource.

Key words: Optimal production and energy control; Energy Efficiency; Resource scheduling

1. Introduction

The global push for decreasing green-house gas emissions to limit its effect on climate change has motivated manufacturers to focus on decreasing their energy consumption since a considerable portion of energy consumption globally can be attributed to manufacturing (Wang et al. 2017). From the energy consumed in manufacturing, a large portion is used in heating processes (Yan and Zheng 2020).

The contributions of this work are introducing a mathematical programming formulation for scheduling the on/off cycles for an energy intensive operation in order to decrease energy consumption while full-filling production requirements and giving a cutting plane algorithm that uses a scenario clustering approach to solve the problem in efficient time.

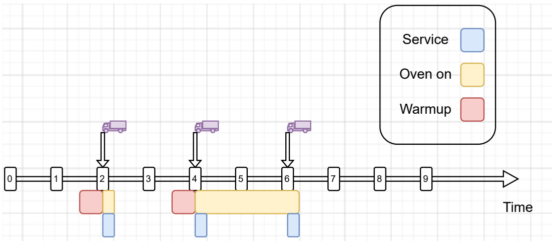

Figure 1. An example schedule that uses two cycles. In this case, the vehicles are admitted to the oven and their service starts as soon as they arrive in the system

2. Problem Definition

The setting of our problem is based on the problem of controlling the energy consumption in a major automotive manufacturing plant in Turkey. The energy consumption in the paint ovens constitutes a large portion of the production cost of the vehicles and an even larger portion of the total energy consumption in the plant. The ovens are among the last stages in the manufacturing process, and for this reason, any planned turning off of these resources can resulting in blocking upstream. The vehicles are categorized into vehicles with regular colors and special colors. The throughput for the regular color vehicles is high enough that any planned turning off of ovens cannot be done without compromising on the throughput. However, the utilization for the special color vehicles is lower, providing the opportunity for planned turning off of the ovens.

We consider a single resource that operates in different modes: on, off, and warmup. When the resource is on, it can be turned off immediately. However, a warmup is initiated when the turned-off resource is turned on. The warmup delay is deterministic and given as t^{WU}. The energy consumption in the off mode is the lowest and is set as 0. The energy consumption in the warmup mode c_{WU} can be higher than in the on mode c_{ON}. An on-off cycle is determined by the time the resource is turned on U_c and then turned off W_c for cycle c. An on-off schedule (U_c, W_c), $c = 1, \ldots, C$ is a collection of on-off cycles during the planning horizon.

The jobs arrive randomly at the resource at different times. The arrival times A_j for J jobs are given. There is a requirement for the number of jobs that must be completed

The problem we consider is determining an on-off schedule for the planning horizon to minimize the total average energy and waiting cost, TC. Figure 1 depicts a sample on-off schedule that uses two cycles in the planning horizon. We incorporate stochasticity in the arrival time of the parts to the system through solving the problem for multiple scenarios simultaneously. To improve the time performance of this mathematical programming formulation, we use a cutting plane algorithm that adds constraints related to preventing production when the oven is off iteratively.

$$\min \quad \text{TC} = c_{WU}\sum_{c\in\mathbb{C}}(V_c - U_c) + c_{ON}\sum_{c=1}^{C}(W_c - V_c) + c_{WT}\sum_{j\in\mathbb{J}}(S_j - A_j) \tag{1a}$$

$$\text{s.t.} \quad U_c \geq 0, \qquad c \in \mathbb{C}, \tag{1b}$$

$$V_c - U_c \geq t^{WU} Z_c, \qquad c \in \mathbb{C}, \tag{1c}$$

$$
\begin{aligned}
&W_c \geq V_c && c \in \mathbb{C}, && \text{(1d)}\\
&S_j \geq V_c - M(1 - X_{cj}), && j \in \mathbb{J}, c \in \mathbb{C}, && \text{(1e)}\\
&S_j \geq A_j, && j \in \mathbb{J}, && \text{(1f)}\\
&F_j = S_j + t_j^{OT}, && j \in \mathbb{J}, && \text{(1g)}\\
&S_{j+1} \geq F_j, && j \in \mathbb{J}, && \text{(1h)}\\
&W_c \geq F_j - M(1 - X_{cj}), && j \in \mathbb{J}, c \in \mathbb{C}, && \text{(1i)}\\
&W_c \leq T_s, && c \in \mathbb{C} && \text{(1j)}\\
&U_{c+1} \geq W_c, && c \in \mathbb{C}, && \text{(1k)}\\
&Z_c \leq \sum_j X_{cj}, && c \in \mathbb{C}, && \text{(1l)}\\
&\sum_j X_{cj} \leq M Z_c, && c \in \mathbb{C}, && \text{(1m)}\\
&\sum_c X_{cj} = 1, && j \in \mathbb{J}, && \text{(1n)}\\
&X_{cj}, Z_c \in \{0, 1\}, && j \in \mathbb{J}, c \in \mathbb{C}, && \text{(1o)}
\end{aligned}
$$

For settings similar to the problem here, threshold-type policies have been shown to perform well (Tan et al. 2023).We use a threshold-type policy as a benchmark. The threshold policy utilizes two thresholds, one for turning the oven off when there are few vehicles to process and one for turning the oven on when many vehicles are waiting to be processed. For the purpose of assessing the performance of these methods we use a simulation model of the system that captures the dynamics of the vehicles in multiple upstream locations. Given the location of each vehicle in the system and the traffic of vehicles in front of it, the arrival time of the vehicle to the oven can be estimated and the mathematical programming framework works based on those estimates.

3. Results

In this work, we devise a mathematical programming formulation that generates the optimal schedule for controlling an energy-intensive operation in order to minimize energy consumption subject to stochasticity in part arrivals and production rate constraints. We have conducted numerical experiments that show the cutting plane algorithm can effectively reduces the solution time for the problem when multiple scenarios are considered.

Acknowledgments
This work was supported by TUBITAK (grant number 221M393).

References

Tan, Barış, Oktay Karabağ, Siamak Khayyati. 2023. Energy-efficient production control of a make-to-stock system with buffer-and time-based policies. *International Journal of Production Research* 1–19.

Wang, Junfeng, Yi Feng, Zicheng Fei, Shiqi Li, Qing Chang. 2017. Markov chain based idle status control of stochastic machines for energy saving operation. *2017 13th IEEE conference on automation science and engineering (CASE)*. IEEE, 1019–1023.

Yan, Chao-Bo, Ziqian Zheng. 2020. Problem formulation and solution methodology for energy consumption optimization in bernoulli serial lines. *IEEE Transactions on Automation Science and Engineering* **18**(2) 776–790.

Exploiting Production and Maintenance Scheduling with Deep Reinforcement Learning in Stochastic Environments

Maria Grazia MARCHESANO, Guido GUIZZI, Giuseppe CONVERSO, Liberatina Carmela SANTILLO

Dipartimento di Ingegneria Chimica, dei Materiali e della Produzione Industriale, Universita degli Studi di Napoli Federico II, ‘ Piazzale Tecchio 80, 80125 Napoli, Italy. {mariagrazia.marchesano, g.guizzi, giuseppe.converso, santillo}@unina.it

The landscape of operations management research is transforming, catalysed by the integration of advanced computational methodologies. Notably, Deep Reinforcement Learning (DRL) emerges as a groundbreaking tool, offering unprecedented capabilities in addressing intricate operational challenges. This study explores the application of DRL within the framework of operations management, particularly focusing on areas traditionally constrained by computational limitations or complexity. We delve into the application of DRL in diverse operational scenarios, including dynamic resource allocation, predictive maintenance, and production scheduling characterized by stochastic processing times. Through a series of experimental investigations, DRL has been proved to not only surpass conventional approaches in handling multifaceted problems, but also unveil new potentials for efficiency and effectiveness in handling operations considering stochastic variables. Our findings provide a comprehensive overview of the practical applications of DRL, its efficacy in capturing the intricacies of stochastic models, and its current boundaries in the field of operations management. We conclude with a balanced and realistic evaluation, recognizing the substantial promise of DRL in advancing more responsive and smart operational systems. At the same time, we emphasize the ongoing need for research and innovation to thoroughly assess the potential of DRL in the multifaceted and dynamic challenges of operations management.

Key words: Deep Reinforcement Learning; Production Planning and control; Maintenance

1. Introduction

In the fast-paced world of modern manufacturing, strategic maintenance planning is crucial to achieving operational optimization. Efficient resource allocation and system reliability are essential for maintaining high performance in manufacturing processes (Akl et al. 2022). In particular in complex flow shop systems, where a fixed sequence of processing tasks is required, production and maintenance scheduling becomes imperative to maintain system performance (Huang et al. 2020). The dynamic nature of these systems requires quick responses to minimize downtime and prevent equipment failures, ensuring a seamless production continuum (Hu et al. 2022).

Artificial Intelligence (AI), and more specifically Deep Reinforcement Learning (DRL), a subfield of machine learning, has emerged as a powerful tool in navigating the complexities of maintenance scheduling. DRL has demonstrated significant potential in optimizing manufacturing systems, enhancing scheduling, system optimization, and control approaches (Panzer and Bender 2021).

For example, Valet et al. (2022) proposed a deep RL-based opportunistic maintenance scheduling approach. Considering the operational status of the production system, this approach schedules

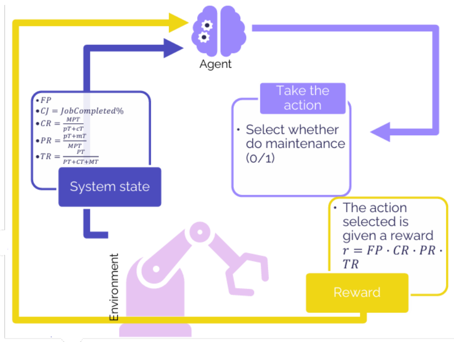

Figure 1. Proposed Approach

maintenance tasks during periods of low demand, minimizing the impact on production and improving maintenance effectiveness. Similarly, Yan et al. (2022) developed a double-layer Q-learning algorithm for dynamic scheduling with preventive maintenance in digital twin-enabled manufacturing systems. This approach considers the uncertainties and dynamics of the production system to maximize production efficiency while minimizing maintenance costs. Furthermore, Mao et al. (2022) developed a hash map-based memetic algorithm to minimize the total flow time of jobs while considering multiple maintenance activities. Although this approach did not explicitly integrate RL, the integration of such techniques could enhance the scheduling optimization process by enabling adaptive decision-making based on real-time data and system dynamics. Moreover, Kosanoglu et al. (2022) proposed a DRL-assisted simulated annealing algorithm for maintenance planning. This approach combines the benefits of both DRL and simulated annealing, allowing efficient exploration of the solution space and finding optimal maintenance schedules.

Despite the advantages of using DRL and AI in maintenance scheduling, several challenges remain. Nunes et al. (2023) reviewed the challenges in predictive maintenance, emphasizing the need for accurate data, effective machine learning models, and appropriate maintenance strategies. Advancements in AI techniques are crucial to overcoming these challenges and improving predictive maintenance practices in manufacturing systems.

2. Contribution of The Work

In this work, we introduce an integrated simulation tool and DRL algorithm to efficiently schedule and plan maintenance events in a production line flow shop. This approach combines the strengths of simulation and DRL to optimize maintenance processes and maximize productivity. The simulation tool creates a virtual environment that replicates the production line flow shop, allowing modeling and simulation of machine operations, job flows, and maintenance events. The DRL algorithm is the core intelligence of the proposed approach as it is displayed in Fgure 1. It learns optimal decision-making policies by interacting with the simulated environment and maximizing cumulative rewards. The algorithm considers machine failure probabilities, job priorities, and scheduling constraints to make informed decisions about maintenance scheduling. This multifaceted strategy not only underscores the novelty of our method but also its potential to contribute significantly to the field, bridging gaps identified in existing literature. By leveraging the strengths of AI and simulation, this approach offers a comprehensive solution to optimize maintenance processes, enhance resource allocation, and improve overall system reliability.

3. Conclusions and Future Research

Future research should continue to address the challenges of predictive maintenance and explore further integration of advanced AI techniques to improve decision making in dynamic manufacturing environments. Our findings provide a comprehensive overview of the practical applications of DRL, its efficacy in capturing the intricacies of stochastic models, and its current boundaries in the field of operations management. We conclude with a balanced and realistic evaluation, recognizing the substantial promise of DRL in advancing more responsive and smart operational systems, while emphasizing the ongoing necessity for research and innovation to thoroughly assess the potential of DRL in the multifaceted and dynamic challenges of operations management.

Acknowledgments

This study was carried out within the MICS (Made in Italy – Circular and Sustainable) Extended Partnership and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1551.11-10-2022, PE00000004). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

References

Akl, Amany M., Sondoss El Sawah, Ripon K. Chakrabortty, Hasan Hüseyin Turan. 2022. A Joint Optimization of Strategic Workforce Planning and Preventive Maintenance Scheduling: A Simulation–Optimization Approach. *Reliability Engineering and System Safety* **219**(November 2021) 108175.

Hu, Jueming, Yuhao Wang, Yutian Pang, Yongming Liu. 2022. Optimal maintenance scheduling under uncertainties using Linear Programming-enhanced Reinforcement Learning. *Engineering Applications of Artificial Intelligence* **109**(December 2021) 104655.

Huang, Jing, Qing Chang, Jorge Arinez. 2020. Deep reinforcement learning based preventive maintenance policy for serial production lines. *Expert Systems with Applications* **160** 113701.

Kosanoglu, Fuat, Mahir Atmis, Hasan Hüseyin Turan. 2022. A deep reinforcement learning assisted simulated annealing algorithm for a maintenance planning problem. *Annals of Operations Research* .

Mao, Jia Yang, Quan Ke Pan, Zhong Hua Miao, Liang Gao, Shuai Chen. 2022. A hash map-based memetic algorithm for the distributed permutation flowshop scheduling problem with preventive maintenance to minimize total flowtime. *Knowledge-Based Systems* **242** 108413.

Nunes, P., J. Santos, E. Rocha. 2023. Challenges in predictive maintenance – A review. *CIRP Journal of Manufacturing Science and Technology* **40** 53–67.

Panzer, Marcel, Benedict Bender. 2021. Deep reinforcement learning in production systems: a systematic literature review. *International Journal of Production Research* .

Valet, Alexander, Thomas Altenmüller, Bernd Waschneck, Marvin Carl May, Andreas Kuhnle, Gisela Lanza. 2022. Opportunistic maintenance scheduling with deep reinforcement learning. *Journal of Manufacturing Systems* **64**(March) 518–534.

Yan, Qi, Hongfeng Wang, Fang Wu. 2022. Digital twin-enabled dynamic scheduling with preventive maintenance using a double-layer Q-learning algorithm. *Computers and Operations Research* **144**(July 2021) 105823.

Exploring the Frontiers of Generative AI in Classical Production Line Research

Sencer YERALAN
Global Academics Coalition, Gainesville, Florida, USA, yeralan@globalacademicscoalition.net

Cathal HEAVEY
University of Limerick, Limerick, V94 T9PX, Ireland, cathal.heavey@ul.ie

Amir GHASEMI
Institute AIFB, Karlsruhe Institute of Technology, Kaiserstr. 89, Karlsruhe, 76133, Germany, amir.ghasemi@kit.edu

The realm of production line research has recently leveraged machine learning techniques to address its classical problems. However, the integration of Generative Artificial Intelligence in this field remains notably underexplored. Our study aims to fill this gap by examining the applicability and limitations of Generative AI in the context of a conventional tandem queueing model, a staple in production line analysis. We embarked on an experimental investigation where Generative AI was employed to propose empirical functional forms. These AI-generated forms were designed to offer curve-fit-type approximations for the model under study. The core objective was to evaluate how these AI-driven approximations perform in comparison to the exact solutions. This comparison sheds light on the efficacy of Generative AI in capturing the complexities and nuances of classical production models. Our findings offer significant insights into the potential roles and boundaries of Generative AI in this specialized research area. This study not only contributes to the existing body of knowledge by introducing an innovative approach to an established field but also opens new avenues for further research in employing advanced AI techniques in the realm of production line and general queueing theory.

Key words: Generative AI; Classical Production Line Research

1. Introduction

Artificial Intelligence (AI) has been used in the analysis of manufacturing and service operations. For example, machine learning can greatly enhance the efficiency and productivity of production lines (Kang et al. 2020), where it can be used to predict equipment failures and schedule maintenance, thus reducing downtime. Such models are discriminate AI models, rather than Generative AI (GenAI) models. GenAI refers to a subset of artificial intelligence technologies that have the ability to create new content. Any model that produces output can be technically seen as generative, with García-Peñalvo and Vázquez-Ingelmo (2023) defining "'Generative AI' as the production of previously unseen synthetic content, in any form and to support any task, through generative modeling."

GenAI has made great advances, leading in recent years to the development of Large Language Models (LLMs) that can analyse inputs such as text or images (Kumar et al. 2023). These tools, such as ChatGPT (*ChatGPT* Accessed 22 March 2024), are predicted to have a significant effect on businesses and society due to their ability to perform tasks, while others are more skeptical (Rogers 2024). However, even with the current developments with LLMs, how we lecture and work has changed. In this paper, our focus is not on the development of these models but on how they can enhance the efficiency of manufacturing and service systems, the development of analytical formulae or simulation models for the analysis of these systems. In the remainder of this article, ChatGPT 4.0 is used (*ChatGPT* Accessed 22 March 2024).

2. GenAI: Queuing Models of Production Lines

In ChatGPT 4.0 an input is entered, denoted as the prompt (English language used), with new knowledge generated. ChatGPT 4.0 uses the prompt to identify what knowledge to generate. The details of each prompt can cause different output from ChatGPT 4.0, as the outputs are probabilistic. In ChatGPT 4.0 the following prompt is entered:

"*Create a queuing model in Python of a two station production line with blocking.*"

The above prompt generates the following knowledge from ChatGPT 4.0

"*Creating a queuing model of a two-station production line with blocking involves setting up a system where there are interactions between the stations due to limited buffer capacities. This type of model is often referred to in operations research and industrial engineering for analyzing the flow of items through a production process where stations can block each other if one station gets full and cannot accept more items until another station clears some space. Here's a step-by-step setup of such a model:*

1. *Define the Stations*

- *Station 1 (S1): Processes items and passes them to Station 2.*
- *Station 2 (S2): Receives items from Station 1 and further processes or finishes them.*

etc"

This answer is only a partial output from ChatGPT 4.0. In the full reply approximately 45 lines of general information on queuing models is provided for the analysis of production lines.

3. GenAI: DES Models of Production Lines

There are several articles written on the use of ChatGPT to assist in the writing of programming languages. Sun et al. (2023) compare the use of ChatGPT for automatic code summarization against other state-of-the-art code summarization existing models. Bucaioni et al. (2024) evaluating ChatGPT's abilities in two different programming languages, namely C++ and Java. They found that compared to code written by humans, ChatGPT provides lower quality with regards to runtime nad memory usage. Another ariticle is by Kazemitabaar et al. (2023) who examined 69 novice programmers who examined ChatGPT and Codex for Python code tasks and found it improved authoring tasks. Jackson et al. (2024) focused on the use of ChatGPT 3.0 and Codex to develop discrete event simulation models (DES) for logistic models.

Using ChatGPT 4.0 we inserted the following prompt:

"*Write a discrete event simulation model for a two station production line with finite buffers.*"

The output from ChatGPT 4.0 is:

"*To create a discrete event simulation model for a two-station production line with finite buffers, we'll use Python and the simpy library, as it's broadly used and well-suited for such simulations. This model will simulate a production line where parts go through two sequential processing stations with a buffer in between. Here's a step-by-step guide and the necessary Python code:*
etc"

The above answer is only a partial output from ChatGPT 4.0. In addition, it also provided a documented SimPy (*Overview — SimPy 4.1.1 documentation* 2024) model of a two station line of a production line.

4. Observations and Conclusions

GenAI can only provide very high-level documentation and no model for a queueing model of a production line. It can also assist in programming and provide basic DES models for production lines. However, it is necessary for the user to understand key programming and system concepts to successfully implement solutions. GenAI tools such as ChatGPT 4.0 provides basic supports in such tasks. While GenAI needs to be integrated into education, engineering intuition, by grounding the model in its expected use and considering the properties of the system, offers insights currently overlooked by GenAI.

References

Bucaioni, A., Ekedahl, H., Helander, V. and Nguyen, P. T. (2024). Programming with ChatGPT: How far can we go?, *Machine Learning with Applications* **15**: 100526.
URL: *https://www.sciencedirect.com/science/article/pii/S2666827024000021*

ChatGPT (Accessed 22 March 2024).
URL: *https://chat.openai.com*

García-Peñalvo, F. and Vázquez-Ingelmo, A. (2023). What Do We Mean by GenAI? A Systematic Mapping of The Evolution, Trends, and Techniques Involved in Generative AI, *International Journal of Interactive Multimedia and Artificial Intelligence* **8**(4): 7.
URL: *https://www.ijimai.org/journal/sites/default/files/2023-11/ijimai8$_{41}$.pdf*

Jackson, I., Ivanov, D., Dolgui, A. and Namdar, J. (2024). Generative artificial intelligence in supply chain and operations management: a capability-based framework for analysis and implementation, *International Journal of Production Research* pp. 1–26.
URL: *https://www.tandfonline.com/doi/full/10.1080/00207543.2024.2309309*

Kang, Z., Catal, C. and Tekinerdogan, B. (2020). Machine learning applications in production lines: A systematic literature review, *Computers & Industrial Engineering* **149**: 106773.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S036083522030485X*

Kazemitabaar, M., Chow, J., Ma, C. K. T., Ericson, B. J., Weintrop, D. and Grossman, T. (2023). Studying the effect of AI Code Generators on Supporting Novice Learners in Introductory Programming, *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, Association for Computing Machinery, New York, NY, USA, pp. 1–23.
URL: *https://dl.acm.org/doi/10.1145/3544548.3580919*

Kumar, S., Musharaf, D., Musharaf, S. and Sagar, A. K. (2023). A Comprehensive Review of the Latest Advancements in Large Generative AI Models, *in* R. N. Shaw, M. Paprzycki and A. Ghosh (eds), *Advanced Communication and Intelligent Systems*, Springer Nature Switzerland, Cham, pp. 90–103.

Overview — SimPy 4.1.1 documentation (2024).
URL: *https://simpy.readthedocs.io/en/latest/*

Rogers, R. (2024). What's AGI, and Why Are AI Experts Skeptical? Section: tags.
URL: *https://www.wired.com/story/what-is-artificial-general-intelligence-agi-explained/*

Sun, W., Fang, C., You, Y., Miao, Y., Liu, Y., Li, Y., Deng, G., Huang, S., Chen, Y., Zhang, Q., Qian, H., Liu, Y. and Chen, Z. (2023). Automatic Code Summarization via ChatGPT: How Far Are We? arXiv:2305.12865 [cs].
URL: *http://arxiv.org/abs/2305.12865*

How to achieve travel time reliability in a transportation network? Evaluating network structure, service frequency and dispatching strategy

Marion BAUMANN
Institute for Material Handling and Logistics, Karlsruhe Institute of Technology, Karlsruhe, Germany, marion.baumann@kit.edu

Anne LANGE
Frankfurt University of Applied Sciences, Frankfurt, Germany, anne.lange@fb3.fra-uas.de

Kai FURMANS
Institute for Material Handling and Logistics, Karlsruhe Institute of Technology, Karlsruhe, Germany, kai.furmans@kit.edu

Travel time reliability - for passengers and goods - describes the probability, that passengers or goods arrive at the destination latest at a planned time. Hence, reliability is a fundamental precondition to plan efficient transportation processes, avoiding late arrivals on the one hand, but also limiting the necessity to integrate buffer times on the other hand. We develop a mathematical model to capture the most relevant aspects of network design, which affect the reliability of the network. These aspects are the network structure, the transportation frequency, the dispatching strategy and the probability distribution of the transportation times on the legs of the network. We compute the resulting distributions of the travel time and evaluate the reliability of the network from a user's perspective. Additionally, we also derive some generalized insights. Our model is applicable to many planning tasks, where transportation with transfer and regular services are involved, as end-to-end freight transportation networks (like parcel services, groupage, container transport, airfreight), and multi- and intramodal passenger transport using public transportation (rail, air, bus...).

Key words: transportation networks; service quality; service reliability; distribution of travel times; stochastic model

---

1. Introduction

Efficient and reliable transportation networks for passengers and goods are an essential part of prospering economies. Networks allow the consolidation of transportation requirements, thus reducing costs and simultaneously helping to reduce the negative impacts on the environment.

Reliability of travel times is an integral part of transportation network service quality. That is, in an optimal solution, passengers and goods need to reach their destination on-time. In the case of passenger air transportation, for instance, flight delays are one of the most established measurements for the quality of service of air transportation (Prince and Simon 2015). Average delays may indicate the service quality, but the distribution of delays, or travel time, is even more critical for the network's reliability (see for instance Knight (1974) and Rietveld (2005)). The distribution of travel times indicates how precisely the travel time for a single trip in the network can be forecast. A precise expectation of the time spent traveling reduces planning uncertainty for subsequent activities. In contrast, high uncertainty entails inefficient buffers in passenger travel or supply chain plans.

Gaver Jr (1968) for instance has formalized and solved the headstart problem (how much earlier to start traveling than required by the expected travel time) using travel time distributions (in this case travel times by car including congestions) and cost functions for early and late arrival.

However, we observe that typical planning approaches consider a minimization of total travel time or total passenger waiting time, which is equivalent to minimizing the average time. van Oort (2014) identified via an extensive survey of public authorities, that reliability of service is of minor relevance in practice, in cases even neglected, at the stage of network design and timetabling.

We study how three aspects of network-based transportation impact the travel time distribution for the entire trip, namely the dispatching strategy (wait or no-wait), the frequency of recurring trips, as well as the structure of the transportation network.

The decision when to dispatch vehicles is a fundamental one in operating transportation networks. In our analysis, we investigate the role of two fundamental control policies for managing the flow of vehicles in the transportation network: waiting for a delayed vehicle (wait) vs. departing on schedule (no-wait). Ginkel and Schöbel (2007) explain the contrasting objectives of both control policies. In the case of 'wait', no passenger misses her connection, but at the downside that all passengers will arrive late at their destinations. In the case of 'no-wait' all delayed passengers will miss their connections and will have to wait for the next vehicle which establishes a delay for these passengers. However, other passengers arrive on-time at their destinations.

The frequency of trips tightly relates to the dispatching strategy since it will be inefficient to wait for connecting passengers beyond the arrival of the next vehicle to travel along the same line, given there are not capacity constraints. Thus, the trip frequency acts as an upper limit to the waiting time distribution in the network. Rietveld (2005) qualitatively discusses departure frequencies in multimodal chains. For the case of egress from a main station, he argues that low frequencies on the last mile transport worsen increase uncertainty in the total travel time as missing a connection is rather costly in case schedules are not coordinated.

The network structure further impacts the distribution of travel times as, for instance, transfers become more challenging as more vehicles connect to a subsequent vehicle.

2. Modeling Description

We analyze different cases of networks, see Fig. 1. We assume that the vehicles are not a limited resource. We refer to total travel time of a trip in a transportation network as the sum of all individual leg travel times plus intermediate waiting times at transfer stations.

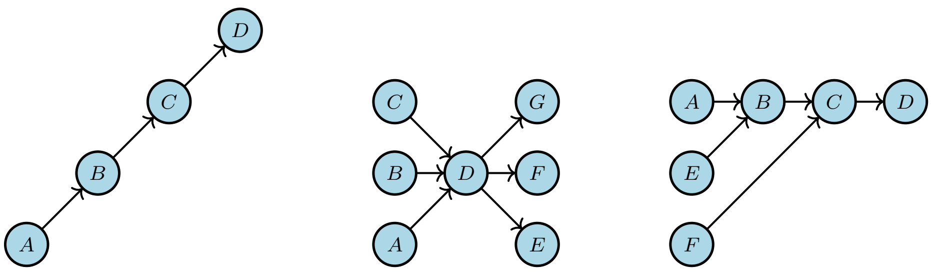

(a) Linear network (b) Hub-and-spoke network (c) Fishbone network

Figure 1. Network structures under consideration

3. Solution Approach

A distinction can be made between the means of transport view and the passenger view. The travel time of the means of transport from a start node to an end node as well as the distribution of the arrival delay per station can be calculated iteratively by cutting off distributions in combination with convolution. Depending on whether the ”no-wait” or ”wait” control policy is applied, additional steps must be performed. For the control policy ”wait”, a maximum function of the travel time distributions is determined from the individual distributions of the means of transport. After calculating the travel time distribution of the means of transport from a start node to an end node as well as the distribution of the delay of arrival per station, the travel time distribution of a passenger can be determined. It is assumed that a passenger changes the means of transport at most once.

4. Numerical Evaluation

We conduct experiments along four essential dimensions, in order to gain insights on the system's behavior as the context changes. The analysis investigates the three common network structures illustrated in Fig. 1. For each of these structures, we run multiple experiments to highlight the effects as the networks become more complex. Demand plays an important role in the model and we model two options. One is uniformly distributed demand between all nodes the network. In contrast, we compare to the case where there is one stream of dominating demand in the model. We compare two delay distributions, a low variance and a high variance case. We compare a policy of 'wait' vs. a policy of 'no-wait' for all cases.

The results show that different network structures have a specific profile in terms of network size and operating frequencies on the threshold for wait / no-wait policy. The expected travel time of no-wait policy is superior to that of wait policy as networks grow. However, the 99% quantile of the travel time with the no-wait policy is significantly higher than the travel time with the wait policy, regardless of the network size and structure. In case the timetable is repeated more frequently, the no-wait policy is attractive even for small networks.

5. Conclusions

Our work contributes strategic insights for reliable vehicle scheduling. We analyze the distribution of travel times as a combination of control policies and network structures. Further research can be conducted on generalising the model and extending the model with additional control policies.

References

Gaver Jr, D. P. 1968. Headstart strategies for combating congestion. *Transportation Science* **2**(2) 172–181.

Ginkel, A., A. Schöbel. 2007. To wait or not to wait? the bicriteria delay management problem in public transportation. *Transportation Science* **41**(4) 527–538. doi:10.1287/trsc.1070.0212.

Knight, T. E. 1974. An approach to the evaluation of changes in travel unreliability: a “safety margin” hypothesis. *Transportation* **3**(4) 393–408.

Prince, J. T., D. H. Simon. 2015. Do incumbents improve service quality in response to entry? evidence from airlines' on-time performance. *Management Science* **61**(2) 372–390.

Rietveld, P. 2005. Six reasons why supply-oriented indicators systematically overestimate service quality in public transport. *Transport Reviews* **25**(3) 319–328.

van Oort, N. 2014. Incorporating service reliability in public transport design and performance requirements: International survey results and recommendations. *Research in Transportation Economics* **48** 92–100.

How to increase emergency responsiveness of a disrupted healthcare network in a large-scale disaster by locating Mobile Care Centers?

Canan PEHLIVAN
Industrial Engineering Center, IMT Mines Albi, Albi, France, canan.pehlivan@mines-albi.fr

Franck FONTALINI
Industrial Engineering Center, IMT Mines Albi, Albi, France, franck.fontalini@mines-albi.fr

This study focuses on strategically pre-positioning mobile care centers, determining their optimal capacity and locations, to reinforce the existing healthcare network and facilitate an efficient and reliable emergency response. The objective is to minimize congestion and maintain control over waiting times within the care network. This includes considering various uncertainties and probabilistic disruptions in infrastructure that may occur during a disaster. As a case study, we center our attention on Istanbul, Turkey, a city bracing for a potentially destructive earthquake in the foreseeable future.

Key words: Mobile Care Centers; Healthcare Network; Emergency; Large-scale disaster; Location Allocation; Stochastic Programming; Data

1. Introduction

This paper concerns the study of a natural disaster (specifically an earthquake) and seeks to devise a methodology for the selection and placement of temporary care centers, such as mobile units or tents, in response to such crises. Large-scale disasters can profoundly affect societal well-being and disrupt humanitarian relief networks. In the aftermath of such disasters, there is a surge in demand for healthcare services, resulting in overloaded hospitals. Moreover, the disaster may damage critical infrastructure, rendering hospitals, roads, bridges, and airports partially or completely nonfunctional. The confluence of increased demand and disrupted healthcare networks inevitably leads to significant congestion and prolonged waiting times in operational and accessible hospitals, substantially diminishing service quality at a time when urgent care is crucial (Pouraliakbari-Mamaghani et al. (2023)).

Our work aims to strategically pre-positioning mobile care centers, determining their optimal capacity and locations, to reinforce the existing healthcare network and facilitate an efficient and reliable emergency response. The objective is to minimize congestion and maintain control over waiting times within the care network by strategically locating temporary care centers in the disaster zone. This necessiates considering various uncertainties and probabilistic disruptions in infrastructure that may occur during a disaster.

2. The Scientific Challenges and Proposed Methodology

2.1. The Study Case and Data Analysis

This project centers its attention on Istanbul, Turkey, a city bracing for a potentially destructive earthquake in the foreseeable future. Istanbul's dense population of 16 million residents raises significant concerns about its readiness for such seismic events. Acknowledging this vulnerability, a comprehensive disaster prevention and mitigation plan was collaboratively devised in 2002 by the Istanbul Metropolitan Municipality (IMM) and the Japan International Cooperation Agency JICA (2002). A subsequent study (IMM and DEZIM (2019)) delves into an extensive analysis of earthquake damage and loss estimates for urban superstructures and infrastructure across Istanbul.

Current projections from the IMM suggest that a seismic event with a magnitude of 7.5 could potentially result in the destruction of approximately 90,000 buildings in the city, with an additional 260,000 buildings facing significant damage. Various studies in the literature, including those Parsons et al. (2000), Acar and Kaya (2019), further contribute to discussions on the likelihood and effects of a potential earthquake in Istanbul.

The initial phase of this project involves conducting a comprehensive analysis of existing data, sourced from the aforementioned references, to formulate input scenarios encompassing healthcare demand and structural damage. Istanbul comprises districts with varying degrees of vulnerability. We possess data outlining the quantity, type, and age of buildings, population statistics, and risk of damage for each district. Additionally, we have extensive data and analyses on potential earthquake scenarios and their impact on the city. The challenge lies in effectively harnessing all this information to model healthcare demand. Hence, the primary objective of the data analysis phase is to predict the behavior of the demandx for diverse healthcare services across different districts of Istanbul following a seismic event. This will depend on several variables, including earthquake specifics, district vulnerability, and structural integrity. Given the inherent uncertainties associated with earthquakes, including magnitude, epicenter, time, and their impact on vulnerable areas and city infrastructure, careful data analysis becomes imperative.

2.2. Positioning mobile care centers

The idea of locating mobile hospitals was first suggested by Blackwell and Bosse (2007). Recently, Acar and Kaya (2019) considered the location and re-location decisions of emergency medical centers (EMC) with various earthquake scenarios. Their model is a combination of pre-disaster and post-disaster location models. They introduced a network design problem to determine the number and capacities of EMCs that need to be acquired before the disaster, the re-location, and allocation decisions in post-disaster situations considering the travel times, expected waiting times, and possible facility damages in the post-disaster environment. The authors proposed a two-stage stochastic programming model for this purpose. Several scientific challenges motivate us to further address a similar problem of locating mobile hospitals in a large-scale disaster.

1. Timely response is crucial in saving lives and impacts the survival rates in a disaster situation. In conventional disaster logistics literature, emphasis is typically placed on travel times between demand sites and selected facility locations, with waiting times often overlooked. This oversight stems from the fact that disaster logistics literature traditionally prioritizes the location of the relief distribution centers, where waiting at the facility is not a primary concern compared to travel times. However, Acar and Kaya (2019) innovatively incorporated waiting times into their model by integrating each Emergency Medical Center as an M/M/1 queuing model, effectively addressing the waiting time factor within their mathematical framework to achieve targeted expected waiting times.

We believe that in such strategic level decisions, it is important to ensure that each emergency facility has enough capacity to accommodate its demand. Thus, it is reasonable to adopt a loss system approach when determining the location and capacity of these emergency facilities, as commonly assumed in the literature when dealing with urgent patients requiring immediate medical attention (Pehlivan et al. (2014)). To address this, we propose to model each mobile care center as a loss queuing system M/G/c/c from which the analytical relationship between service capacity and patient acceptance rate can be obtained by Erlang Loss formula. By working with rejection probabilities (Erlang loss probability), we aim to strategically locate and plan the capacity of emergency medical facilities to ensure a specified maximum rejection probability across the network, thereby optimizing disaster response efforts.

2. Disrupted network: We consider probabilistic damages to the critical urban structures (including the existing hospital network, the roads, bridges, etc.) by assigning a survival probability according to the risk level of the district under an expected disaster scenario. Our goal is to develop a methodology to assess the vulnerability of the existing network and design a robust and resilient network by strategically locating mobile care centers effectively.

3. Robustness: In addressing the challenge of uncertainty, two-stage stochastic programming is a widely adopted approach in humanitarian relief logistics literature. However, two-stage stochastic optimization primarily focuses on expectations of uncertainties, often overlooking variabilities. To address this limitation, we intend to test the robustness of our results using a digital twin model. This will allow us to assess the resilience and reliability of our strategies under various scenarios, providing valuable insights into the effectiveness of our disaster planning.

References

Acar, Müge, Onur Kaya. 2019. A healthcare network design model with mobile hospitals for disaster preparedness: A case study for istanbul earthquake. *Transportation Research Part E: Logistics and Transportation Review* **130** 273–292.

Blackwell, Thomas, Michael Bosse. 2007. Use of an innovative design mobile hospital in the medical response to hurricane katrina. *Annals of emergency medicine* **49**(5) 580–588.

IMM, DEZIM. 2019. Olasi deprem kayip tahminleri raporu. `https://depremzemin.ibb.istanbul/`. İstanbul Büyükşehir Belediyesi (IMM) ve Deprem Zemin Araştırmaları Merkezi (DEZIM) tarafından yayınlanan bir rapor.

JICA, IMM. 2002. The study on a disaster prevention/mitigation basic plan in istanbul including microzonation in the republic of turkey. Tech. rep., Technical report. Japanese International Cooperation Agency.

Parsons, Tom, Shinji Toda, Ross S Stein, Aykut Barka, James H Dieterich. 2000. Heightened odds of large earthquakes near istanbul: an interaction-based probability calculation. *Science* **288**(5466) 661–665.

Pehlivan, Canan, Vincent Augusto, Xiaolan Xie. 2014. Dynamic capacity planning and location of hierarchical service networks under service level constraints. *IEEE Transactions on Automation Science and Engineering* **11**(3) 863–880.

Pouraliakbari-Mamaghani, Mahsa, Ahmed Saif, Noreen Kamal. 2023. Reliable design of a congested disaster relief network: A two-stage stochastic-robust optimization approach. *Socio-Economic Planning Sciences* **86** 101498.

Impact of retrial distribution in stationary and time-dependent queues

Ömer Özümerzifon
Chair of Production Management, University of Mannheim, Mannheim, Germany, oezuemerzifon@uni-mannheim.de

Benjamin Legros
EM Normandie Business School, Métis Lab, Clichy, France, benjamin.legros@centraliens.net

Raik Stolletz
Chair of Production Management, University of Mannheim, Mannheim, Germany, stolletz@uni-mannheim.de

Queueing systems are used in various service systems, such as repair facilities, health care, and call centers. In many of these service systems, customers leave the queue before being served due to a lack of patience. However, these customers may re-enter the system after some time as retrials. This research project focuses on the stationary and time-dependent performance evaluation of multi-server queueing systems with retrials, and the impact of distribution of retrial time, i.e., the time after a customer returns to the system.

We analyze empirical data on the retrial behaviour of customers from several large call center. Novel stationary approaches to approximate main performance measures are developed. A comparison against simulation demonstrates the accuracy of the approaches. Moreover, analytical results and numerical insights demonstrate the impact of the retrial time distribution based on the empirical analysis.

Key words: retrials; time-dependent queues; call centers

Integration of Machine Learning and Optimization Models for a Data-Driven Newsvendor Problem with Random Yield

Bijan BIBAK
Department of Industrial Engineering, Koç University, Sariyer, Istanbul, Turkey
bbibak20@ku.edu.tr
Fikri KARAESMEN
Department of Industrial Engineering, Koç University], Sariyer, Istanbul, Turkey
fkaraesmen@ku.edu.tr

We investigate a data-driven lot-sizing problem within the random yield context, inspired by semiconductor manufacturing. In this setting, the yield rate of a production process is influenced by numerous observable features prior to making lot-sizing decisions. Additionally, demand variability is considered, which may also depend on various features. Addressing the data-driven lot-sizing problem is complex due to the reliance on a vast array of features with limited data. To navigate these complexities, we introduce methods that blend machine learning techniques with stochastic optimization. Leveraging both a publicly accessible semiconductor yield data set and a crafted synthetic data set, we assess various estimation and optimization strategies. Our findings reveal the substantial benefits of incorporating feature information for reducing costs. Moreover, the best approach for tackling the problem merges machine-learning based estimation techniques with the theoretical optimization principles specific to random yield inventory problems.

Key words: Inventory Control: Lot-sizing; Yield Uncertainty; Data-Driven Optimization; Machine Learning

1. Introduction

In many sectors, firms encounter substantial planning difficulties due to yield uncertainty, where the actual quantity produced or replenished may deviate from what was initially planned. This uncertainty in yield adds a layer of complexity to production and replenishment processes, making effective planning challenging. Historically, the research community in operations research has dedicated efforts to understanding and addressing these challenges of production and inventory planning under uncertain yields with significant contributions to the field. A survey by Yano and Lee (1995) has meticulously categorizes the existing literature, highlighting key discoveries and methodologies developed to tackle these issues.

The motivation for our research stems from a specific challenge faced in the semiconductor industry, known for its struggles with yield variability during wafer production. Despite advancements in technology and design, yield unpredictability remains a critical concern. Data from leading manufacturers indicate particularly low yield rates for some products during the initial stages of production, emphasizing the need for planning models that accommodate yield uncertainty amidst increasing product variety and decreasing lead times. With the advent of machine learning, predictive methods for yield have shown promise, suggesting that features related to design and production environment could significantly enhance yield rate forecasts. Our study explores a data-driven approach for estimating and optimizing planning problems under yield uncertainty, using both real and synthetic data sets to test the efficacy of these methods. By addressing the complexities of lot-sizing problems affected by both yield and demand uncertainties, our approach demonstrates the practical value of integrating data-based decision-making strategies from a cost reduction standpoint.

2. The Problem and Models

We consider the following problem: a manufacturer has a random demand D to satisfy by some deadline. The manufacturing process is subject to random yield and whenever Q items are released, a random proportion turns out to be of acceptable quality. Let us denote by Y the random yield rate (the proportion of acceptable items). We assume that whenever Q items are released, QY is good for delivery. This is the well-known multiplicative (or proportional) model of yield uncertainty (Henig and Gerchak 1990). We assume that Y takes values between 0 and 1 and consider the following standard formulation each unit short of D incurs a penalty cost of b, and each unit above the demand has a cost of h. The objective is choose the lot size Q that minimizes the expected cost. This results in the following formulation:

$$\min_{Q} z(Q) = E\left[\left(b(D-QY)^{+} + h(QY-D)^{+}\right)\right] \tag{1}$$

The above is a well-studied problem in the random yield literature. The solution to the problem (1) with full distributional information on D and Y dates back to Shih (1980), Gerchak et al. (1988) and Henig and Gerchak (1990).

2.1. A Model with Yield and Demand Features

Our main focus in the paper is on the case where there are additional feature data that can possibly explain part of the variability in the random yield and demand. The feature information is available before determining the lot size Q. Let us assume that $\mathbf{X}$ is a vector consisting of r features that can potentially provide information on D and $\mathbf{U}$ is a vector consisting of s features that can provide information on Y. We can now choose the lot size conditional on the observed features, leading to:

$$\min z(Q) = E\left[\left(b(D|\mathbf{X}=\mathbf{x}\ -\ QY|\mathbf{U}=\mathbf{u})^{+} + h(QY|\mathbf{U}=\mathbf{u}\ -\ D|\mathbf{X}=\mathbf{x})^{+}\right)\right] \tag{2}$$

The above formulation is a direct extension of (1) and shows that if the conditional distributions are available, optimization proceeds as in the case without features. However, in many real applications, the number of features may be very large and the observed sample that is available may be relatively small and conditional distributions cannot be estimated easily. Next, we turn our attention to the data-driven case with a large number of features to develop viable solutions to (2).

3. Data-Driven Case with Many Features

Now, let us assume that we have a sample of data that includes the yield and demand realizations (d_i and y_i) along with corresponding set of feature values $\mathbf{X}_i$ and $\mathbf{U}_i$ for each observation. To provide an example of the potential challenges of data-driven optimization, a publicly available semiconductor data set has more than 500 features for each observation for a relatively small sample of 1567 observations. This underlines the importance of model reduction and robust estimations for conditional uncertainty from a large number of features.

Some recent research considers a number of promising approaches for predictive analytics that combine estimation capabilities of machine learning methods with cost-based stochastic optimization (see for example Ban and Rudin (2019), Bertsimas and Kallus (2020))). To address the challenge of solving (2) with a large number of features, we employ and test a number of the proposed approaches along with some new problem-specific approaches. At the end of a comprehensive numerical study, We reach two main conclusions: i) feature information when used smartly leads to a significant decrease in the costs ii) the best performing method for the random yield problem combines the analytical solution of (2) with machine learning methods for esimating the conditional distribution.

4. Conclusions

We develop and systematically test a number data-driven strategies for calculating optimal lot sizes in manufacturing, taking into account uncertainties in yield and demand. Machine learning-based predictions are integrated with an expected cost minimization approach to derive rules for determining optimal lot sizes within a training dataset to learn the best lot-sizing rule for a given set of features.

A comprehensive version of the results can be found Bibak and Karaesmen (2024)

Acknowledgments
This research was supported by the Scientific and Technological Research Council of Türkiye through a TUBITAK 1001 Research Grant (project no: 123M871)

References

Ban, G., C. Rudin. 2019. The big data newsvendor: Practical insights from machine learning. *Operations Research* **67**(1) 90–108.

Bertsimas, D., N. Kallus. 2020. From predictive to prescriptive analytics. *Management Science* **66**(3) 1025–1044.

Bibak, B., F. Karaesmen. 2024. Integration of machine learning and optimization models for a data-driven lot sizing problem with random yield. *Available at SSRN 4760990* .

Gerchak, Y., R.G. Vickson, M. Parlar. 1988. Periodic review production models with variable yield and uncertain demand. *IIE Transactions* **20** 144.

Henig, I.M., Y. Gerchak. 1990. The structure of periodic review policies in the presence of random yield. *Operations Research* **38** 634–643.

Shih, W. 1980. Optimal inventory policies when stockouts result from defective products. *International Journal of Production Research* **18** 677–686.

Yano, C.A., H.L. Lee. 1995. Lot sizing with random yields: A review. *Operations Research* **43** 311–334.

Machine Learning Based Approximations for Some Production Control Problems with Product Returns and Remanufacturing

Arash GANDOMKAR GHALHAR, Bijan BIBAK, Fikri KARAESMEN,
Dept. of Industrial Engineering, Koç University, Sariyer, Istanbul, Turkey aghalhar22@ku.edu.tr , bbibak20@ku.edu.tr, fkaraesmen@ku.edu.tr,

Oktay KARABAG
Dept. of Industrial Engineering, Izmir University of Economics, Balçova, Izmir, Turkey oktay.karabag@ieu.edu.tr

Barış TAN
Dept. of Industrial Engineering, Özyegin University, Çekmeköy, Istanbul, Turkey baris.tan@ozyegin.edu.tr

Flow control for manufacturing systems with return loops of recycled products that can be refurbished or remanufactured is an important current problem. Under strong assumptions one can build Markov Decision Process (MDP) models for analyzing such systems. These models can be solved numerically at the expense of a significant computational effort if the state space is not too large. In order to develop fast and accurate methods that can approximately compute the optimal control policy and its parameters, and the corresponding cost performance, we employ a machine-learning based approach. To this end, we first build a labeled data set using numerical solutions of Markov Decision Process models of such systems for different combinations of input parameters and then test different learning algorithms that use the inputs as features, obtaining the hyper-parameters by cross-validation. The resulting approximations are fast enough to integrate into simulation-based on-line learning algorithms.

Key words: Manufacturing Flow Control; Remanufacturing; Markov Decision Processes

1. Introduction

Circular Economy encompasses supply chains that incorporate product returns, refurbishing and remanufacturing decisions as possible options or sometimes requirements (Dekker et al. (2013)). Making these decisions efficiently has significant economic and environmental consequences. A recent European research project AUTO-TWIN (https://www.auto-twin-project.eu/) proposes to investigate and aid such green gateway decisions by the aid of a digital twin of a manufacturing system.

Motivated by the potential of having a digital twin that can simulate an existing system, we explore methods for dynamic decision making for flows including returns and remanufacturing decisions. In particular, we investigate ideas for rapid and practical analysis for dynamic optimization that can be later integrated with simulation optimization.

A number of papers investigate basic models that include return flows and dynamic remanufacturing/coordination decisions. For instance, Vercraene and Gayon (2013) consider product returns and Vercraene, Gayon and Flapper (2014) investigate coordination of remanufacturing and returns acceptance.

The above papers and many others use Markov Decision Process (MDP) models of these systems. The MDP framework is formal and enables an understanding of the system but becomes impractical when the state space of the system grows or when Markovian assumptions may not hold. In this paper, we present some preliminary results that may eventually help obtain approximations for models with larger state spaces employing machine learning methods to enable learning of rules that estimate performance measures and policy parameters from inputs without the need for a numerical solution.

2. Models

2.1. Benchmark: A Capacitated Production/Inventory System without Returns

We consider a single-stage make-to-stock manufacturing system with a single processor (machine). We assume that demand arrives randomly according to a stationary Poisson process with rate λ, the processing times of the machine are exponentially distributed with rate μ. We assume that $\lambda < \mu$ for stability. When a product completes processing, it is placed in a finished goods inventory. When demand arrives, it is immediately satisfied from inventory whenever inventory is available. If no inventory is available, the arriving demand is backordered and waits in the backorder queue until a product becomes available. We assume that physical inventory incurs holding costs of h per item per time and backorders incur a cost of b per backorded item per unit time. This is a standard setup studied for instance by Buzacott and Shanthikumar (1993).

The typical dynamic production control problem is to decide when to continue production and when to stop depending on the inventory level. This well-established benchmark case has many advantages for testing purposes. The optimal policy is known to be a threshold type policy (a base stock policy). In addition, both the optimal threshold and the optimal cost have closed form expressions (Buzacott and Shanthikumar (1993)). In addition, the relative value function $w(x)$ can be numerically obtained by value iteration.

Machine Learning Approach for Quick Estimation

Since the important performance measures and policy parameters can be obtained expicitly, we can generate a labeled training data set where the input features are the parameters λ, μ, h and b and the targets are the corresponding optimal cost and the optimal threshold. One important remark is that both the optimal threshold and the optimal cost are highly non-linear functions of the inputs. We, therefore, test some non-linear estimation tools from machine learning to learn rules from the training data that estimate the optimal base stock level and the corresponding cost.

For the numerical setup, we generated the main dataset by systemically changing the parameter values (λ, μ, h, and b). The dataset was subsequently divided into training (80%) and testing (20%) segments. Utilizing 5-fold cross-validation, we tuned the optimal configuration for the Random Forest algorithm. The robustness of the model was tested across datasets with varying features and sizes, as summarized in the Table 1 below. In the table, we experiment with different training-test sizes and report the Root Mean Squared Error (RMSE) and the Mean Absolute Percentage Error (MAPE) for the cost estimation. We also report the threshold estimations that are within $\pm$ 1 unit of the true optimal threshold. Results indicate that, for this model the estimations perform well. More importantly, the model still produces satisfactory results even when only 50% of the data is used for training, and the rest for testing.

	(Train,Test) Size	(Train, Test) RMSE (g)	(Train,Test) MAPE (g)	$\lvert\hat{S} - S^*\rvert \leq 1\ (\%)$
Instance 1	(576,144)	(0.21,0.43)	(1.82,5.41)	(99.31,97.92)
Instance 2	(164,41)	(0.41,1.17)	(1.33,3.93)	(97.56,87.80)
Instance 3	(504,126)	(0.23,0.65)	(1.70,3.81)	(99.40,95.24)
Instance 4	(360,360)	(0.27,1.20)	(3.26,7.38)	(98.89,95.00)
Instance 5	(102,103)	(0.59,1,20)	(2.46,7.21)	(96.08,83.50)
Instance 6	(315,315)	(0.35,1.02)	(2.95,7.29)	(98,73,94.60)

Table 1. Performance of the Benchmark Model Across Diverse Datasets with Different Sizes.

2.2. A Capacitated Production/Inventory System with Returns

As the next test case, we consider a system with returns for which the performance measures cannot be computed analytically. Let us assume that the used items are returned at rate δ according to a Poisson process. These items can be accepted and placed in an input inventory for refurbishing or rejected at a cost. We assume that refurbishing requires the same processing as producing a new item and a refurbished item is as good as a new product.

Figure 1 presents the model without returns on the left and the one with the returns on the right.

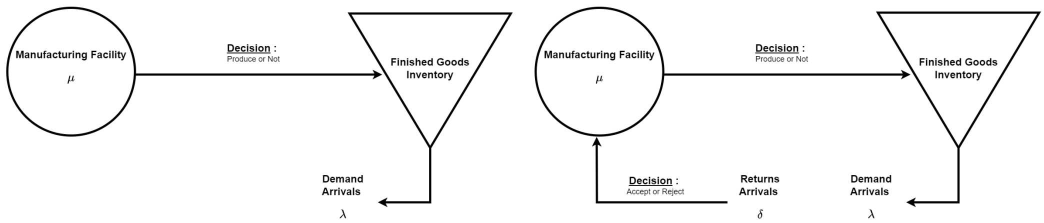

Figure 1. Models without (left) and with returns (right)

3. Conclusions

We propose and test a supervised learning approach to estimate performance measures of a production/inventory system. If such approximations can be developed robustly, they can be integrated within simulation optimization for dynamic decision making. The initial results look promising.

Acknowledgments

European Union's Horizon 2020 Research and Innovation Programme [AutoTwin project under grant agreement No. 101092021].

References

Buzacott, J. A., J. G. Shanthikumar. 1993. *Stochastic models of manufacturing systems*. Prentice-Hall.

Dekker, R., M. Fleischmann, K. Inderfurth, L. N. van Wassenhove. 2013. *Reverse logistics: quantitative models for closed-loop supply chains*. Springer Science & Business Media.

Vercraene, S., J.-P. Gayon. 2013. Optimal control of a production-inventory system with product returns. *International Journal of Production Economics* **142**(2) 302–310.

Vercraene, S., J.-P. Gayon, S. D. Flapper. 2014. Coordination of manufacturing, remanufacturing and returns acceptance in hybrid manufacturing/remanufacturing systems. *International Journal of Production Economics* **148** 62–70.

Maximizing the throughput capacity of mixed load multi-deep storage and retrieval systems

Timo LEHMANN
Institute for Material Handling and Logistics, Karlsruhe Institute of Technology, Karlsruhe, Germany, timo.lehmann@kit.edu

René De KOSTER
Rotterdam School of Management, Erasmus University, Rotterdam, Netherlands, rkoster@rsm.nl

Multi-deep storage systems are space efficient storage solutions for e-commerce and spare part logistics. Popular versions are robotic compact storage and retrieval (RCS/R) and multi-deep automated storage and retrieval (AS/R) systems. In such systems, multiple loads can be stored behind or above each other in a single lane, which leads to high space utilization. However, blocking loads must be reshuffled if they block a requested load. This increases the command cycle time. We present models to estimate the throughput capacity of these systems, which can be used in the design phase. We build these models using four different storage assignment strategies, three different load reshuffling strategies and two retrieval load selection strategies, incorporating the access frequency of products and allowing multiple stored loads per product for some strategies. Seven strategy combinations are analyzed which include among others the current AutoStore strategy and two state of the art strategies. The throughput capacity is determined using closed queuing networks and the model's quality is validated with simulation. We find that the class-based storage assignment strategy, where different classes share the same lanes, yields the highest throughput capacity and that relocation of reshuffled loads to other storage channels is superior to temporarily buffering and bringing back these loads. Furthermore, when information about the access frequency and number of loads per product is available, the throughput capacity can be increased significantly by properly storing and reshuffling loads to better positions. Based on the throughput model, we optimize the rack layout yielding maximum throughput capacity and provide the throughput capacity for a given industry rack layout. Furthermore, we provide managerial insights on storage assignment, reshuffle and retrieval load selection strategies for multi-deep storage systems.

Key words: Deep-lane storage; Multi-deep storage; Throughput capacity; RCSRS; AS/RS; travel time model; performance estimation

---

1. Introduction

In recent decades, the introduction of various automated compact warehouse storage systems has transformed the logistics landscape. These storage systems employ multi-deep storage lanes to efficiently store products. Examples of these systems are multi-deep automated storage and retrieval (AS/R), robotic compact storage and retrieval (RCS/R), and multi-deep autonomous vehicle storage and retrieval (AVS/R) systems. A comprehensive overview of such systems is provided by Azadeh et al. (2019). Multi-deep (or deep-lane) storage systems are characterized by high space utilization and, for some systems, high throughput capacity. Consequently, they find applications in many distribution centers and spare part warehouses. In these systems, stored loads (totes, bins, pallets, containers) are arranged in multi-deep storage, either behind one another in a storage lane (as in AS/R and AVS/R systems), or they are stacked vertically (as in RCS/R systems or sea-container stacks). Load handling devices (cranes, robots, shuttles) cannot directly access each load. When a blocked load needs retrieval, the system must reshuffle the blocking storage loads, thereby adding to the command cycle time and decreasing system throughput capacity, a critical performance indicator.

The operator's challenge lies in efficiently storing and retrieving loads while minimizing the number of reshuffles. When a new system must be designed, it is important to estimate up front the impact of the design and operating strategies on the throughput capacity. Operating strategies include the storage assignment, reshuffle, and retrieval load selection strategy for a requested product, as well as channel (i.e. storage lane) policies such as shared storage of loads of different products in a channel or dedicated storage of only loads of the same product. When information on access frequency of different products and the number of loads per product in the system are known, more sophisticated storage assignment policies, like class-based storage or smart retrieval selection strategies can be used. However, storage planners have to lean on simulations and empirical data to devise new multi-deep storage systems, because analytic models for multi-deep storage systems employing these strategies are hardly available in literature. Additionally, comparisons between strategies have hardly been explored.

Therefore, in this paper we focus on RCS/R and multi-deep AS/R systems and address the following research question: How do the different storage assignment, reshuffle, and retrieval load selection strategies perform and which strategy combination maximizes the long-run throughput capacity? We also investigate the strategy used by the market leader for RCS/R systems (The AutoStore (2023) strategy is shown by Meller (2023)) and compare it to other strategies and investigate how information on the product access frequency and the stored number of loads per product can be used to improve the throughput capacity. We use this knowledge to find the optimal rack layout of a multi-deep RCS/R or AS/R system of a given storage capacity that maximizes the throughput capacity.

Our throughput models are applicable to two different multi-deep storage systems: RCS/R and multi-deep AS/R systems. We primarily focus on RCS/R systems. In the appendix, we illustrate how our results can also be applied to multi-deep AS/R systems.

2. Modeling Approach

To estimate the throughput capacity of a given RCS/R system, we use a closed queuing network (CQN) with two server nodes (see Figure 1). In a CQN, retrieval jobs are always available and robots

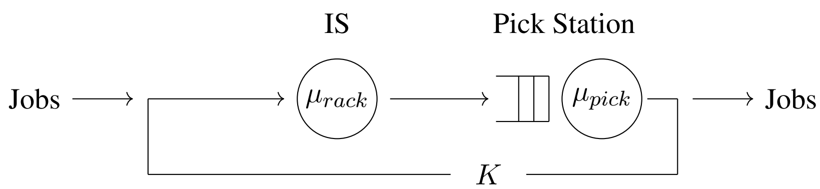

Figure 1. Closed Queuing Network with K Robots and two Server Nodes.

do not have to wait for new command cycles. The time needed for storage, reshuffles and retrieval of a command cycle is assumed to be an exponentially distributed delay process (an infinite server, denoted by IS). The service time of the pick station is assumed to be exponentially distributed and the queuing discipline is FCFS. Robots may have to queue at the pick station while the operator is still handling a previous load and there are K robots in the network.

3. Managerial Insights

The assessment of seven distinct operational strategies yields valuable insights for storage planners. Our findings demonstrate that the integration of supplementary data, including the number of loads per product and the frequency of access for products across different classes, significantly enhances the throughput efficiency of a given storage system. Notably, we conduct a comprehensive comparison and analysis of the prevailing AutoStore strategy against both random selection and a novel class-based approach. Our results reveal that while the AutoStore strategy performs comparably to random selection, the class-based approach demonstrates promising improvements. Moreover, we ascertain upper and lower bounds for the length-to-width ratio in RCS/R systems, along with identifying the optimal positioning for pick stations—preferably located at the midpoint of the longer side within an RCS/R system.

4. Conclusions

Our study introduces diverse models for calculating throughput capacity in multi-deep AS/R and RCS/R systems, offering practical utility during the design phase of new storage facilities. By comparing the AutoStore strategy with a newly proposed class-based alternative, we advocate for enhancements to the AutoStore approach based on our comparative analyses. Furthermore, our research lays the groundwork for developing throughput models applicable to other storage systems, such as the PowerCube (Jungheinrich 2023) or AVS/R systems, fostering broader insights into optimal storage system design and operational efficiency.

References

AutoStore. 2023. Autostore - company website. URL `https://www.autostoresystem.com/company`. Accessed on: 17.12.2023.

Azadeh, Kaveh, René De Koster, Debjit Roy. 2019. Robotized and automated warehouse systems: Review and recent developments. *Transportation Science* **53**(4) 917–945.

Jungheinrich. 2023. Jungheinrich - powerstore powercube. URL `https://www.jungheinrich.de/systeme/automatische-lagersysteme/powercube`. Accessed on: 17.12.2023.

Meller, Russell D. 2023. Considerations when designing an autostoretm system. *Progress in Material Handling Research* **16th proceedings**.

Modeling and analyzing the effect of stochastic production on product rollovers

Niklas Owens
University of Regensburg, Regensburg, Germany, niklas.owens@ur.de

Justus Arne Schwarz
University of Regensburg, Regensburg, Germany, justus.schwarz@ur.de

Barış Tan
Özyeğin University, Istanbul, Turkey, baris.tan@ozyegin.edu.tr

Oktay Karabağ
Izmir University of Economics, Izmir, Turkey oktay.karabag@ieu.edu.tr

In many industries, frequently introducing new generations of already existing products is a common tool for manufacturers to maintain their market share. Before the market introduction of a new product generation, decisions on production and sales quantities as well as prices must be made. The resulting product rollover strategies can be divided into single and dual rollover strategies. Phasing the old product generation out of the market before the introduction of the new generation is called a single product rollover strategy while selling both generations simultaneously is known as a dual rollover strategy. Stochastic production capacity consumption regarding the new generation is common due to the lack of knowledge with its production process. We analyze the effect of stochastic capacity consumption required by the new product generation due to new production processes associated with the new product generation on optimal sales and production decisions and the resulting rollover strategies. This talk proposes a stylized multi-period model for determining optimal product rollover decisions under stochastic production processes. Preliminary numerical results about structural properties of optimal sales and production decisions during product rollovers are presented.

Key words: Product Introduction; Product Rollover; Stochastic Capacity Consumption

1. Introduction

A product rollover refers to the process of transitioning from one version of a product to the next generation with updated or new features. This typically involves either a direct retiring of the older version, i.e., a single rollover strategy (SRS), or temporary sales of both generations, i.e., a dual rollover strategy (DRS) (Billington et al. 1998). The goal of a product rollover is to maintain customer satisfaction and market shares by offering consumers the latest innovations and advancements. The observed frequency of rollovers depends on the industry, e.g., in the home appliance industry new washing machines are introduced every two to three years.

When executing a product rollover, firms need to decide on the amounts to be produced of both product generations, the product offering, i.e., the rollover strategy, and the prices to be charged for the products to maximize their profits. These decisions need to be made under consideration of stochastic customer behavior with cannibalization if both products are offered simultaneously as well as under consideration of stochastic production processes.

There is a large body of literature that analyzes product rollovers with unlimited production capacity under different customer demand models, e.g., Lim and Tang (2006), Liang et al. (2014), Çelik et al. (2024). Limited but deterministic production capacities are considered by Schwarz and Tan (2021). We propose a new model to analyze the impact of stochastic capacity consumption required by the new product generation due to new production processes associated with the new product generation.

2. Model

The model considers a monopolist with finite production capacity that offers its old product generation on the market and plans to introduce a new generation to the market. The dynamics of the transition are captured in three periods (see Figure 1). The manufacturer's objective is to maximize the expected profit from the revenue during the pre-introduction, the introduction period, and the post-introduction period subtracted by the production and holding costs by deciding on production quantities, offered products, and prices.

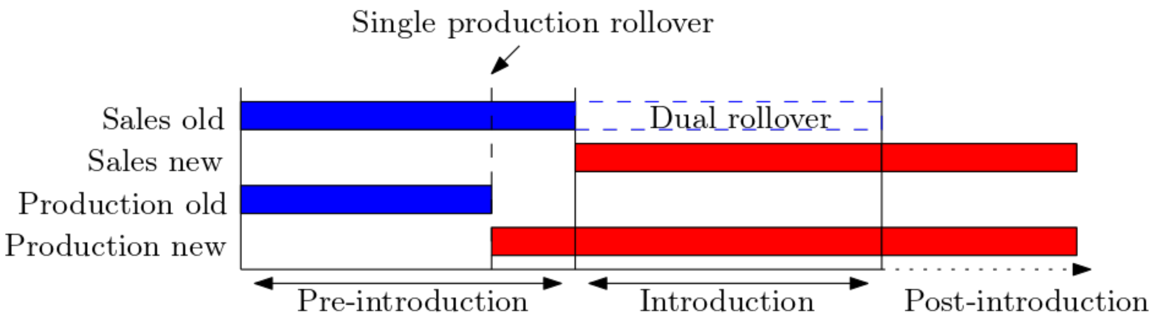

Figure 1. Product rollover timeline.

During the pre-introduction period, only the old product generation is offered to the customers. Old products produced during the pre-introduction period can be used to fulfill a stochastic and price-dependent demand during the pre-introduction period. It is assumed that the new generation is produced on the same production line but only one generation can be produced at a time. Consequently, if the manufacturer also wants to produce new products during the pre-introduction period, a single production rollover has to be executed. That means the manufacturer has to decide when to switch from producing the old generation to producing the new generation. The capacity consumption is known for the old product generation. However, the potentially increased capacity consumption of the new generation is stochastic due to new production technologies and processes. After making the production decisions and observing the realized capacity consumption the corresponding production output is received. The manufacturer decides on how many of the old product units to offer on the market during the pre-introduction period and the respective price. Old products that were not offered or not sold due to the stochastic customer behavior, as well as any new products can be carried over to the introduction period under the consideration of inventory holding cost. Any units not carried over are salvaged.

At the beginning of the introduction period, the manufacturer decides on how many units of the new product generation will be produced during the introduction period while considering the inventory levels of both product generations. After observing the realization of the still stochastic capacity consumption of the new generation, the manufacturer also decides on the rollover strategy and the prices. After the realization of the stochastic demand process, the manufacturer observes the remaining units in stock.

It is assumed that in the post-introduction period, there is still positive demand for the new generation but the old generation becomes obsolete. Hence, there is no demand for the old generation. During the post-introduction the manufacturer has the production processes under control so that there is no longer an additional stochastic capacity consumption. Because the production capacity is typically chosen to match the long-term demand, it is sufficient to satisfy the demand for the new product. Thus, leftover units of the new generation from the introduction period may be used for sales in the post-introduction period, saving the manufacturer production costs but incurring holding costs. Leftovers of the old generation can only be salvaged at the end of the introduction period.

3. Solution approach and numerical findings

Solutions for the proposed model are obtained numerically by backwards induction. Optimal production, sales, and carry over decisions on each stage of the dynamic program are determined via complete enumeration. Prices are discretized in a reasonable price range that is determined by the cost structure and the demand model.

The numerical study shows that the optimal production and pricing decisions feature structural properties in the available stock levels as well as the model parameters. Moreover, the numerical study compares optimal decisions under deterministic and stochastic capacity consumption. The results indicate that the manufacturer can choose from multiple mitigation actions at hand to cope with the stochastic production process. These actions range from minor changes in pricing to a change in the rollover strategy. In particular, using a DRS to leverage the predictable production of old products is observed.

4. Conclusions

We propose a multi-period model to obtain optimal production and sales decisions for product rollovers under stochastic production capacity consumption. First numerical studies show that optimal decisions under stochastic production processes differ from those made under deterministic conditions. Future research will focus on formally establishing numerically observed structural properties and the development of more efficient solution approaches.

References

Billington, C., H. L. Lee, C. S. Tang. 1998. Successful strategies for product rollovers. *MIT Sloan Management Review* **39** 23–30.

Liang, C., M. Çakanyildirim, S. P. Sethi. 2014. Analysis of product rollover strategies in the presence of strategic customers. *Management Science* **60** 1033–1056.

Lim, W. Shi, C. S. Tang. 2006. Optimal product rollover strategies. *European Journal of Operational Research* **174** 905–922.

Schwarz, J. A., B. Tan. 2021. Optimal sales and production rollover strategies under capacity constraints. *European Journal of Operational Research* **294** 507–524.

Çelik, B., J. A. Schwarz, B. Tan. 2024. New product introductions with selection of unique and common features in monopoly markets. *Omega* **125**.

Modelling and Analysis of A Remanufacturing System with Different Quality Returned Material and Finished Goods

Oktay KARABAĞ
Department of Industrial Engineering, İzmir University of Economics, Sakarya Caddesi No:156, Balçova, İzmir, Turkey
oktay.karabag@ieu.edu.tr

Fikri Karaesmen, Barış Tan
College of Engineering, Koç University, Rumeli Feneri Yolu, İstanbul, Turkey
fkaraesmen@ku.edu.tr, btan@ku.edu.tr

Effective decision-making in a circular value chain is pivotal for achieving environmental sustainability and cost efficiency. This study delves into the complexities of returned material purchasing and refurbishing decisions within a circular value chain. Specifically, it focuses on a refurbisher operating in a make-to-stock fashion, considering the refurbishment of first- and second-quality returned materials into corresponding finished goods. The refurbisher must make critical decisions regarding purchasing returned materials and determining the refurbishing processes, encompassing five distinct refurbishing options and two purchasing decisions. The optimal control problem is formulated as a Markovian Decision Process, and a linear programming approach is employed to derive the optimal policy. Analysis of the optimal policy for a wide range of system parameters shows that the optimal refurbishment policy is driven by the per-unit profits for different refurbishment options and also by the returned material, demand rates, and production times for different conversion options. Hence, this result emphasizes the necessity of a comprehensive model capturing the dynamics of material and demand flows, together with the prices and costs, for making the right refurbishment and purchasing decisions in a circular value chain.

Key words: Circular Economy; Refurbishment; Markov Decision Process

1. Introduction

Product reuse is not a recent phenomenon but has gained significant traction in recent decades. Product reuse is now recognized as a profitable and sustainable strategy for many companies. Motivations for adopting such reusing strategies include economic, legislative, and environmental factors. For such systems, making the right decisions is more complicated due to the differences in the quality of returned material flows and refurbished products, different demand patterns for refurbished products, and differences in the purchasing and refurbishing costs, among others. Extensive systematic overviews of such systems and possible challenges related to their operational decisions have been proposed by Dekker (2004) and Khan et al. (2021).

Motivated by a business model of packing returned battery cells of different remaining capacities into different quality refurbished packs intended for different uses, here, we consider a discrete-material flow remanufacturing system in which a single facility with a limited production capacity operates in a make-to-stock manner to fulfill the demand for a single product's high- and low-quality versions. The system is portrayed in Figure 1 and its details are given in the next section. In the relevant literature, Vercraene et al. (2014) and Nadar et al. (2023) are the most pertinent studies to our work. However, our work differs from these studies by using a stochastic modeling approach to address an integrated purchasing and remanufacturing decision for a remanufacturing system operating in a make-to-stock fashion.

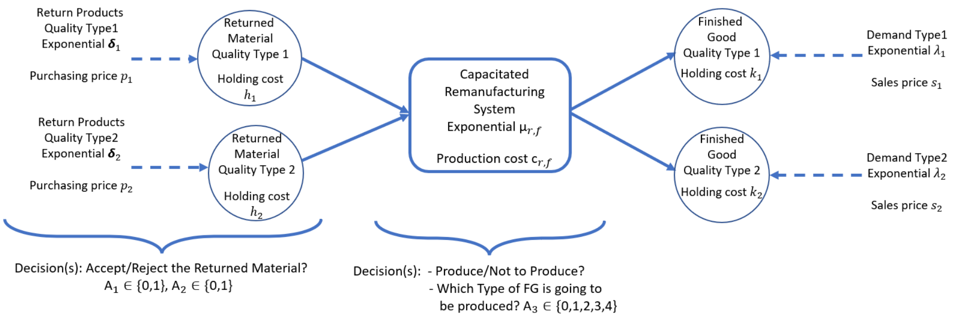

Figure 1. Purchasing, production and remanufacturing decisions for a system with two returned material and two finished good quality levels

2. Problem Definition

The system has two different forms of inventory to perform its operations: (*i*) returned material and (*ii*) finished goods inventories. The remanufacturer observes the first- and second-quality returned material levels directly, as well as the inventory levels of first- and second-quality finished goods.

The arrival processes of the returned material types are modeled as two independent Poisson processes, i.e., the inter-arrival times of the first- and second-quality returned materials are considered to be independent of each other, and they are exponentially distributed with rates δ_1 and δ_1, respectively. Each time a returned material unit at any quality level arrives at the system, the remanufacturer decides whether to buy it. If the remanufacturer decides to purchase a first-quality (second-quality) returned material that has arrived at the system, she pays a purchasing cost of p_1 (p_2). The costs of keeping one unit of stock in the first- and second-quality returned material buffers are assumed to be h_1 and h_2, respectively.

The remanufacturer makes five different types of decisions related to the production process. She may decide (*i*) not to produce, (*ii*) to produce a first-quality finished product from a first-quality returned material, (*iii*) to produce a first-quality finished product from a second-quality returned material, (*iv*) to produce a second-quality finished product from a first-quality returned material, and (*v*) to produce a second-quality finished product from a second-quality returned material. The system has a limited production capacity, processing one returned material unit at a time. The system's returned material processing times follow an exponential distribution whose rate depends on the remanufacturer's production decision. The average time being spent to process a unit of returned material of quality r to a unit of finished good of quality f is denoted by $1/\mu_{r,f}$ where $r, f \in \{1, 2\}$. Furthermore, the processing cost of a unit of returned material is $c_{r,f}$.

Two customer types arrive at the system, demanding the same finished goods but at different quality levels. The inter-arrival times of the customers demanding first- and second-quality finished goods are considered to be independent of each other, and they are exponentially distributed with rates λ_1 and λ_2, respectively. The manufacturer earns s_1 and s_2 for each sales transaction of first-and second-quality finished goods, respectively. It is assumed that $s_1 > s_2$. We also assume a stock-out-based substitution. Suppose a demand for a high-quality (low-quality) product arrives. In such a case, if the high-quality (low-quality) stock is empty and a low-quality (high-quality) product is available, $\varphi_{1,2}$ ($\varphi_{2,1}$) percent of the customers substitute the high-quality (low-quality) product with a low-quality (high-quality) one. Additionally, the costs of keeping one unit of stock in the first- and second-quality finished goods buffers are assumed to be k_1 and k_2, respectively.

Given that the remanufacturer fully observes her inventory levels and market prices for her finished goods, she attempts to determine the optimal returned material purchasing and refurbishment strategies that maximize her average profit over an infinite planning horizon.

The inter-event times in the model are considered to be exponentially distributed random variables; therefore, the system dynamics can be characterized as a continuous time Markov Chain. This enables us to model our problem as a Markovian decision process and derive its optimal policies using well-known solution methods. In this study, we employ the Linear Programming (LP) approach, among several other alternatives, to obtain the optimal policy for the system. Using the LP approach empowers us to leverage state-of-the-art optimization software packages, enabling a more expedient problem resolution. As a result, the identification of the optimal policies can be achieved in significantly less time compared to classical methods like value/policy iterations. For the sake of brevity, we refrain from presenting our Markov Decision Process model and Linear Programming formulation here.

3. Results

Using the LP formulation we devised, we conducted extensive numerical experiments to examine the factors that affect the optimal policy structure and various performance measures. We show that the optimal refurbishment policy depends on the sales prices of different quality products, costs of different quality returned materials, production costs for different refurbishment conversions, the returned material, demand rates, and production times. Therefore, the dynamic stochastic model presented in this study allows managing the refurbishment operations more effectively. Furthermore, following the event-based dynamic programming framework, we analytically show that the optimal purchasing and conversion decisions are of state-dependent threshold policies.

Acknowledgments

European Union's Horizon 2020 Research and Innovation Programme [AutoTwin project under grant agreement No. 101092021].

References

Dekker, Rommert. 2004. *Reverse logistics: quantitative models for closed-loop supply chains.* Springer Science & Business Media.

Khan, Syed Abdul Rehman, Zhang Yu, Heris Golpira, Arshian Sharif, Abbas Mardani. 2021. A state-of-the-art review and meta-analysis on sustainable supply chain management: Future research directions. *Journal of Cleaner Production* **278** 123357.

Nadar, Emre, Mustafa Akan, Laurens Debo, Alan Scheller-Wolf. 2023. Optimal procurement in remanufacturing systems with uncertain used-item condition. *Operations Research* .

Vercraene, Samuel, Jean-Philippe Gayon, Simme Douwe Flapper. 2014. Coordination of manufacturing, remanufacturing and returns acceptance in hybrid manufacturing/remanufacturing systems. *International Journal of Production Economics* **148** 62–70.

MTO vs. MTS: Customer Preferences and Firm Strategies in Competitive vs. Cooperative Markets

George Liberopoulos, Myron Benioudakis, Michalis Deligiannis
Department of Mechanical Engineering, University of Thessaly, Volos, Greece, glib@uth.gr, benioudakis@aueb.gr, deligiannis@uth.gr

Apostolos Burnetas
Department of Mathematics, National and Kapodistrian University of Athens, Athens, Greece, aburnetas@math.uoa.gr

We study a market model of two firms offering substitute products to price-and-time-sensitive customers. One firm produces customized MTO products on a first-come, first-served basis, while the other is a MTS supplier offering standard products on demand. Customers decide which product to purchase based on their valuations of the products, the prices, and the estimated waiting time for the customized products, or they go elsewhere. We analyze the equilibrium strategies of the firms competing on output rates and compare these strategies to the optimal joint strategies under a cooperation scenario. Additionally, we explore the effect of the MTO firm's production capacity and the MTS firm's storage capacity on market outcomes. Through numerical analysis, we investigate how various parameters, including those related to product valuation distribution, influence market dynamics.

Key words: Customized products; Standard products; Competition; Cooperation; Market share; Output rate; Pricing; Strategic customers

1. Introduction and Problem Formulation

Customer perceptions of product value vary significantly, influenced by functional, emotional, life-changing, and social impact needs. To address this diversity, companies increasingly use technology-enabled mass customization to tailor goods to individual needs, though these customized products are more expensive and have longer wait times compared to standard products that benefit from economies of scale and immediate availability. Despite higher prices, customized products offer greater benefits, with many consumers willing to pay more and wait longer for them. However, excessive waiting times or high prices can drive customers to substitutes, as seen during the COVID-19 pandemic in the car market. Companies must strategically price their products to balance attractiveness and profitability, while many mass-production manufacturers are incorporating customization services, and luxury brands are embracing standardization for profitability. Effective customization poses significant challenges for mass producers, often leading to the discontinuation of such services. Leveraging partnerships, acquisitions, and outsourcing can provide solutions for offering effective customization without starting from scratch.

This study develops a model analyzing competitive and cooperative strategies between a make-to-order (MTO) supplier and a make-to-stock (MTS) supplier in a market of price-and-time-sensitive strategic customers. Customers decide which product to buy based on their valuations of the products, the product prices, and the estimated waiting time for customized products. The study investigates the Nash equilibrium strategies of both firms when competing on output rates and compares them to optimal joint strategies if they cooperate, exploring the impact of production and storage capacities on market outcomes and customer behavior through numerical analysis. The analysis draws mainly from the stream of research on rational queueing (Hassin and Haviv 2003, Hassin 2016). A recent work in this area by the authors is Deligiannis et al. (2024), while a related work is Zhou et al. (2023).

The model that we consider features a market with two firms, F1 and F2, offering substitute products to randomly arriving strategic customers who may demand one unit of a product. F1 produces customized MTO products at unit cost c_1, following a FIFO M/M/1 queueing model with an average service time of $1/\mu$. F2 sells standard MTS products at unit cost c_2, with negligible stock replenishment lead times, ensuring prompt service. For most of our analysis, we consider c_2 to be fixed except in our sensitivity analysis where we assume that it is decreasing in F2's storage capacity, u, due to economies of scale. Customers arrive according to a Poisson process with rate Λ and independently choose between F1, F2, or balking based on their individual product valuations, R for F1 and θR for F2, where $\theta \in (0,1)$. In steady-state, the arrival rates to F1 and F2 are λ_1 and λ_2 respectively. The expected waiting cost for F1 customers is $W(\lambda_1) = \frac{w}{\mu-\lambda_1}$, and the expected utilities for joining F1 and F2 are $U_1(\lambda_1) = R - W(\lambda_1) - p_1$ and $U_2 = \theta R - p_2$, respectively. The firms' expected profits are given by $G_i(\lambda_i, p_i) = \lambda_i(p_i - c_i)$ for $i = 1, 2$.

2. Analysis and Numerical Experiments

Customers act independently to maximize their utility, leading to a symmetric aggregate equilibrium strategy. They choose between joining F1, F2, or balking based on comparing $U_1(\lambda_1)$, U_2, and zero (utility of balking). The thresholds $r_{1,0}$, $r_{2,0}$, and $r_{1,2}$ indicate customer indifference points, where $r_{1,0} = W(\lambda_1) + p_1$, $r_{2,0} = \frac{p_2}{\theta}$, and $r_{1,2} = \frac{W(\lambda_1)+p_1-p_2}{1-\theta}$. These thresholds determine the best response: joining F1 if $R \geq r_{1,2}$, joining F2 if $r_{2,0} \leq R \leq r_{1,2}$, or balking if $R \leq r_{2,0}$. Two market cases arise: a duopoly (when $W(\lambda_1) + p_1 > p_2/\theta$) with positive arrival rates for both firms, and a monopoly (when $W(\lambda_1) + p_1 \leq p_2/\theta$) where only F1 has positive arrivals. The equilibrium arrival rates, $\lambda_1^e(p_1, p_2)$ and $\lambda_2^e(p_1, p_2)$, depend on these conditions and are derived from solving specific equations involving the valuation distribution and waiting costs. The equilibrium prices p_1^e and p_2^e are characterized, showing dependence on the arrival rates and illustrating the dynamic pricing strategies under different market conditions.

The firms' strategies depend on market structure. In a competitive market, each firm operates independently with the aim of maximizing its own expected profit. This setting creates a competitive game where each firm's optimal strategy depends on the output rate of the other firm. The best response functions of the firms are characterized by properties such as decreasing behavior and constraints based on the distribution of customer arrivals. These properties ensure the uniqueness of the Nash equilibrium, where both firms' strategies are mutually optimal. The conditions for equilibrium differ depending on factors such as cost structures and customer arrival rates, with the equilibrium shifting between a duopoly and monopoly scenario based on certain thresholds.

On the other hand, in a cooperative market, firms collaborate to optimize joint outcomes, aiming to maximize their combined total payoff. This approach involves jointly setting prices or output rates to leverage each firm's strengths. The optimal strategies in a cooperative setting lead to a unique global maximizer for the firms' joint payoff. However, the conditions for optimality are complex and depend on factors such as cost parameters and the shape of the distribution of customer arrivals. The cooperative model contrasts with the competitive one by emphasizing joint profit maximization over individual firm competition.

Comparing firm competition and cooperation, it is evident that cooperation yields higher total payoffs for the firms, as it leverages synergies between them. In contrast, competition focuses on individual profit maximization, which can lead to suboptimal outcomes in terms of overall industry performance. The analysis underscores the trade-offs between competition and cooperation in different market structures, highlighting how the choice between them can significantly impact firm profitability and industry dynamics.

We explore the impact F1's production capacity (μ) and F2's storage capacity (u) on market outcomes. We establish propositions regarding the conditions under which the equilibrium market is a duopoly or a monopoly based on μ and u. Moreover, we derive monotonicity properties of the Nash equilibrium arrival rates and equilibrium prices with respect to μ and u, illustrating the intricate interplay between production and storage capacities in determining market dynamics. These findings provide insights into the strategic decisions firms must make to enhance their competitiveness and navigate market conditions effectively.

Additionally, we conduct numerical experiments to delve deeper into the implications of various parameters, including μ, u, and customer heterogeneity, on market outcomes, firm profitability, customer benefit, and system efficiency under both competitive and cooperative scenarios. By considering a normal distribution for customers' valuations, we illustrate the impact of their heterogeneity. The experiments reveal subtle insights, such as the initial decrease in total profit under competition due to low storage capacity for F2, and the imbalance between increased firm profitability and decreased customer surplus under cooperation. Sensitivity analyses on storage and production capacities validate analytical findings and highlight the differential impact of these capacities on prices and customer benefits under competition and cooperation. Furthermore, exploring the effect of customer valuation similarity parameter (θ) uncovers scenarios where F2 is effectively excluded from the market under both competitive and cooperative settings, emphasizing the complex dynamics between firms and customers.

3. Conclusions

This study sheds light on the complex dynamics of competition and cooperation between MTO and MTS suppliers in a market of price-and-time-sensitive strategic customers. By analyzing equilibrium strategies, firm behaviors, and market outcomes, valuable insights emerge regarding the impacts of customer heterogeneity, production and storage capacities, and strategic pricing on firm profitability and industry dynamics. The findings highlight the trade-offs between competition and cooperation, with cooperation offering higher total payoffs for firms. Numerical experiments deepen our understanding of these dynamics, revealing the subtle implications of various parameters on market outcomes, firm profitability, customer benefit, and system efficiency.

References

Deligiannis, M., M. Benioudakis, M. Liberopoulos, A. Burnetas. 2024. Optimal pricing and capacity management in service systems with delay-sensitive mixed-risk customers. *International Journal of Production Research* (online) 1–27. doi:10.1080/00207543.2024.2333936.

Hassin, R. 2016. *Rational Queueing*. Chapman & Hall book. CRC press. doi:10.1201/b20014.

Hassin, R., M. Haviv. 2003. *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*, vol. 59. Springer Science & Business Media. doi:10.1007/978-1-4615-0359-0.

Zhou, W., W. Huang, V. N. Hsu, P. Guo. 2023. On the benefit of privatization in a mixed duopoly service system. *Management Science* **69**(3) 1486–1499. doi:10.1287/mnsc.2022.4424.

On the generalization of Markov models for single-echelon inventory systems

Elisa Gebennini
Facoltà di Scienze Tecnologiche e dell'Innovazione, Universitas Mercatorum, Roma, RM, Italy elisa.gebennini@unimercatorum.it

Andrea Grassi, Liberatina Carmela Santillo, Emma Salatiello
Dipartimento di Ingegneria Chimica dei Materiali e della Produzione Industriale (DICMAPI), Università degli Studi di Napoli Federico II, Napoli, NA, Italy andrea.grassi@unina.it, liberatinacarmela.santillo@unina.it, emma.salatiello@unina.it

At the present, modern supply chains deal with increased variability in the demand patterns and lead times. Hence, the need for accurate stochastic models arises, specifically for the lost-sale case which occurs when any demand in excess of current inventory is lost. Early studies showed that Markov processes can be adopted in order to model inventory systems with different inventory control policies. Specially, the building block composed of two machines and one buffer can be rethought to model a single-echelon inventory systems. The objective of this work is to investigate directions and strategies for the development and generalization of this kind of modelling approaches with the aim of better addressing real industrial conditions.

Key words: inventory control, single-echelon inventory system, stochastic lead time, stochastic demand

1. Introduction

Inventory models have garnered the attention of both practitioners and researchers over an extensive period. Early studies relied on deterministic assumptions, whereas contemporary investigations have incorporated the impact of variability on both demand and supply lead time (Zipkin 2000). Variability in demand and lead time may lead to stock-outs, resulting in either lost sales or back-orders. Accordingly, carefully calibrated safety stocks are necessary to mitigate stock-out risks. The foremost methodologies for inventory management are (i) the periodic-review policy, wherein orders are triggered at fixed intervals, and (ii) the continuous-review policy, whereby a fixed-size order is released each time the inventory position descends to a predefined reorder point. In general, as evidenced in the scientific literature (Bijvank and Vis 2011, see, e.g.), the modeling of the lost-sales scenario poses the greatest challenge. This is because the amount of demand lost during an inventory cycle results in an additional inventory availability in the next cycle. This highlights a gap in research on lost-sales models within the continuous-review inventory policy.

Another important aspect related to inventory modelling in supply chains is the length of the supply chain itself. Basically, contemporary literature (de Kok et al. 2018) distinguishes among three main scenarios for supply chain planning and control: single-echelon, dual-echelon and multi-echelon. A single-echelon system is a supply chain system with only one inventory location; in a dual-echelon system there are two inventory locations; in a multi-echelon system the inventory locations are more than two. Early in the '60s, a seminal paper by Clark and Scarf (1960) put the attention on the fact that inventories are distributed along the actors forming a supply chain and that the optimal inventory policy must be defined by taking into account the interactions between them. This work generated a research thread that produced a number of works. As a matter of fact, the evolution of supply chain design in recent years put the attention on the responsiveness as a key factor to increase profit. Hence, a now-common approach is to split the supply chain into two parts (Chopra 2019), the responsive (*pull*) one and the efficient (*push*) one. The *pull* part, usually confined to the last inventory location of the supply chain, works on a profit basis, accepting a cost increase where a more than proportional growth in revenues can be generated by improving the service level (e.g. increased variety, reduced delivery time, etc.). On the other hand, the *push* part works on a cost reduction basis (i.e. it is requested to be efficient) since it faces a reduced variability scenario, being the most of the variability absorbed by the responsive part. Thus, it is once again a relevant and pressing issue to properly model single-echelon inventory systems characterized by high variability in demand and lead times, which are suitable to represent the responsive part of a supply chain. Specifically, the relevance of random lead times studies, which has been observed since the end of the '90s (see, e.g., Bookbinder and Çakanyildirim 1999), is even more remarkable in recent supply chain contexts.

2. Problem definition

Since, as discussed above, single-echelon inventory systems are experiencing an increasing interest (especially for modelling the responsive part of a supply chain), the idea arises of treating a single-echelon inventory system similarly as a *building block* in a production line by exploiting the achievements obtained in years of research on the topic of production line modelling. In the production line theory, the *building block* is composed of two machines and a buffer in between. A single-echelon inventory system can be viewed as the counterpart of what a building block in the production line modelling theory is (see reviews such as Dallery and Gershwin 1992, Papadopoulos et al. 2019), where the first and second machines represent the supply and demand processes, respectively, and the buffer is the inventory location.

Moreover, in recent years authors such as Ang et al. (2013) and Lagodimos et al. (2018) introduced a novel approach for one of the most famous models for a single-echelon inventory system, i.e. the continuous-review EOQ model. They proposed to discretize the time line allowing for discrete demand and discrete inventory level to be treated. Thus, the concept of discrete-time EOQ (DT-EOQ) has been defined which can fit with modelling approaches based on discrete-time discrete-state Markov chains. Figure 1 shows a sample of the on hand inventory over time generated by a simulation model which moves over time in discrete steps (i.e., the time units) and represents the dynamics of a DT-EOQ model. In this case, in each time unit transitions are evaluated as a function of the current system state (i.e., the inventory level), the consumption probability p_2 and the supply probability p_1. It can be noted that different inventory cycles have different time lengths and different lead times while maintaining the same supply lot of a fixed size Q, while r is the reorder point.

Consequently, innovative discrete-time discrete-state Markov chains can be designed to tackle the inventory control issue, considering stochastic demand and lead time. Specifically, in initial investigations, binomial daily demand and geometrically distributed lead time may be modeled. By appropriately manipulating the chain's structure and applying partitioning techniques and other approaches derived, e.g., from the field of phase-type distributions, the authors are confident in further generalizing both demand and lead time.

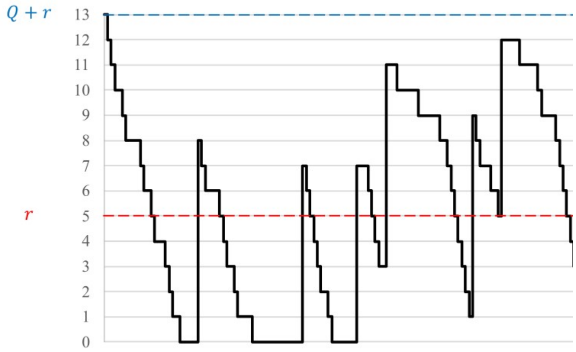

Figure 1. Example of on hand inventory level in a DT-EOQ system.

3. Conclusion

This paper investigates the opportunity of applying Markov approaches to the modelling of single-echelon inventory systems. The adoption of a discrete-time framework, such as in the DT-EOQ model, allows us to exploit discrete-time discrete-state Markov chains in order to provide exact solutions for the steady-state probability distribution of the inventory level and compact formula for the main performance measures (such as the fill rate, the average inventory level, the average length of the inventory cycle).

Acknowledgments

This study was carried out within the MICS (Made in Italy – Circular and Sustainable) Extended Partnership and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1551.11-10-2022, PE00000004). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

References

Ang, Marcus, Jing-Sheng Song, Mingzheng Wang, Hanqin Zhang. 2013. On properties of discrete (r, q) and (s, T) inventory systems. *European Journal of Operational Research* **229**(1) 95–105.

Bijvank, M., I. F. A. Vis. 2011. Lost-sales inventory theory: A review. *European Journal of Operational Research* **215**(1) 1–13.

Bookbinder, James H., Metin Çakanyildirim. 1999. Random lead times and expedited orders in (Q,r) inventory systems. *European Journal of Operational Research* **115**(2) 300–313.

Chopra, Sunil. 2019. *Supply Chain Management: Strategy, Planning, and Operation*. 7th ed. Pearson Education, Boston, USA.

Clark, A.J., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Science* **6**(4) 475–490.

Dallery, Yves, Stanley B Gershwin. 1992. Manufacturing flow line systems: a review of models and analytical results. *Queueing Systems* **12**(1) 3–94.

de Kok, T., C. Grob, M. Laumanns, S. Minner, J. Rambau, K. Schade. 2018. A typology and literature review on stochastic multi-echelon inventory models. *European Journal of Operational Research* **269**(3) 955–983.

Lagodimos, A.G., K. Skouri, I.T. Christou, P.T. Chountalas. 2018. The discrete-time EOQ model: Solution and implications. *European Journal of Operational Research* **266**(1) 112–121.

Papadopoulos, C. T., J. Li, M. E. J. O'Kelly. 2019. A classification and review of timed Markov models of manufacturing systems. *Computers and Industrial Engineering* **128** 219–244.

Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw-Hill, New York.

Slowing down e-commerce deliveries - Empirical findings and results of a stochastic-dynamic solution approach

Stefan Voigt, Markus Frank, and Heinrich Kuhn
Catholic University Eichstätt-Ingolstadt, Chair of Supply Chain Management & Operations, Auf der Schanz 49, 85049 Ingolstadt
stefan.voigt@ku.de, markus.frank@ku.de, and heinrich.kuhn@ku.de

Slow logistics is a new logistical concept to reduce logistics costs and to improve eco-efficiency. In e-commerce, this concept can be realized by offering customers the option of accepting an extended delivery time for their online purchases. This enables greater bundling options for a parcel service provider. However, this requires that customers are at all willing, or at least can be motivated, to wait longer for their online orders. The paper presents the results of an online survey in German-speaking countries, which investigates whether customers are willing to wait longer for their online orders or whether certain incentives are required, e.g., savings on shipping costs or information on reducing CO2 emissions, to increase customers' willingness to wait. The empirical findings lead to a new last-mile delivery problem that allows customers to choose from different delivery speeds when placing an order in a stochastic-dynamic order arrival environment. To solve the entire decision problem model, we develop a stochastic-dynamic solution method that solves prize-collecting vehicle routing subproblems. It becomes obvious that a slower delivery option could reduce logistics costs without significantly increasing the waiting times of an average delivery order.

Key words: Last-mile delivery; Sustainability; Consumer behavior; Vehicle routing problem; Stochastic optimization

1. Motivation and problem setting

Parcel deliveries are at an all-time high, amounting for 161 billion parcels worldwide in 2022. This is accompanied by a growing demand for even faster deliveries in the business-to-consumer (B2C) sector. However, short delivery times put tremendous pressure on transportation networks and often lead to less efficient distribution processes. We explore a delivery concept that deliberately slows down the logistics processes involved in parcel delivery, thereby allowing for the consolidation of more shipments over an extended time period. For instance, Amazon already provides a "free no-rush" delivery option in many U.S. regions. Customers who select this option accept a longer delivery time frame and, in return, e.g., receive a discount. Alternatively, customers may be nudged by information highlighting the reduced emissions associated with slower delivery, which could encourage them to accept the potential delay Dietl et al. (2023). We assume the point of view of an e-commerce retailer who operates its own delivery fleet and offers a range of delivery options, including a notably slower *some-day* alternative. Each customer's order is subject to specific delivery date constraints, comprising both an earliest and a latest possible delivery date. The retailer must determine the most cost-effective delivery day for each customer, as well as the clustering and routing. We term this problem setting *some-day delivery problem* (SDDP). Our modeling assumptions align with the broader category of multiple period VRPs (MPVRP) (Archetti et al. 2015). In practice, customer orders arrive continuously throughout the planning period, creating a dynamic setting. While a similar dynamic MPVRP, has been proposed by Wen et al. (2010), our situation differs in that not all relevant information is known in advance but may be available in a stochastic manner.

2. Analysis on customers' willingness to wait (w2w)

Dietl et al. (2023) investigate whether customers in German-speaking countries are willing to wait (w2w) longer for their online orders in the fashion, shoes, and accessories sector when offered incentives in return. The research employed an online survey using a within-subjects design. The study shows that a significant share of participants (49.3 %) would not like to put on the ordered goods until the next week or even not until the next month. This group of customers could be offered an extended delivery period. The control group of the study participants were asked whether they would be w2w 4 to 7 days for the delivery of their order instead of 2 to 3 days. Other randomly selected participants were offered various incentives if they were willing to accept the longer delivery time. Results indicate that incentives significantly extend customers' w2w. Savings on shipping costs and information on the reduction in CO_2 emissions greatly impact customers' patience. Additionally, the research highlights that customers' responses to the incentives vary based on individual characteristics, such as gender, age, environmental awareness, and the urgency of their orders.

3. Stochastic and dynamic solution approach

We use a deterministic model, named SDDP, as the basis for our stochastic-dynamic modeling approach. The deterministic model SDDP is inspired by an MPVRP proposed by Archetti et al. (2015). In its dynamic setting, only a subset of customer orders is known at planning instant, i.e., only customers revealing their demand in the previous period. The customer demands within the following periods are uncertain and stochastic. New orders arrive at the end of each period. A periodic re-planning is carried out each period to account for the newly arrived orders and updated demand forecasts. We define a benefit measure for each known customer indicating the value of serving the customer in the current period rather than postponing. We henceforth call this benefit measure "*prize*." The *prize* combines positive and negative effects on the advantageousness of servicing an order in the current period, including the following aspects: urgency of the order, future capacity utilization, the probabilities of emerging nearby customers in future periods, as well as inventory costs. As intuition, the prize should be high if the order is urgent or we expect a high demand in future periods. Contrary, the prize should be low if we expect nearby customers emerging in future periods. This prize is then used to solve an auxiliary prize-collecting VRP with time windows (PCVRPTW), that decides which customers to deliver in that period and the corresponding routing. The PCVRPTW is solved heuristically with a hybrid adaptive large neighborhood search with granular insertion operators (HALNS-G). The HALNS-G extends the original version Voigt et al. (2022, 2023) with problem-specific operators, in particular granular insertion operators. The concept of granular insertion operators is inspired by the granular tabu search Toth and Vigo (2003). In the granular insertion operators, the insertion positions are confined to customers located in close proximity to the customer being inserted.

4. Sample results and contribution

In our experiments we introduce a simulator based on VRPTW instances with 1,000 customers. In each period, the simulator randomly samples a set of 100 customers from the respective instance. This set is then revealed as the pending customer set for the current planning period. For each instance, we generate 30 periods and calculate several performance measures (e.g., the average costs per customer, the average and maximum number of vehicles used per period). We conduct several experiments to generate managerial insights such as the cost reduction against several benchmark policies, the impact of the length of the delivery interval, and the impact of the share of customers selecting the some-day option. As an exemplary experiment, Table 1 shows the results on varying the length of the delivery interval of the some-day option $(1, 2, \ldots,$ or 5 days). We compare the average costs against the costs achieved with an earliest policy. In this policy, we serve each customer as quickly as possible, i.e., within the period following the order. Compared to the earliest policy, we can reduce costs to 78.9 % with a delivery interval of just 1 day. The costs savings increase with longer delivery intervals but become increasingly marginal, indicating that a moderate interval length of 3 days seems sufficient.

Delivery interval [days]	1	2	3	4	5
Average costs vs. earliest policy [%]	78.9	68.1	61.3	56.8	54.0

Table 1. Average costs for various delivery interval lengths of the some-day option

Contribution The paper contributes by (1) describing a novel slow logistics concept for B2C parcel delivery, (2) reviewing and categorizing MPVRPs with delivery dates, (3) introducing a solution approach for a dynamic MPVRP with stochastic information, (4) implementing a hybrid adaptive large neighborhood search with granular insertion operators for solving prize-collecting VRPTWs, and (5) showing by simulation that a slow delivery option significantly improves costs.

References

Archetti, Claudia, Ola Jabali, M Grazia Speranza. 2015. Multi-period vehicle routing problem with due dates. *Computers & Operations Research* **61** 122–134.

Dietl, Melanie, Stefan Voigt, Heinrich Kuhn. 2023. From rush to responsibility - evaluating economic and ecological incentives on customers' willingness to wait for online orders. *SSRN Electronic Journal* doi:10.2139/ssrn.4551860. URL `https://doi.org/10.2139/ssrn.4551860`.

Toth, Paolo, Daniele Vigo. 2003. The granular tabu search and its application to the vehicle-routing problem. *INFORMS Journal on Computing* **15**(4) 333–346. doi:10.1287/ijoc.15.4.333.24890.

Voigt, Stefan, Markus Frank, Pirmin Fontaine, Heinrich Kuhn. 2022. Hybrid adaptive large neighborhood search for vehicle routing problems with depot location decisions. *Computers & Operations Research* **146** 105856. doi:https://doi.org/10.1016/j.cor.2022.105856.

Voigt, Stefan, Markus Frank, Pirmin Fontaine, Heinrich Kuhn. 2023. The vehicle routing problem with availability profiles. *Transportation Science* **57**(2) 531–551. doi:10.1287/trsc.2022.1182.

Wen, Min, Jean-François Cordeau, Gilbert Laporte, Jesper Larsen. 2010. The dynamic multi-period vehicle routing problem. *Computers & Operations Research* **37**(9) 1615–1623.

Time-Dependent Shipment Options and Shipment Fees for E-Commerce Fulfillment Centers

Uta Mohring
Rotman School of Management, University of Toronto, Toronto, Canada, uta.mohring@rotman.utoronto.ca

Melvin Drent, Ivo Adan, Willem van Jaarsveld
Eindhoven University of Technology, Eindhoven, the Netherlands, {m.drent, i.adan, w.l.v.jaarsveld}@tue.nl

Key words: Pricing; shipment options; e-commerce; Markov chain; demand management

---

Given the fast growth of e-commerce, fulfilling online orders faster and cheaper becomes more relevant than ever. In that respect, e-commerce companies increasingly promise same-day shipment. Same-day shipment is a powerful marketing tool to attract and retain customers. Nowadays, customers place orders in the afternoon or evening and expect these products to be shipped immediately so that they arrive the next day. From an operational perspective, however, same-day shipment promises are challenging to fulfill, especially by the end of the day when the time window for same-day shipment becomes tighter. Hence, same-day shipment requires a well-coordinated effort between marketing and operations to ensure a smooth and efficient order fulfillment process.

Companies that sell their products online typically fulfill the orders from so-called fulfillment centers. In these centers, orders are picked from the shelves and subsequently prepared for shipment. At fixed predefined deadlines, typically at the end of the day, the orders are consolidated into large batches and handed over to the external parcel delivery company responsible for the actual delivery of the products (e.g., FedEx or UPS). Meeting these strict deadlines is crucial for the fulfillment process as missing the departure of a delivery vehicle by only a few minutes may result in a delay to the customer of at least one day (Doerr and Gue 2013). In their pursuit of offering increasingly faster service, e-commerce companies might run the risk that they over-promise their shipment services. Indeed, the fulfillment center may not be able to keep up with the inflow of orders that require same-day shipment and need to be ready before the set deadline. This risk is particularly high when companies charge only flat fees for the different shipment options. Under such flat fees, customers cannot be incentivized to choose next-day shipment over same-day shipment when the operations in the fulfillment center would require so.

Time-dependent shipment options and shipment fees

As a natural mitigation strategy to the risk described above, e-commerce companies may want to introduce a time- and fulfillment center load-based dynamic shipment policy. Under such a dynamic shipment policy, the offered shipment options and the corresponding shipment fees can be used as levers to better balance the inflow of orders that require same-day and next-day shipment with the available capacity and remaining time to do so. While such a dynamic shipment policy has merit from an operational perspective, it may be perceived as unfair by customers (Klein et al. 2019), and it is arguably difficult to implement from a marketing perspective. Static shipment options and corresponding fees that depend only on the remaining time to the set deadline seem to be more appropriate and promising as well, especially as customers are rather sensitive to small price changes (Acimovic and Graves 2017). Such well-targeted static time-dependent shipment options and fees have the potential to influence customers' shipment preferences in order to balance the overall workload and increase the capacity utilization in fulfillment centers, very much akin to a dynamic shipment policy, but without eliciting perceptions of unfairness among customers. Further, static time-dependent shipment options and fees can be easily and transparently communicated to customers on the order website.

In this research, we study the potential of time-dependent shipment options and fees in managing the distribution of customers demanding same-day or next-day shipment. When e-commerce companies operate such time-dependent shipment policies, they need to decide for each moment in time which shipment options they offer as well as the fees that correspond to these shipment options. We study precisely these decisions. Because of its practical appeal, we focus on so-called cutoff-based shipment options under which same-day and next-day shipment are offered to all customers placing their orders until a certain cutoff point while customers ordering after the cutoff point will then only be offered next-day shipment (Mohring et al. 2024). Given the cutoff point decision, shipment fees need to be set for each moment in time. Thus, the following interrelated questions arise:

(1) Cutoff-based shipment options: Until what moment in time should same-day shipment still be offered to customers?

(2) Shipment fees: How much should the shipment fees be?

(3) Differentiated fees: When and how should the shipment fees be adapted over time?

Methodology

We build a parsimonious model that provides answers to the questions posed above. We consider a multi-period model, where each period ends with a deadline upon which orders that are due for shipment in that period need to be handed over to the parcel delivery company. Throughout each such period, online customers arrive according to a Poisson process. Customers who complete their online transaction before the cutoff point are offered the choice between same-day or next-day shipment together with their corresponding fees. We assume that customers make this choice by trading off the utility that they derive from both shipment options. Customers that arrive after the cutoff point cannot choose same-day shipment, and hence their products are handed over to the parcel company upon the next deadline at the end of the following period. The processing capacity of the fulfillment center within each period is randomly distributed. If the available capacity is insufficient to process all orders due for shipment in a given period, these orders are then said to be late and carried over for shipment in the subsequent period.

We are interested in finding the optimal time-dependent shipment options and fees such that the long-run average profit consisting of revenue minus costs for late orders is maximized. By developing a discrete-time Markov chain model for the steady-state analysis of the stochastic e-fulfillment centers described above, we provide closed-form expressions for the relevant performance measures as well as structural properties of the optimal time-dependent shipment policy.

Main results and contributions

We are the first to study time-dependent shipment options and fees for stochastic e-fulfillment centers. We characterize how e-commerce companies may derive substantial benefits from introducing a static time-dependent shipment policy. More precisely:

(1) We present an exact analysis of the steady-state performance measures for the stochastic e-fulfillment centers described above based on the underlying discrete-time periodic Markov chain model.

(2) We show that the optimal time-dependent shipment policy exhibits an intuitive structure: The fee for same-day shipment increases as the time until the cutoff point of the current period decreases.

(3) Motivated by its practical appeal, we numerically study a simple instance of our time-dependent shipment policy with a two-level fee structure. We compare this simple time-dependent shipment policy with two benchmark policies prevalent in practice, which rely on static time-independent fees for both shipment options but differ in whether they include a cutoff point for the same-day shipment option and how they set the shipment fees. We find that (i) including a cutoff point under static shipment fees increases profits substantially, and that additionally (ii) moving to time-dependent shipment fees (i.e., our approach) increases profits by another considerable margin.

The key contribution of this research is that these results provide E-commerce companies with valuable and easily implementable managerial guidelines for the design of shipment options and shipment fees in E-commerce to better balance online demands for same-day or next-day shipment with the available capacity in the fulfillment centers responsible for collecting and shipping those orders.

References

Acimovic, J., S.C. Graves. 2017. Mitigating spillover in online retailing via replenishment. *Manufacturing & Service Operations Management* **19**(3) 419–436. doi:10.1287/msom.2016.0614.

Doerr, K.H., K.R. Gue. 2013. A performance metric and goal-setting procedure for deadline-oriented processes. *Production and Operations Management* **22**(3) 726–738. doi:10.1111/j.1937-5956.2012.01375.x.

Klein, R., M. Neugebauer, D. Ratkovitch, C. Steinhardt. 2019. Differentiated time slot pricing under routing considerations in attended home delivery. *Transportation Science* **53**(1) 236–255. doi:10.1287/trsc.2017.0738.

Mohring, U., C. Jacobi, K. Furmans, R. Stolletz. 2024. Managing cutoff-based shipment promises for order fulfilment processes in warehousing. *OR Spectrum* 1–31.doi:10.1007/s00291-023-00739-7.

Validation of call center workforce scheduling models

Ger KOOLE
Vrije Universiteit Amsterdam, the Netherlands, ger.koole@vu.nl

Siqiao LI
CCmath bv, Amstelveen, the Netherlands, siqiao@ccmath.com

Sihan DING
Apple Inc, China, dingsihan@hotmail.com

As part of an effort to validate call center performance models we address some technical issues to be able to perform this validation, especially the role of noise and the way to eliminate forecasting errors. We also give the main results of the validation study.

Key words: validation, call centers, simulation

1. Introduction

Many models to predict call center performance exist in the literature (see, e.g., Gans et al. (2003)), but little is known about their accuracy. This study tries to fill this gap by comparing different simulation models to the realized performance of one particular call center. The models differ in a number of ways, especially in the way arrivals, patience, agent heterogeneity, and breaks are modeled. We would like to identify the model that minimizes the mean absolute error (MAE) between the predicted and realized service level.

We focus on daily performance measures. There are a number of problems with comparing daily realized performance and the outcome of simulation models. The source of these problems is that every day is different: we do not have i.i.d. replications of the same day as is usually the case (see, e.g., Kleijnen (1995)). The issues we tackle are:

- eliminating the **system noise**. Even an exact model would not have error 0, because of random fluctuations during the day. Note that this noise is substantial, even at the daily level, as shown in Roubos et al. (2012). Therefore we would like to make a distinction between that part of the error that is due to noise and that part that is due to the imperfectness of the model, the *model error*;
- eliminating the **forecasting error**. No forecast is exact. To focus on the model error and eliminate the forecasting error we would like to use the actual rate of the inhomogeneous Poisson arrival process. However, this rate is unknown, therefore we use the actual instead. Unfortunately this creates an error by itself, because service level and actuals are negatively correlated, the actuals give more information than we are allowed to use. Therefore we have to find a way to eliminate this *cheating error*.

The contribution of this study is twofold: we develop theory on how to validate service models (Section 2) and we obtain insights in which features are crucial to model in call centers (Section 3).

2. Validating inhomogeneous models

Let us formulate our model mathematically. The r.v. Λ represents the parameters that change from day to day, which are the rate of the non-homogeneous Poisson process and the agents that are scheduled and their shifts. The performance, typically the service level obtained during a day, is denoted by X. Because X depends on Λ we write $X(\Lambda)$. Note that X is a r.v., even for fixed $\Lambda = \lambda$: its value depends on the realization of the Poisson process, the handling times, times at which agents take breaks, etc. We can also simulate various models. The service level estimation given by the simulation is written as $S(\Lambda)$. However, the arrival rate is not observed. We could replace that part of Λ by a forecast, but they are usually quite bad. Instead, we use use the actual instead of the rates. Λ in which the rates are replaced by the random realizations of the arrivals is written as $A(\Lambda)$, the corresponding simulation $S(A(\Lambda))$.

With $\mathbb{E}_\bullet$ we indicate the expectation with respect to the corresponding r.v. For example, $\mathbb{E}_S S(A(\Lambda))$ is the expected simulated performance of a random day for a random realization of the rates; $\mathbb{E}_\Lambda X(\Lambda)$ is the random performance "averaged" over the days. Note that we can interchange expectations, e.g., $\mathbb{E}_X \mathbb{E}_\Lambda X(\Lambda) = \mathbb{E}_\Lambda \mathbb{E}_X X(\Lambda)$.

We are interested in estimating $\mathbb{E}_\Lambda |\mathbb{E}_X X(\Lambda) - \mathbb{E}_S S(\Lambda)|$, which corresponds to the mean absolute error (MAE) of the service level. However, we measure $X(\Lambda)$ and $\mathbb{E}_S S(A(\Lambda))$, the latter by averaging over a sufficiently high number of simulations. We will show how to get an estimate of the MAE based on the actuals and simulations.

We can show that

$$\mathbb{E}_S S(A(\Lambda)) - X(\Lambda) \approx \mathbb{E}_S S(A^2(\Lambda)) - S(A(\Lambda)) + \mathbb{E}_S S(\Lambda) - \mathbb{E}_X X(\Lambda), \tag{1}$$

where A^2 means taking a sample twice. The part $\mathbb{E}_S S(A^2(\Lambda)) - S(A(\Lambda))$ approximates the model noise and the cheating error, and can be obtained only using simulations. The remainder is what we measure.

Simulations show that $\mathbb{E}_S S(A^2(\Lambda)) - S(A(\Lambda))$ is approximately normally distributed with a mean equal to 0.4% and a standard deviation of 3.7%. Its MAE is $\mathbb{E}_{\Lambda,A} \big| \mathbb{E}_S S(A^2(\Lambda)) - S(A(\Lambda)) \big| = 3\%$.

Our goal is to compute the MAE of $\mathbb{E}_S S(\Lambda) - \mathbb{E}_X X(\Lambda)$, $\mathbb{E}_\Lambda \big| \mathbb{E}_S S(\Lambda) - \mathbb{E}_X X(\Lambda) \big|$. We cannot simply subtract 3% from $\mathbb{E}_\Lambda \big| \mathbb{E}_S S(A(\Lambda)) - X(\Lambda) \big|$. To get a better understanding we start with computing the first two moments of the model error. Our simulations show $\mathbb{E}_{\Lambda,A} \big(\mathbb{E}_S S(A(\Lambda)) - X(\Lambda) \big) = 3.9\%$ and $\sigma_{\Lambda,A} \big(\mathbb{E}_S S(A(\Lambda)) - X(\Lambda) \big) = 5\%$.

The first moment of the model error follows directly by taking expectations in Equation (1):

$$\mu := \mathbb{E}_\Lambda \big(\mathbb{E}_S S(\Lambda) - \mathbb{E}_X X(\Lambda) \big) = 3.5\%.$$

This value is of interest in itself: it tells us to which extend the model is *biased*, to which extent there is a systematic error. But even if μ is small, the errors can be big but fluctuating, sometimes positive, sometimes negative. That is why we defined the performance measure as $\mathbb{E}_\Lambda \big| \mathbb{E}_S S(\Lambda) - \mathbb{E}_X X(\Lambda) \big|$.

Next we compute the standard deviation of the model error. As we expect $\mathbb{E}_S S(A^2(\Lambda)) - S(A(\Lambda))$ to be independent of the model error, we have

$$\sigma^2 := \sigma^2_{\Lambda,A} \big(\mathbb{E}_S S(\Lambda) - \mathbb{E}_X X(\Lambda) \big) \approx$$
$$\sigma^2_{\Lambda,A} \big(\mathbb{E}_S S(A(\Lambda)) - X(\Lambda) \big) - \sigma^2_{\Lambda,A} \big(\mathbb{E}_S S(A^2(\Lambda)) - S(A(\Lambda)) \big).$$

In our case $\sigma = \sqrt{(0.05^2 - 0.037^2)} = 3.4\%$. We conclude that the first moments of the model error and will also be quite similar to the measured error.

We can go a step further if we assume the measurements to be normally distributed. For the simulated noise/cheating factor this is the case, for the measured error this is a rough approximation. For $\mathbb{E}_S S(\Lambda) - \mathbb{E}_X X(\Lambda) \sim \mathrm{N}(0.035, 0.034)$, using a straightforward calculation, we find

$$\mathbb{E}_\Lambda \big| \mathbb{E}_S S(\Lambda) - \mathbb{E}_X X(\Lambda) \big| \approx \frac{2\sigma}{\sqrt{2\pi}} \mathrm{e}^{-\frac{1}{2}(\frac{\mu}{\sigma})^2} + \mu \Big(2\Phi\Big(\frac{\mu}{\sigma}\Big) - 1 \Big) = 4.0\%,$$

with Φ the standard normal distribution function. Note that the MAE without the correction is $\mathbb{E}_{\Lambda,A} \big| \mathbb{E}_S S(A(\Lambda)) - X(\Lambda) \big| = 4.5\%$. We conclude that the added precision given by the correction is small compared to the measured error, around 10%.

The results of this section are derived for the “HT & Patience Model”, one of the models we analyzed. Other models give similar results. Note that even the best model will give a non-zero measured error. The lower bound to the measured error is approximated by $\mathbb{E}_S S(A^2(\Lambda)) - S(A(\Lambda))$. Its MAE is 3%.

3. Validation results

We did an extensive data analysis of a call center of which we had data about the calls and the agents: who did which calls, but also when agents where available to take calls, when they took breaks, etc. Based on this analysis we compared several models each having a different set of features. The most important ones are:
- **Handling times**: can be taken empirical or exponential;
- **Average handling times**: the averages can be taken all the same or weighted based on the mix of agents available that day;
- **Patience**: empirical or exponential;
- **Breaks**: are yes or no taken into account, proportional to the length of the breaks.

We find that all models overestimate the SL. The best one in terms of the MAE is the model with handling times exponential, AHTs adapted to the agent mix, empirical patience, and breaks taken into account. The MAE is $\approx 3\%$, for a SL that is usually around 80%. The worst models are the ones that have a yearly overall AHT and do not take breaks into account. Although it looks evident, these two features are often not taken into account.

Acknowledgments

We thank VANAD Laboratories for supplying the data and Rob van der Mei and Raik Stolletz for useful discussions.

References

Gans, N., G.M. Koole, A. Mandelbaum. 2003. Telephone call centers: Tutorial, review, and research prospects. *Manufacturing & Service Operations Management* **5** 79–141.

Kleijnen, J.P.C. 1995. Verification and validation of simulation models. *European Journal of Operational Research* **82** 145–162.

Roubos, A., G. Koole, R. Stolletz. 2012. Service-level variability of inbound call centers. *Manufacturing & Service Operations Management* **14** 402–413.