



Michele Urbani

MAINTENANCE POLICIES OPTIMIZATION IN THE INDUSTRY 4.0 PARADIGM



Michele Urbani

MAINTENANCE POLICIES OPTIMIZATION IN THE INDUSTRY 4.0 PARADIGM

Dissertation for the degree of Doctor of Science (Economics and Business Administration) to be presented with due permission for public examination and criticism in the room A203, Polo Scientifico e Tecnologico Fabio Ferrari at the University of Trento, Italy, on the 10th of December 2021, at 2 PM local time.

The dissertation was written under a joint supervision (cotutelle) agreement between University of Trento, Italy and Lappeenranta-Lahti University of Technology LUT, Finland and jointly supervised by supervisors from both universities.

Acta Universitatis
Lappeenrantaensis 994

Supervisors Professor Mikael Collan
LUT School of Business and Management
Lappeenranta-Lahti University of Technology LUT
Finland

Professor Matteo Brunelli
Department of Industrial Engineering
University of Trento
Italy

Reviewers Professor Kari Koskinen
Faculty of Engineering and Natural Sciences
Tampere University
Finland

Professor Yuri Lawryshyn
Faculty of Applied Science and Engineering
University of Toronto
Canada

Opponent Professor Yuri Lawryshyn
Faculty of Applied Science and Engineering
University of Toronto
Canada

ISBN 978-952-335-742-6
ISBN 978-952-335-743-3 (PDF)
ISSN-L 1456-4491
ISSN 1456-4491

Lappeenranta-Lahti University of Technology LUT
LUT University Press 2021

Michele Urbani

Maintenance policies optimization in the Industry 4.0 paradigm

Lappeenranta 2021

87 pages

Acta Universitatis Lappeenrantaensis 994

Diss. Lappeenranta-Lahti University of Technology LUT

ISBN 978-952-335-742-6, ISBN 978-952-335-743-3 (PDF), ISSN-L 1456-4491,

ISSN 1456-4491

Maintenance management is a relevant issue in modern technical systems due to its financial, safety, and environmental implications. The need to rely on physical assets makes maintenance a *necessary evil*, which, on the other hand, allows achieving a high quality of end products, or services, and a safety level that is adequate for the regulatory requirements. The advent of the fourth industrial revolution offers meaningful opportunities to improve maintenance management; technologies such as Cyber-Physical Systems, the Internet of Things, and cloud computing enable realizing modern infrastructure to support decisions with advanced analytics. In this thesis, the optimization of maintenance policies is tackled in this renewed technological context.

The research methods employed in this thesis include interviewing of subject experts, literature research, and numerical experiments. Mathematical modelling is used to model network effects in complex technical systems, and simulations are used to validate the proposed models and methodologies. The problem of maintenance policies comparison is addressed in one of the publications; using the proposed bi-objective analysis, an effective maintenance policy was identified. Maintenance of complex systems organized in a networked fashion is studied in another project, where maintenance costs and system performances are considered. The proposed model allowed to identify a set of non-dominated (in the Pareto sense) maintenance policies, and an efficient resolution procedure was developed. The possibility to use a digital twin to replicate a Cyber-Physical System for maintenance policies optimization is addressed in another publication. The main hurdles in realizing such a complex infrastructure are analyzed, and managerial implications are presented. Finally, following a qualitative research approach, the opportunities offered by additive manufacturing are identified and presented in a book chapter. The opportunities for both maintenance efficiency gains and new business models are identified and discussed.

Keywords: Maintenance, Optimization, Industry 4.0, Digital Twin, Heuristic methods

My first acknowledgement goes to Professor Matteo Brunelli for his guidance and his patience. His openness to discussion and his availability were priceless to me. Moreover, his encouragement to travel and his trust let me grow both as a young researcher and a person. To Professor Mikael Collan, it goes my acknowledgement for his support and guidance. His ability to connect people and his mentorship were priceless to me during my periods of stay in Finland. Both my advisors were always supportive, they gave me the possibility to work and study in two amazing countries, Italy and Finland, and they always encouraged me to travel and to go to conferences. I want to thank Professor Antti Punkka for his hospitality and the precious collaboration we had during my stay at the System Analysis Laboratory.

For their support and their endless care, I thank my family, my mother Cristina, my father Francesco, and my sister Martina. I thank my grandparents, Giselda and Luigi, for their endless love and for their thoughts, which follow me wherever I go. Finally, I thank my friends, who have been close to me also during my period abroad and with whom I shared moments of joy and several adventures. Thank you Alberto S., Alberto S., Stefano, Riccardo, Giuseppe, Nicola, Nicolò, Alessandro, Edoardo, Giulio, Silvia, Erica, Rossana, Sara, and all who are not mentioned herer but shared part of this journey.

Michele Urbani
November 12, 2021
Trento, Italy

To my Family

Contents

List of publications	11
Nomenclature	13
1 Introduction	15
1.1 Scope and motivation	17
1.2 Goal and research questions	19
1.3 Outline of the thesis	21
2 Foundations and background	23
2.1 Methodological framework	23
2.1.1 Philosophical position of the research	25
2.1.2 Research ethics	27
2.2 Maintenance policies optimization	28
2.2.1 Reliability-centred maintenance	31
2.2.2 Dynamic grouping maintenance	35
2.2.3 Tools and techniques used in this research	41
2.2.4 Beyond Reliability-Centred Maintenance	45
2.3 Industry 4.0 and maintenance	47
2.3.1 On Industry 4.0	47
2.3.2 Industry 4.0-enabling technologies	50
2.3.3 Digital twins	51
2.3.4 Additive manufacturing for maintenance	52
3 Publications and contribution	55
3.1 Publication I	55
3.2 Publication II	56
3.3 Publication III (under review)	59
3.4 Publication IV	62
3.5 Publication V	65
3.6 Positioning of the research	67

3.7	Summary of publications	68
4	Discussion and conclusions	71
4.1	Discussion	71
4.1.1	Answering the research questions	71
4.1.2	Theoretical and practical implications	73
4.1.3	Limitations of the research	74
4.2	Prospective future research questions	75
4.3	Conclusions	77
	References	79

List of publications

This dissertation is based on the following papers and manuscripts. The rights have been granted by publishers to include the papers in the dissertation.

- I. Urbani, M., Petri, D., Collan, M., and Brunelli, M. (2020). “Maintenance-management in light of Manufacturing 4.0”. In: *Technical, Economic and Societal Effects of Manufacturing 4.0: Automation, Adaption and Manufacturing in Finland and Beyond*. Ed. by Mikael Collan and Karl-Erik Michelsen. Cham: Springer International Publishing, pp. 97–111.

Urbani is the primary author. Collan proposed the research topic, and Petri provided the material and the knowledge to write the contents. Urbani contributed to the design and general writing of the chapter supervised by Petri. Urbani carried out the literature study that provided adequate references for the topics treated in the chapter. Collan carried out the editing of the content, and Brunelli supervised the final revision of the artefact.

- II. Urbani, M., Brunelli, M., and Collan, M. (2020). “A comparison of maintenance policies for multi-component systems through discrete event simulation of faults”. In: *IEEE Access* 8, pp. 143654–143664.

Urbani is the primary author. Urbani proposed the research questions and carried out the literature research. Urbani designed and coded the numerical simulation experiments to test the maintenance policies. The design and general writing of the paper were conducted by Urbani with the supervision of Brunelli. Collan contributed to the general supervision and final editing of the manuscript.

- III. Urbani, M., Brunelli, M., and Punkka, A. (n.d.). “An approach for bi-objective maintenance scheduling on a networked system with limited resources”. In: *Manuscript*, 20 pages. Submitted 2021.

Urbani is the primary author. Urbani proposed the research topic and carried out the literature research to motivate the development of the proposed model. The proposition that motivates the grouping approach was developed and proved by Brunelli. Urbani carried out the development of the algorithmic procedure to solve the model

under the guidance of Brunelli. Urbani performed the implementation of the algorithm and numerical analysis. Urbani, Brunelli, and Punkka contributed to the design of the manuscript. Urbani and Brunelli wrote the manuscript. Comments to the results and conclusions are the outcome of the joint effort of Urbani, Brunelli, and Punkka.

- IV. Savolainen, J. and Urbani, M. (2021). “Maintenance optimization for a multi-unit system with digital twin simulation”. In: *Journal of Intelligent Manufacturing*. DOI: 10.1007/s10845-021-01740-z

Urbani is the secondary author. The research questions were formulated by Savolainen. Urbani carried out the literature study. Savolainen provided expertise in the mining industry. Urbani designed and coded the simulation-optimization experiment, to which the SD module written by Savolainen was connected. The design and general writing of the paper, exception made for the results regarding the SD module, was conducted by Urbani, whereas Savolainen edited the contents.

- V. Urbani, M. and Collan, M. (2020). “Additive manufacturing cases and a vision for a predictive analytics and additive manufacturing based maintenance business model”. In: *Technical, Economic and Societal Effects of Manufacturing 4.0: Automation, Adaption and Manufacturing in Finland and Beyond*. Ed. by Mikael Collan and Karl-Erik Michelsen. Cham: Springer International Publishing, pp. 131–148.

Urbani is the primary author. Collan proposed the research topic. Urbani interviewed the subject expert, prof. Paolo Bosetti from the University of Trento, and gathered the data about both the case study. Urbani contributed to the design and wrote sections one and two of the chapter. Collan wrote the third section of the chapter and carried out the editing and supervision of the whole manuscript.

Nomenclature

CBM	Condition Based Maintenance
CM	Corrective Maintenance
CPS	Cyber Physical Systems
CPPS	Cyber Physical Production Systems
DT	Digital Twin
GA	Genetic Algorithm
IoT	Internet of Things
LCC	Life Cycle Cost
MOO	Multi Objective Optimization
NED	Negative Economic Dependencies
PED	Positive Economic Dependencies
PHD	Prognostic Health Management
PM	Preventive Maintenance
RAMS	Reliability Availability Maintainability Safety
PHM	Prognostic Health Management

Chapter 1

Introduction

The reliability of products and services is fundamental to guaranteeing a steady and resilient growth of society. Despite the efforts of generations of researchers and practitioners, how to achieve and ensure the desired reliability level of an engineered product is still a challenge. Examples of bridges collapsing due to lack of maintenance, flights that must be interrupted due to engine failure, and electric car accidents due to fire ignition populate the news rather frequently. The failure of safety-critical systems is a threat not only to the safety of customers but also to the confidence of ordinary people in the power of science and engineering. With the recent advent of the fourth industrial revolution, society is reaffirming its confidence in technology to deliver economic growth and well-being. This paradigm shift is expected to deliver, among other things, extremely reliable products. However, the higher the number of parts that compose a technical system, the higher the probability is of one of the components failing. And since every engineered object is unreliable in the sense that it degrades with age and/or usage and ultimately fails (Ben-Daya, Kumar, and Murthy, 2016), ensuring the reliability of complex systems remains a major concern and a challenge for engineers.

A reliable product, or system, is the result of several decisions made during the design, production, and operational phases of the product lifecycle (Saaksvuori and Immonen, 2008). A lot can be done to improve reliability during the design phase when prior knowledge and learned lessons guide to achieve high reliability during the operational phase. However good the design is, the operative phase will be characterized by wear and tear phenomena; therefore, a product must be constantly monitored and maintained. Maintenance is indeed the key element to preserve reliability, and it has been defined by Pargar, Kauppila, and Kujala (2017) as “the work performed to keep a system in an appropriate condition and working order”. How to optimize maintenance the organization of complex technical systems is the objective of this research project.

Maintenance is part of the broader discipline called asset management (ISO 55000, 2014), which aims at aligning business objectives to asset performance. A common business objective concerns maximizing the return on investment (ROI) of a manufacturing system, whose performance is determined by its reliability, availability, maintainability, and safety (RAMS) characteristics. A good maintenance strategy steers decisions at an operative level to improve RAMS, whereas it strives to achieve high-level business objectives.

Physical assets degrade due to their use, which may yield unexpected failures and prolonged system downtimes. The latter can compromise the achievement of business objectives, and they may harm the health of workers and the environment. Such undesired events can be avoided by carrying out *preventive* maintenance, that is, by inspecting and restoring items to an acceptable reliability. This *modus operandi* is justified by the lower cost of preventive maintenance compared to corrective maintenance, which usually concerns a contingency situation where there is no choice but to pay a high cost to resume operations.

Due to the aleatory nature of degradation phenomena, drafting out a preventive maintenance strategy is challenging and requires a systematic approach. Information about the state of assets should be regularly gathered and stored in a maintenance management system; then, based on the available data, a decision-making model can be developed to help find the ideal preventive maintenance time and action. A peculiar hallmark of such a decision-making problem is the presence of uncertainty, which makes it challenging to find the trade-off between intervening early and waiting until failure precursors show up. The problem has been studied for decades in the scientific literature, and great progress was made thanks also to continuous technological development.

The advent of the fourth industrial revolution is setting a new pace in the research and development of solutions for preventive maintenance. The Internet of Things (IoT) is enabling real-time monitoring of assets at a fraction of the cost. Cyber-Physical Systems (CPS) allow a seamless connection of the physical and virtual worlds, thus making monitoring of machines and control of production accessible from everywhere. Powerful and computationally demanding simulation-optimization processes can benefit from cloud technology, which enables the execution of software on distributed infrastructures with high availability. Additive manufacturing (AM) technology is starting to mature for maintenance applications; hybrid machines integrating additive and subtractive manufacturing can perform repair tasks in a way that matches and exceeds the quality that can be reached manually. The application of the above-mentioned tools to maintenance management is relatively new to several industrial sectors. From an organizational viewpoint, there is the need for the research and development of new models and methodologies to achieve the seamless integration of operations and business objectives (GTAI, 2014), and

to increase the competitiveness of companies.

1.1 Scope and motivation

The scope of this work is to develop new models and methods for maintenance management optimization in light of the new technologies offered by the fourth industrial revolution. The goal of the developed models is to increase decision-maker awareness from an organizational viewpoint and minimize the cost of maintenance while delivering performance.

Figure 1.1 shows how maintenance policy optimization is found at the intersection of computer science, management science, operations research, and engineering. Knowledge of these four areas is required to realize the fourth industrial revolution. Reliability of hardware parts is a primary concern of engineering, both during the design phase and in control of the assets. Managing a portfolio of assets requires making rational decisions, which is the primary concern of management science. Operations research is called to provide the models that hold the information together and provide decision support. In turn, decision-making models rely on computer science artefacts to be efficiently solved; heuristic algorithms are an example of frequently chosen tools that provide good solutions to hard problems, and they are part of the focus of this research.

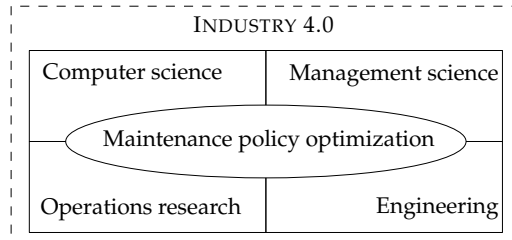


FIGURE 1.1: The map representing the fields covered in this thesis.

According to Rausand and Høyland (2003), there are two approaches to reliability analysis: the *structural* and the *actuarial* approach. The structural, also called physical, approach deals with the reliability analysis of structural elements, such as buildings and bridges. The strength $S(t)$ of an element and the applied loads $L(t)$ are modelled as random variables, which change as a function of the age t of the structure. The role of designers and system managers is to ensure that $\Pr(S(t) > L(t)) > \rho$, where ρ is the systems reliability threshold. The actuarial approach is followed in this thesis, whereby the information about the operating loads and the strength of components is summarized by the probability distribution $F(t)$ of the time to failure (Rausand and Høyland, 2003). No

explicit modelling of physical aspects is considered, and the focus is on the optimization of maintenance dates rather than the type of action to undertake.

The organization of maintenance is a function of the destination of a product (Ben-Daya, Kumar, and Murthy, 2016, p. 4), i.e., for retail, industrial, or defence applications. The range of models, techniques, and business objectives that apply to each group are different, and only industrial products are considered in the following. Industrial products can be standard or custom artefacts that are usually traded among companies, and which cover a role as parts of larger investment plans. A typical business objective of an industrial agent is to exploit the available assets to maximize the ROI, which also covers a fundamental role in maintenance optimization. In practical terms, high reliability and availability of the assets are required to maximize profitability, which, on the other hand, is threatened by the degradation of machines and the consequent need for maintenance. Moreover, industrial companies are characterized by the scarcity of resources, which limits both the production capacity and the possibility to carry out maintenance. How to balance these two factors to achieve profitability is one of the goals of this research.

Industrial products, or systems, are in turn made of four types of components, i.e., hardware, software, organizational, and human components (Zio, 2009). Despite the primary role of software in modern technological applications, the reliability of software tools is not investigated due to the substantial differences between reliability analysis methods of hardware parts; the same applies to humans. The organizational part is the focus of this thesis because it deals, among other things, with preventive maintenance of hardware components. There are two approaches to drafting out a preventive maintenance strategy of hardware parts. The one adopted in this thesis is Reliability Centred Maintenance (RCM) (Rausand and Vatn, 2008), which is complementary to the Risk-Based Maintenance approach. The goal is to develop novel reliability-based models for scheduling preventive maintenance activities, which can deliver better system performance in terms of RAMS. In particular, the contribution of this research is relative to group/block/cannibalization/opportunistic models (Cho and Parlar, 1991; Nicolai and Dekker, 2008), whereby the overall cost of maintenance can be minimized by jointly servicing components. The preventive maintenance problem can be studied at the element level, or at a system level; whereas preventive maintenance of single machines has been thoroughly studied in the past, the context of systems still offers opportunities to optimize maintenance (De Jonge and Scarf, 2020). An industrial system is an ensemble of parts connected in a networked fashion, which show a peculiar behaviour that is not observable when the parts are considered separately. The existence of such behaviour motivates the study of maintenance policy optimization for this specific application: Being able to exploit positive effects and to avoid the negative ones is a source of competitive advantage.

The competitiveness of a company in a global market is fundamental for survival and to operate profitably. Achieving competitiveness is a major reason for optimizing preventive maintenance scheduling. A single breakage event can lead to costly corrective maintenance, which in extreme cases can compromise years of future revenue, in addition to threatening human lives and the environment. Another major source of competitiveness is the digitalization of processes, which can improve efficiency and enhance the control of operations. Although digitalization seems to offer great upside potential for increasing competitiveness, it requires an investigation of how maintenance management models can be integrated with new digital technologies.

Finally, the undergoing technological shift is biased towards to an incremental change in the direction of a new economic model. To foster the sustainability of their business, several companies are redesigning their business model according to the principles proposed by the *circular* economy (Stahel, 2016). The latter encourages the reintroduction of goods in the production cycle through reuse, recycling, and remanufacturing when these are at the end of their operative life. The prospected change of economic paradigm makes possible a shift towards service-oriented businesses, according to which products used belongs to companies and customers purchase their use as a service, typically for a contracted period at a time. This paradigm shift has consequences for maintenance management: According to Stahel (2016), “services liberate users from the burden of ownership and maintenance and give them flexibility”. This usually means that companies selling products as a service must take care of any involved maintenance as a part of the service-contract. Maintenance becomes both a new source of revenues and a burden to be managed. The increasing attention towards the performance of products promotes the development of maintenance policies that can balance reliability, performance, and the availability of resources.

1.2 Goal and research questions

The goal of this research is to investigate preventive maintenance policies for complex systems, and to study how maintenance optimization can benefit from the technologies of the fourth industrial revolution. This goal is reached by finding answers to the following research questions.

Question 1 How is maintenance optimization evolving in light of the fourth industrial revolution? Preventive maintenance is already the standard in several industries. However, there are different approaches to preventive maintenance, which can rely more or less heavily on technology. Increasing the amount of technology means a great upside potential, not only for reliability and maintenance optimization, but also for several other

applications. On the other hand, there are downsides linked to the complexity of the adopted technological solutions, which may in turn be unreliable.

Question 2 How can we balance preventive maintenance and system performance? In complex systems, network effects may arise. How can these be exploited to optimize maintenance and system throughput simultaneously? Optimizing maintenance in complex systems is often a multi-objective problem (Zio, 2009). Reliability, availability, maintainability, and safety are four examples of optimization criteria, which might be conflicting to increasing system performance. Balancing productivity and the specific maintenance needs of multi-unit systems requires a holistic model, otherwise opportunities to carry out preventive maintenance could be missed, and poor performance periods could compromise the production targets.

Question 3 How can a maintenance management system be integrated into a Cyber-Physical System (CPS)? How can heterogeneous models be connected to improve a maintenance policy? What are the challenges and limitations of CPS? CPS are expected to gather, collect, and deliver data to/from different sources and stakeholders in real-time. A CPS aims at solving high-level tasks, e.g., to control production, to optimize energy consumption, to manage the warehouse, to implement condition-based maintenance, and to detect abnormal behaviours. These objectives are sometimes conflicting, and at other times cooperating. Controlling and balancing these objectives is a complex task, which can either lead to finding successful solutions and improving efficiency, or failing to reach the target business objective.

Question 4 How can additive manufacturing (AM) be exploited to improve preventive maintenance processes? What are the benefits and the drawbacks of using AM for preventive maintenance? And what AM-based business models can be envisioned in maintenance services? AM is commonly known for its ability to print objects with complex shapes, which could not be obtained through traditional subtractive manufacturing. However, early applications of AM also include the possibility of repairing and of refurbishing worn or damaged objects. Nowadays, such functionality has been extended and, thanks to the plethora of materials that is currently available to be printed and to the new printing technologies, AM is showing the potential to be used for repair of mechanical parts and for preventive interventions in the healthcare sector. The technological know-how required to use AM is still in the hands of a niche of technicians, whereby it is possible to imagine several ways to monetize such expertise.

1.3 Outline of the thesis

The research outcomes that have been published in international scientific journals, and in the book by Collan and Michelsen (2020) are presented in the following. Figure 1.2 shows the contents of this thesis and how they are connected.

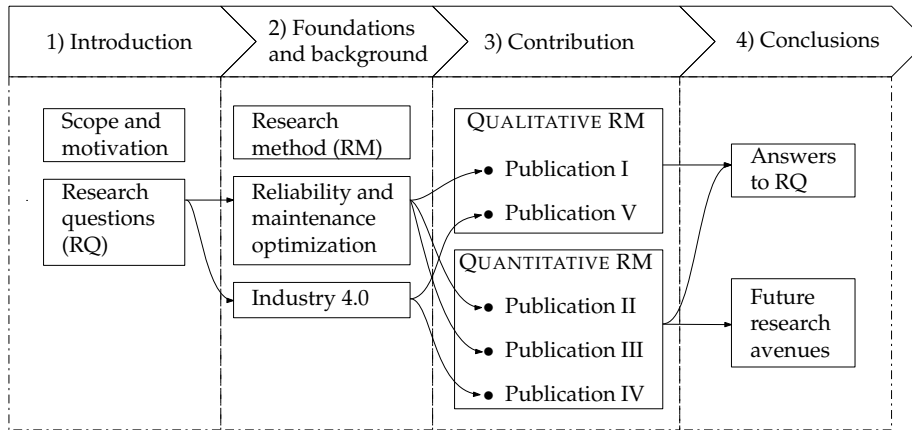


FIGURE 1.2: Contents of this thesis.

Chapter 1 introduces the reader to the scope and motivations of the doctoral project, and to the research questions. Chapter 2 begins by introducing the philosophical position or the view of the world the thesis has and defining the ethical position of the author. Then, the fundamental notions on reliability and maintenance strategies are presented, followed by an introduction to the central concepts that characterize the fourth industrial revolution. Furthermore, the implications of the latter on maintenance management are introduced and discussed. Chapter 3 briefly lists the contribution of the published papers. Finally, the research questions are answered and the results are discussed in Chapter 4; the thesis ends with a section about future research avenues and conclusions are laid out.

The outcomes of this research target different types of readers. Figure 1.3 shows a map of the publications, where these are positioned according to the intended audience, and according to the relevance to the fourth industrial revolution's technologies. The journal papers Publication II and Publication IV, and the manuscript Publication III are intended for a technical audience, i.e., researchers or practitioners who work in the field of maintenance optimization. The book chapters Publication I and Publication V are less technical, and can be easily read by undergraduate and graduate students, as well as non-technical readers. Publications I, II, and III propose a “traditional approach” to maintenance in

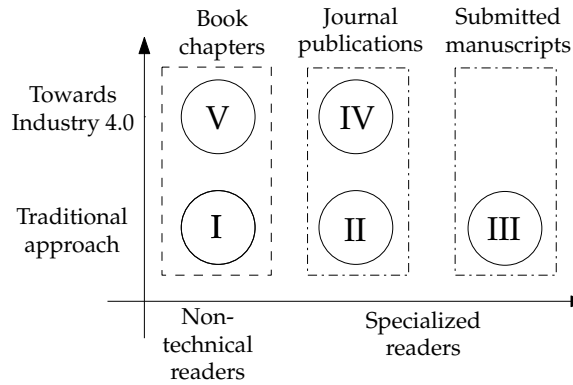


FIGURE 1.3: The research outcomes are represented within a conceptual map, and grouped according to the development approach. I) Publication I, II) Publication II, III) Publication III, IV) Publication IV, V) Publication V.

the sense that they do not deal explicitly with technologies of the fourth industrial revolution, but rather that they propose operations research models. Publications VI and V concern the use of cyber physical systems in maintenance policies optimization and the use of additive manufacturing for maintenance efficiency respectively. Therefore, these contributions are labelled as “towards Industry 4.0”.

Chapter 2

Foundations and background

In this chapter, the theoretical foundations of the work are laid out. The research methods, the philosophical position of this thesis, and the ethical position of the author are briefly presented in Section 2.1. The fundamental concepts about maintenance management and maintenance policy optimization are summarized in Section 2.2; the latter is also an analysis of the most relevant literature on the topic, and it identifies the research gaps that this thesis is going to address. Finally, Section 2.3 introduces how the field of maintenance can benefit from the fourth industrial revolution, and which are the main technologies that are enabling this transition.

2.1 Methodological framework

In the field of engineering management and in the context of this research, scientific investigation is a problem-solving task that concerns different aspects of science and several activities that connect them. Mitroff et al. (1974) proposed a systemic view of the scientific activity, which may eventually fit the research activities that were carried out during this research. Figure 2.1 introduces Mitroff et al.'s system view of the scientific activity: Science is seen as a system, within which four sub-systems can be identified—i.e., “Reality”, “Conceptual model”, “Scientific model”, and “Solution”. The cycles that can be realized by moving from one circle to another identify different ways to carry out a scientific problem-solving process; that is, they represent a solution to an identified real-world problematic issue. A scientific investigation can involve any of the activities and sub-systems in Figure 2.1, and there is no univocal start or end point. The choice of where to start and where to end is relative to the boundary conditions of the problem and the psychology of the investigator. A researcher is free to move among, or to stop on, any of the circles in the diagram as long as this activity increases the awareness of the problem, or it allows learning more about the problem, or it helps to produce an “artefact ” that solves a real-world problematic issue.

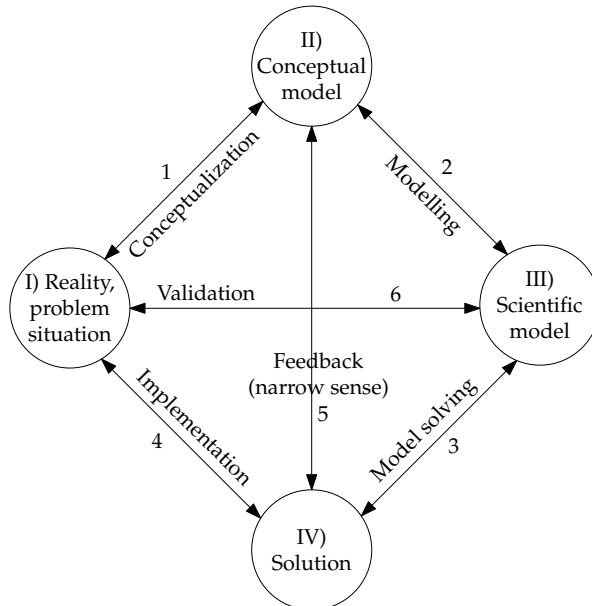


FIGURE 2.1: The systemic approach to problem-solving proposed by Mitroff et al. (1974).

The sub-system “Reality” represents the real world, where a problematic situation can be identified and can trigger a scientific activity. “Reality” can also be the arrival point of scientific activity, whereby the focus is commonly on validation of a “Scientific model”, that is, on the ability of the model to produce a usable and effective solution to the real-world problem. The “Conceptual model” aims at providing a conceptual description of the problem to be solved, and to set out the level of detail that is adopted; the field variables and the constraints of the problem are also defined. Starting from a real situation, the conceptual model can be drafted and it provides a natural starting point for the modelling process, which in turn contributes to the creation of a scientific model of the problem. The “Scientific model” is a formal description, usually based on mathematics, that is used in OR to represent a problem. Three arrows depart from the “Scientific model” in Figure 2.1; firstly, the model can be validated; secondly, the model can be “solved”, e.g., by applying an algorithmic procedure that produces a solution to the problem; thirdly, the scientific model can be used to refine the conceptual model through further modelling activities. Finally, starting from the “Solution”, one may feedback to the conceptual model to modify or refine it; alternatively, a solution can be implemented to produce a change in the real world. The implementation process shows how the activities and the processes that have been presented separately are in fact interrelated: It is misleading and false

to limit implementation to path 4) in Figure 2.1, because the difficulties found during the implementation might be the result of poor conceptualization, modelling, or model solving, just to mention a few.

2.1.1 Philosophical position of the research

Since the research outcomes of this thesis regard the proposal of novel models, methodologies, and (limited) theoretical contributions, it is important to discuss the philosophical foundation of the work. Models and methodologies can in turn be thought of as parts of a theory because they underpin the thesis of a theory or they are used for validation. Said with the words of Weber (2003), “a theory is an *account* that is intended to explain or predict some *phenomena* that we perceive in the world.” Assuming that the world is made of artefacts and that artefacts have properties, the set of properties of an artefact are its *state* (Weber, 2003). The state of an artefact may change at discrete points in time called *events*, and both states and events are properties of an artefact in that they “belong to” a thing. Phenomena are both the states of artefacts or the events that may occur to artefacts. When a theory is built, it attempts to connect two or more phenomena through a (set of) statement(s); in other words, a theory is the articulation of a *law* that describes or predicts how the components of an artefact are related.

The focus of scientists is often on the predictive power of a specific theory; that is, a theory is reliable as long as it can generalize on a large number of similar phenomena. How a theory can be validated is a long-debated topic in the philosophy of science (Smith, 2003), which for the sake of brevity is not discussed here. To test their models and methodologies, and hence their theories, researchers in OR typically make use of simulation tools. Simulations turn out to be particularly useful when a general statement needs to be tested, but observations of the empirical phenomena are limited. Since recognizing that simulations should be validated is akin to state that simulation models are similar to miniature scientific theories (Kleindorfer, O’Neill, and Ganeshan, 1998), what is relevant in the context of this thesis is *how can simulations be validated?* How one can infer that the proposed model captures the essential structure of the observed phenomena is in turn a debated topic. The goal is to develop “defensible decision models” (Kleindorfer, O’Neill, and Ganeshan, 1998) rather than to validate simulation models according to the well-known and opposite philosophical traditions of *empiricism* and *rationalism*.

Empiricism and rationalism are two *foundationalist* positions (Kleindorfer, O’Neill, and Ganeshan, 1998). A foundationalist believes that a model or a theory should find a basis either in direct experience (empiricism) or through self-evident ideas (rationalism). For a rigorous foundationalist, the validation process must be carried out until a foundation, i.e., a set of elementary propositions, cannot be stated. However, practitioners and

academics implicitly recognize that the foundationalist approach often fails as a validation method in the everyday use of simulations. Conversely to foundationalist positions, *anti-foundationalists* believe that if no grounds for a theory can be found, judgement and decision-making cannot be avoided. According to Kuhn (2012, p. 199), values such as fruitfulness and consistency of a theory or a model should be involved in the process of determining its adequacy. Involving values in the validation process means that there must be a recognized basis of common values; however, the latter cannot be easily established and it may require us to debate what this common basis is. The validation of theories through a common basis of values is known as *objectivism*. An objectivist believes that the validation process can be separated from the model builder, and that validation is an algorithmic procedure that is not open to debate. Since objectivism appeals to some external principles, it holds something of the foundationalist position, in that it seeks a common evaluation framework. Conversely to objectivism, *relativism* claims that a model cannot be separated from its builder and the context, and that model validation is a matter of opinion. According to the relativist position, a model is equally valid or invalid depending on the opinion of its stakeholders, and its adequacy is established through a dialogue between model builders and other model stakeholders. A model builder cannot carry out the validation process alone unless they are also the user of the model; the communication and discussion with the client are fundamental to validate and to assign credibility to a model.

The modern debate about validation in the philosophy of science evolved far from *either/or* positions between foundationalism and anti-foundationalism. Several authors agree that model-builders should strive for model credibility and that it should be less of a concern which of the two positions is embraced, as long as model credibility is reasonably increased. The kind of activity carried out in this thesis is regarded as objectivist in that the degree of adherence to commonly recognized concepts is used to validate the proposed theory—e.g., the concepts of reliability, profitability, and availability. However, the validation process of the proposed models was also influenced by the peer review process, which can be regarded as a relativist type of activity. To some extent, a model is credible as long as it exceeds the review process, which represents a form of social acceptance and it is therefore a purely relativist position.

The kind of philosophical activity carried out in this thesis lies in between the objectivist and the relativist positions. The opinion of the author is that as the validation process is based on judgements and decision-making, the ethics of the model builder must be discussed in the validation process. In an anti-foundationalist setting, the validation problem can be converted into an ethical problem, where the model builder and its stakeholders are called to warrant the credibility of the proposed theory (Kleindorfer, O'Neill, and Ganeshan, 1998).

2.1.2 Research ethics

Operations research (OR) concerns the use of mathematics to make decisions that have implications for reality. Whenever these decisions impact the lives of other individuals, or on society and the environment at large, they involve ethical judgements.

The role of ethics in OR has long been debated among scholars and practitioners and the development trends from 1966 to 2009 have been reviewed by Wenstøp (2010). The definition of ethics is not unique among operations researchers and three ethical categories are identified, i.e., virtue ethics, duty ethics, and consequentialism. Virtue ethics deals with the moral character of the agents, who value actions according to their intent; for instance, “to help the others” is a benevolent and charitable activity for virtue ethics. Duty ethics adopts a normative approach, whereby there are norms and duties to be respected; to act according to duty ethics means following a norm, e.g., “do unto others as you would have them do unto you”. Finally, according to consequence ethics, actions can have good or bad effects and an ethical behaviour pursues actions with good effect.

The debate about ethics and OR began with discussing the relevance of ethics in OR, which was regarded as science and as such free from values. It was soon recognized that, since the final goal of OR is to support the decision-making process, ethics is relevant to OR. Recently, the debate focused on the creation and the role of research ethics committees (White, 2009); on the responsibility of OR, and the role of sharing and co-operation (Gallo, 2004); and on responsibility and sustainable development (Brans and Kunsch, 2010).

The application of OR to maintenance optimization and risk management has clear ethical implications. A peculiar hallmark of decision-making in risk management is the presence of *uncertainty*; that is, the decision outcome could not be known a-priori, and thus could lead to undesirable effects. The exposure of human beings to risk due to decisions made by others suggests the adoption of deontological and consequentialist theories.

The deontological approach is part of duty ethics, whereby actions are permitted or forbidden up-front. A deontological view does not care about the consequences of an action and it rules out whether an action is good or bad according to a norm. According to deontological theory, any exposure of a human to a risk that may harm the personal or societal benefit is wrong. Moreover, the translation of human values into monetary value that is often used in risk-informed decision-making models is not acceptable from the deontological point of view. Only if the stakeholders of the decision-making process are willing to be exposed to a risk can deontology accept the use of a person as a means to achievement of the benefit of another entity (Ersdal and Aven, 2008). A company should act deontologically concerning to its employees and stakeholders. Assuming that zero-risk

work environments do not exist, it should be the aim of any company to reduce the risks for its workers to as low as reasonably possible (ALARP) level. The ALARP principle should be part of the deontology of a company, i.e., it is good up-front to lower the risks caused by the working condition to human stakeholders and the environment. This view is following what Brans and Kunsch (2010) claim.

In practical terms, however, the utilitarian approach could be preferred. *Utilitarianism* (Mill, 1998) is part of the consequentialist theories and it “regards an action as good if the action yields value in form of pleasure to humans, and right if the action yields the greatest net value for the society” (Ersdal and Aven, 2008). The assumptions about the possibility and the effectiveness of the utility approach are quite strong, and to make it operational is difficult. Decision aid models, such as the cost-benefit analysis, can help an agent to make risk-informed decisions (Ersdal and Aven, 2008) in the sense that 1) a set of future consequences can be identified, 2) a probability can be attached to each of them, and 3) the lowest risk scenario can be actuated. Although the future outcomes of an action can be described by a model, this does not provide hard decisions but only a decision aid. The decision remains subjective and it is demanded of the decision-makers.

2.2 Maintenance policies optimization

Modern maintenance management is the result of almost a century of development in the management of industrial assets. The research outcomes of this thesis are founded on the long tradition in research on maintenance that puts reliability and maintainability at the centre of maintenance optimization. To understand the adopted approach, and to properly introduce the problem of maintenance optimization, the fundamental concepts in maintenance management and their evolution are presented in the following.

Reliability is the “ability of an item to perform a required function, under given environmental and operational conditions and for a stated period of time” (ISO 8402, 1986). An item can be a component, a sub-system, or a system that is designed to perform one or more functions. If the function of a component is not specified, its reliability and maintainability cannot be measured (Rausand and Høyland, 2003). On the other hand, maintainability is the “ability of an item, under stated conditions of use, to be retained in, or restored to, a state in which it can perform its required functions when maintenance is performed under stated conditions and using prescribed procedures and resources” (BS 4778, 1991).

Maintenance is in turn the practical declination of maintainability. The origin of the word maintenance dates back to the year 1369 when the French word *maintenir* was used with

the meaning of “bearing”. A few years later, in 1389, there is a clue that the word maintenance indicated “the action of providing a person with the necessity of life”; in 1413, the word maintenance indicated the “action of upholding or keeping in being”, which resembles the meaning that it holds today. According to the IEC 30600 (1992), maintenance is the “set of actions that ensure the ability to maintain equipment or structures in, or restore them to, the functional state required by the purpose for which they were conceived”.

Up to the 1940s, the most widespread and almost unique maintenance policy was the *run to failure* policy: This consists of running machines until their failure makes them unavailable, then performing corrective maintenance (CM). A change of pace occurred in the 1950s, when OR models spread to the industry from the field of defence, where they were largely used during World War II. An ever-increasing number of models for the evaluation of preventive maintenance (PM) policies were developed and deployed for single components. Since the 1970s, the impact of maintenance on business objectives was more commonly considered: The Life Cycle Costing (LCC) approach started to take hold and it allowed the integration of financial aspects into maintenance models, thus filling the gap between reliability models of single components and their maintainability. Later, in the 1990s, the spread of microprocessor- and computer-based instrumentation for monitoring of machines allowed the development of the so-called *condition-based* maintenance (CBM), which aims at reducing (or even eliminating) unnecessary interventions by doing maintenance on-demand. Since the 2000s, CBM was further developed into prognostics and health management (PHM), which is a proactive approach striving to foresee the future maintenance needs of a component.

Maintenance actions, costs, and approaches

Maintenance interventions present a twofold nature: That is, an intervention can be corrective (CM) or preventive (PM) depending on whether it is carried out before or after a component fails. CM actions consist of the repair or replacement of components, and they are usually costly due to i) the potential consequences on the safety of the system’s stakeholders, ii) the creation of waste material, and iii) the high cost of missed production. A CM action may need to be carried out immediately or it can be deferred, if system operation is not compromised. Conversely, PM actions aim to be proactive to failure events, which means intervening before components fail and to possibly restore them to an “as good as new state”. The rationale behind PM can be time-based, whereby actions are scheduled at specific intervals, condition-based, or *on-demand*, and opportunistic, namely an unforeseen intervention is exploited to carry out PM on several items jointly. Examples of PM actions range from visual inspections of machines to lubrication of moving parts, or from the replacement of worn parts to the overhaul of turbine blades in an aircraft engine. A PM action is usually cheaper than a CM action on the same component, therefore

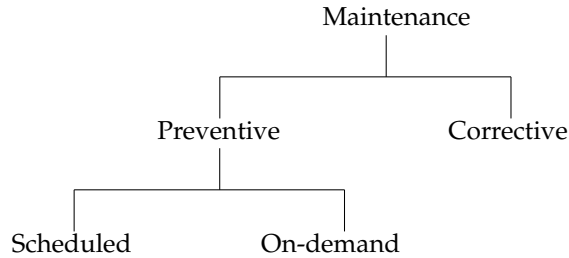


FIGURE 2.2: A schematic representation of maintenance approaches.

the objective of a maintenance manager should be that of avoiding CM in favour of the less expensive PM. If ensuring reliability is costly, not having reliability is even costlier. Figure 2.2 shows a simplified map of the maintenance approaches mentioned above.

Maintenance costs can be divided into two main categories. The first is that of *direct costs*, which are deterministically known and consist of direct cash disbursements. Examples of direct costs are the cost of labour, the cost of material, the cost of spare parts, the cost of contractors, and the cost of infrastructures and related tax. Direct costs may not be known in advance, but they can always be known ex-post. On the other hand, there are *indirect costs*, such as the costs associated with the failure of components, or the cost of unavailability (or downtime) of a system. These include, e.g., loss of revenue, the cost of accidents, and insurance policies; they are unknown and they often have to be estimated, thus leaving room for subjective judgements. Because of the convenience of PM maintenance over CM and of the uncertainty connected to indirect costs, the selection of the optimal maintenance approach is the subject of a lively debate among scholars and practitioners.

Maintenance approaches are in turn corrective, preventive, or a mix of the two. The corrective approach par excellence is the already-mentioned run to failure approach and it presents only little variations, whereas the range of PM approaches is broader. The goal of preventive approaches is to minimize reliability, availability, maintainability, and safety objectives and the life cycle cost of the system. The three best-known PM approaches are reliability-centred maintenance (RCM), risk-based maintenance (RBM), and total productive maintenance (TPM).

TPM is a Japanese-born approach that aims at maximizing equipment effectiveness. According to TPM, maintenance and production are organized jointly, therefore not only downtimes are minimized, but also equipment utilization is maximized. TPM's most peculiar hallmark is likely to be that every employee is involved in continuous improvement

processes, both vertically, i.e., from top managers to workers on the floor, and horizontally, that is among different company's departments. The work is carried out by small groups of employees in charge of specific activities and it requires a high level of motivation and engagement of workers. Because of this, TPM was successfully adopted by several Japanese manufacturing industries, but it is less common outside Japan. To further deepen the topic of TPM, the interested reader can refer to the books by Wireman (2004) and Nakajima (1988).

The main objective of RBM is to quantify and reduce the risk that may originate from failure consequences to acceptable levels, by implementing corrective or preventive actions. The three main steps of the RBM approach are 1) accident scenario S identification, 2) failure probability p assessment, and 3) evaluation of the consequences x . Then, a risk R_i can be defined by the tuple $R_i = \{S_i, p_i, x_i\}$ (Aven, 2012), and the identified risks can be ranked and compared. The expected practical result is that components yielding a high risk are to be inspected and maintained more frequently. Common techniques of analysis in RBM are the well-known Failure Mode Effect Analysis (FMEA) and Failure Mode Effect and Criticality Analysis (FMECA) (Rausand and Høyland, 2003, p. 88), hazard analysis, and the HAZard and OPerability (HAZOP) (Zio, 2007, p. 19) analysis. The book by Zio (2007) provides an introduction to the previously mentioned techniques.

The research work carried out in this thesis has been developed according to the RCM setting, which is presented in depth in the following.

2.2.1 Reliability-centred maintenance

Reliability-centred maintenance is a methodological approach to maintenance planning, whose aim is to maintain the system function at the minimum expenditure of resources (Ben-Daya, Kumar, and Murthy, 2016). The RCM approach was chosen over the RBM and TPM approaches because it focuses on drafting a maintenance schedule, the optimization of which is the main goal of the research outcomes presented in the next chapter.

The RCM methodology foresees the following steps. Firstly, *initiation and planning* are carried out and the system, sub-system, or components that are the subject of the analysis are identified. Then, a *functional failure analysis* identifies a set of Functional Significant Items (FSI) that are critical to the system operation and the related maintenance costs. Several techniques for functional failure analysis are also common to the RBM approach; three examples of formal approaches to identify FSIs are Fault Trees (Bedford and Cooke, 2001, p. 99), FMECA, and Reliability Block Diagrams (Rausand and Høyland, 2003, p. 118). Commonly-used practical examples of FSIs are the delivery of a flow of water to a reactor, the containment of a fluid within a tank, or the connection of a pump to a system of pipes. The following step consists of *consequences evaluation*, whereby the severity of

unforeseen failures is defined through expert judgements elicitation and cost estimation. The severity of items failure is then used to select the most effective maintenance actions to both criticality and cost minimization. Once the previous steps have been addressed, the *implementation* phase can begin: Costs and benefits of different portfolios of maintenance actions are traded-off, and a schedule of maintenance activities can be drafted. The essence of such a techno-economical trade-off is summarized by the plot in Figure 2.3, where the cost of preventive and corrective maintenance is plotted as a function of the preventive maintenance frequency. By balancing the amount of preventive maintenance,

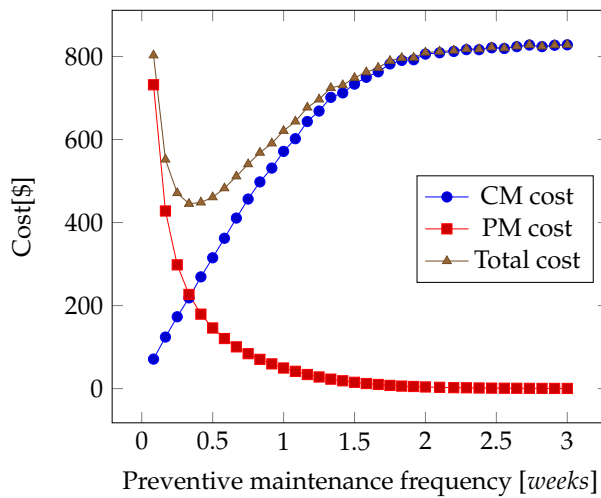


FIGURE 2.3: The cost of CM, PM, and the total maintenance cost as a function of the maintenance frequency. Reprinted from Publication IV.

which is less expensive but causes downtimes more often, and the amount of corrective maintenance, which is costlier and might have catastrophic consequences, it is possible to achieve the minimum of the total cost curve in Figure 2.3; reaching the minimum of the brown curve should be the objective of maintenance managers. Concurrently to the previous steps, the effectiveness of maintenance interventions is measured and data are gathered for *continuous improvement* purposes and to control system performance.

The total cost of running a technical system can be evaluated using the Life Cycle Cost (LCC) model (Ben-Daya, Kumar, and Murthy, 2016, p. 506), which accrues for the cost of the asset, the cost of spare parts, the cost of work, and possibly indirect costs such as the cost of missed production and waste material. The philosophy behind the LCC model is that not only should the cost of the single maintenance event be minimized, but the whole life of an asset is considered and the cost of all maintenance events is minimized

overall. Using the information resulting from the LCC model, it is possible to optimize the time intervals between interventions, refine the design of the system, estimate the long-term capital requirements, and improve the whole RCM process using field data. The manufacturer of a machine, or a productive system, may use the LCC information of the former version of a machine to improve the LCC of the next version during the design phase, which means to design for lower LCC; from the customer viewpoint, the LCC model is instead focused on optimizing maintenance costs and capital expenditure. From a practical perspective, a limitation of the LCC model concerns the estimation of the indirect costs, which are a major cost item and are subject to uncertainty.

The RCM framework also presents some limitations to be aware of. The first concerns the use of manufacturer-declared failure rate parameters, which are usually collected through test campaigns in a controlled environment. However, true operative conditions may be harsher, or milder, than test conditions and failure rates should be used carefully and possibly be re-parametrized using updated data. In the case of a new machine design, failure data might not be available at all, thus increasing the importance of monitoring and inspections. The access to field data and working conditions by manufacturers is also a major hurdle to improve the design of a machine and its reliability.

Modelling the failure behaviour of components is a fundamental task in reliability engineering. The most widespread approach is utilizing probability theory, which allows representing the uncertainty connected with aleatory degradation phenomena. When studying an aleatory phenomenon such as the failure of gear, usually the access to failure data is limited by the possibility to observe the phenomenon. To know the true failure behaviour of an item, one should theoretically observe an infinite number of failures, which is impossible. The solution is to observe a limited number of events and to approximate the true time to failure (TTF) distribution through a parametric equation. When following this approach, one should be aware of its limitations. The first concerns the selection of the right model, i.e., the equation that approximates more closely the distribution of available data. This kind of uncertainty is known as *epistemic uncertainty* and can be resolved by searching through the available equations. On the other hand, our knowledge of a failure phenomenon can be improved by observing a larger number of failures; the kind of uncertainty addressed with this approach is known as *aleatory uncertainty*.

Two common models for TTF representation are the exponential and the Weibull model. The exponential model $\lambda e^{-\lambda x}$ requires knowledge of only the failure rate λ of a component and the failure frequency $f(x)$, which is a function of the working time x . In turn, the probability that the item fails before x is described by the equation $F(x) = 1 - e^{-\lambda x}$. This model provides good accuracy in representing the TTF probability of electronic components, but it is not accurate for mechanical systems; in the latter case, the Weibull model

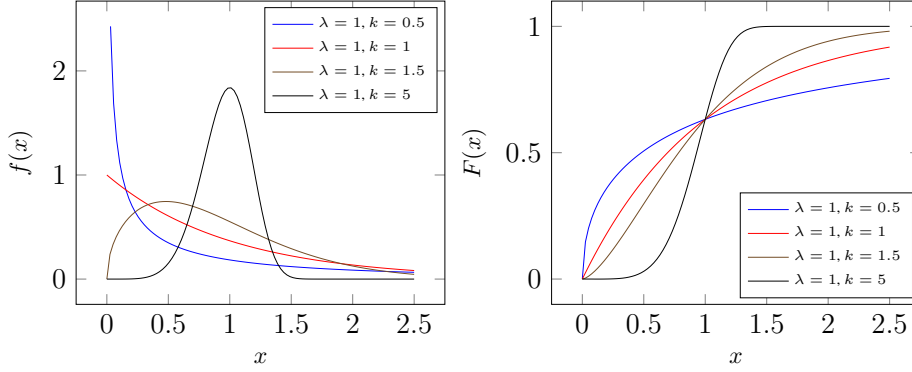


FIGURE 2.4: A few examples of Weibull density and cumulative functions.

The case $\lambda = 1$ and $k = 1$ is equivalent to the exponential model.

is known to be more representative. The Weibull distribution is characterized by two (seldom three) parameters λ and k , which are known as the “scale” and “shape” factors respectively. The higher representativeness capacity of the Weibull model comes at the expense of a higher number of parameters, the value of which need to be known. The probability that a component fails before x work time units according to Weibull is

$$F(x) = \begin{cases} 1 - e^{-(x/\lambda)^k} & x \geq 0 \\ 0 & x < 0 \end{cases} \quad (2.1)$$

and the failure frequency $f(x)$ is the derivative of $F(x)$ with respect to the work time x . One can observe that if $k = 1$, the Weibull model is equivalent to the exponential model. Figure 2.4 shows a few examples of Weibull frequency and probability distributions; the case $\lambda = 1$ and $k = 1$, i.e., the exponential model, can be compared to a few examples of Weibull distributions. The accuracy of model parameters is equally important to the choice of the right model and it is the starting point for the implementation of any RCM approach.

So far the RCM approach has been presented as a single-component approach; however, industrial systems are often ensembles of non-identical components that present specific maintenance needs. If systematically addressed, the possibility of maintaining multiple components jointly is an opportunity that can be exploited to save money and reduce the duration of maintenance interventions.

2.2.2 Dynamic grouping maintenance

Approaches aimed at maintaining multiple components simultaneously are known as *grouping* approaches, or grouping strategies. Grouping approaches aim at answering questions like how to group maintenance tasks, when to carry out maintenance on a group of components, how to handle opportunities for preventive maintenance when a sudden failure occurs, and how to dynamically optimize a group of activities? Although the dynamic grouping maintenance problem has been largely addressed in the literature, how to optimize maintenance by using dynamic information is still an open challenge.

The grouping problem is especially relevant in the context of *multi-unit* systems, which can in turn be divided into single- or multi-asset systems (Petchrompo and Parlikad, 2019). The two classes of systems differ in that multi-asset systems present an indistinct asset configuration and different maintenance tasks, which are instead clearly defined for single-asset systems. System reliability also affects the stakeholders differently in single- and multi-asset systems: In single-asset systems, the owner of the system and the user are the same entity, whereas in multi-asset systems reliability affects the user and the owner differently. An example of a multi-asset system is a portfolio of motorways. The company that owns the assets is interested in maintaining high asset availability, because that is the primary source of revenue; in this specific case, maintenance is both a burden that worsens user experience and a major item of expenditure. On the other hand, customers see the motorway as a service and they pay to travel on a safe and reliable piece of infrastructure.

Modelling multi-asset systems requires considering heterogeneous assets and the interests of different stakeholders; multi-asset models are indeed further classified into models for the management of fleets and portfolios of assets. Multi-asset management is an active area of scientific research; for a review of the literature we refer the interested reader to Petchrompo and Parlikad (2019). Single-asset systems are also referred to as *multi-component* systems—i.e., an array of elements that cannot be further decomposed into subsystems or components that are in turn target of maintenance (De Jonge and Scarf, 2020). Maintenance models for multi-component systems are an active research area and a great number of papers has been published on the topic; the results achieved by the scientific community have been reviewed several times in the past, see, e.g., Cho and Parlar (1991), Wang (2002), and De Jonge and Scarf (2020).

Maintenance models are seldom comprehensive enough to include all of the several aspects that influence the management of a real plant, and they usually focus on a limited number of issues that are typically the most critical from the point of view of safety, reliability, or profitability. Cho and Parlar (1991) and Nicolai and Dekker (2008) classify multi-component models into the following topical categories:

1. Machine interference/repair models that investigate the interference among machines in the same environment;
2. Group/block/cannibalization/opportunistic models that identify the components that should be preventively or correctively maintained to minimize the system LCC;
3. Inventory/maintenance models that account for joint maintenance and spare parts inventory planning;
4. Maintenance/replacement models that aim at helping the decision-maker to select the right maintenance action;
5. Inspection/maintenance models the goal of which is to determine the right interval of time between inspections.

The models developed in this thesis (Publications II and IV) mainly contribute to the class of group/block/cannibalization/opportunistic models. The latter hinges on the idea that system components are linked to each other through so-called *component dependencies*. Component dependencies occur when multiple units are considered as a whole and the system performance is influenced by the joint maintenance of these units. Dependencies of different types exist: Economic, stochastic, and structural dependencies were recognized by several authors in their reviews of the literature. In a chronological order, the reviews about multi-unit systems models that leverage on component dependencies are Cho and Parlar (1991), Dekker, Wildeman, and Duyn Schouten (1997), Wang (2002), and Nicolai and Dekker (2008). Recently, the resource dependencies were recognised as the fourth class of dependence by Olde Keizer, Flapper, and Teunter (2017), and they were also accepted in the later reviews of Petchrompo and Parlikad (2019) and De Jonge and Scarf (2020).

Economic dependencies

Economic dependencies can be positive or negative. A positive economic dependence (PED) occurs when the joint execution of more than one maintenance task leads to more efficient use of resources than the separate execution of such activities. PEDs take place because of the existence of economies of scale or downtime opportunities (Dekker, Wildeman, and Duyn Schouten, 1997). Preventive maintenance interventions commonly require some preliminary operations, which could be shared among several different activities. For instance, to access a remote part of a building it might be necessary to install a scaffold, independently of the number of parts that are accessed. Since the cost of the scaffold must be paid in any case, it might be convenient to carry out maintenance also on other parts that require the payment of the same setup cost; such occurrence is known as an *economy of scale*. On the other hand, the occurrence of a failure obliges the

plant manager to carry out corrective maintenance. The contingency situation justifies the payment of the setup cost, hence it triggers the opportunity to carry out other preventive maintenance actions.

A negative economic dependence (NED) occurs when the simultaneous execution of maintenance activities results in a higher cost than the execution of the activities separately. NEDs may be due to manpower restrictions, safety requirements, and redundancy or production losses.

Stochastic dependence

Stochastic dependence between the elements of a multi-component system is the ability of some components to influence the lifetime distribution of other components. Nicolai and Dekker (2008) proposed the following classification of stochastic dependencies: Type I failure of a component may cause both the failure of other components or of the whole system. Type II failure of a component can induce the failure of a second component with a given probability, whereas the failure of the second component act as a shock on the first component—i.e., the failure rate is influenced without causing instantaneous failure. Type III failure causes a shock to other components, affecting their failure rate.

A condition-based maintenance policy with stochastic dependencies and economic dependencies was proposed by Do, Scarf, and Iung (2015). The conditions to trigger maintenance were based on the current state of components, which were inspected only at specific points in time. If compared to other models with economic dependencies only, the main limitations were the number of considered components, which were only two, and the system configuration. Actual limits of stochastic dependence modelling are, first of all, the complexity (Van Horenbeek and Pintelon, 2013), which is a function of the number of components and their configuration, and also the difficulty to assess the effect of failures and degradation of one component on the others. These difficulties were partially overcome by Shi and Zeng (2016), who used stochastic filtering theory to make predictions on the remaining useful life of components in multi-component systems. Using PED and NED in addition to stochastic dependencies, Shi and Zeng's model opportunistically optimized the maintenance period and grouping structure of components. A promising development in modelling stochastic dependencies was provided by data-driven methodologies for remaining useful life estimation as showed by Peng, Dong, and Zuo (2010).

Structural dependence

Structural dependencies concern the influence of physical connections between components on maintenance. Originally, a structural dependence was intended to occur when “the disassembly sequence of a maintenance action influences the maintenance duration

and cost” (De Jonge and Scarf, 2020). A few examples are the removal of some system modules to access the damaged component in case of limited space to operate, e.g., in gearboxes (Dinh, Do, and Iung, 2020); alternatively, the precedence relations in the disassembly sequence could be the relevant aspect to model (Zhou et al., 2015; Dao and Zuo, 2017); by the same token, the inspection of a component can have consequences on the operation of the neighbouring components, such as in chemical plants. The above-mentioned examples are referred to as structural *technical dependencies* to distinguish them from structural *performance dependencies*.

A performance dependence regards the influence that the maintenance of a component shows on the performance of the system. This depends both on the performance of the single units, and on the configuration of the system. Some common system configurations are the serial, parallel, and k -out-of- N structures, the possible combinations of the previous, and arbitrary system structures. Consider, for example, a manufacturing system where machines and human operators are connected in an arbitrary structure. System units process the due jobs and deliver the end product to the next step of production; if a machine fails unexpectedly, it may cause starvation of downstream machines. The inability of a machine to operate gives rise to setup cost discounts and contiguity discounts (Papadakis and Kleindorfer, 2005)—i.e., other machines might be maintained without further affecting the cost of the intervention or the performance of the system. Further examples of models concerning performance dependencies have been reviewed by Olde Keizer, Flapper, and Teunter (2017), Petchrompo and Parlikad (2019), and De Jonge and Scarf (2020).

Resource dependence

Formally accepted as a separate class of dependence for the first time in the paper by Olde Keizer, Flapper, and Teunter (2017), resource dependencies were present in the literature since the early maintenance models. Resource dependencies deal with the limitation to carrying out maintenance according to the number of available resources. Modelling this type of dependence allows the solving of logistics problems of high practical importance. Several types of resources can be modelled: First of all, maintenance workers’ restrictions are considered. The limited availability of workers is an upper limit to the number of activities that can be carried out simultaneously; this aspect is particularly relevant to the dynamic grouping problem, and it requires scheduling ability to be solved. Similarly to workers, tools and equipment availability must be coordinated with the schedule of operations. Moreover, workers and tools present specific skills that make them eligible to be used for the maintenance of specific machines. Spare parts restrictions are considered as a resource dependence, when these are shared among multiple components; for instance, Nguyen, Do, and Grall (2017) proposed a predictive maintenance and inventory

strategy for multi-component systems using importance measures. Finally, maintenance programs are subject to budget restrictions, which may vary from the gross yearly budget to the monthly budget dedicated to each of the above-mentioned resources. The paper by Mild and Salo (2009) provides an example of dynamic budget allocation for infrastructure maintenance.

Grouping policies and opportunistic grouping

Component dependencies are system properties that can be effectively exploited to minimize the cost of maintenance and to avoid severe production losses. Understanding and systematically addressing component dependencies are the keys to achieving the overall minimum of maintenance costs; as long as such dependencies are not recognized, any attempt to optimize maintenance is doomed to fail. The mean by which dependencies are exploited is through the execution of multiple maintenance activities simultaneously, a practice also known as *grouping*. A grouping policy is a function that maps the system state to a set of maintenance actions to undertake. A policy can be a simple heuristic rule, or a complex mathematical function, e.g., a neural network, which can be respectively defined by a rule of thumb, or by algorithmic procedures. Usually, the main goal of a grouping policy is to provide a decision aid that encodes the available knowledge of component dependencies to improve the system performance.

To enhance the ability of a grouping policy to harness component dependencies, the *opportunistic grouping* of corrective and preventive interventions can be put in place. That is, when the failure of a component occurs unexpectedly, it may cause a (partial) system shutdown, which represents an opportunity to carry out preventive maintenance on those components that are forced to stop. The interested reader can refer to the paper by Geng, Azarian, and Pecht (2015) for a case study on opportunistic maintenance.

The dynamic grouping problem

The dynamic grouping problem was formulated for the first time by Wildeman, Dekker, and Smit (1997), who proposed a dynamic programming algorithm to group maintenance activities by leveraging the existence of economic dependence between components. In Wildeman et al.'s model, a set of maintenance activities, one per component, is preventively scheduled to draft a maintenance plan. The execution of an activity triggers the payment of a setup cost, which is shared in case multiple activities are carried out jointly. Therefore, the algorithm reschedules the activities in groups to save on setup costs. On the other hand, to shift activities from their ideal date, an incremental expected cost of corrective maintenance is potentially paid. This expected cost is a convex function of the shifting time from the ideal maintenance data, at which the cost is zero. Figure 2.5

depicts the expected cost of maintenance for two separate activities and for the same activities fulfilled in a group; the time that minimizes the expected cost of maintenance is the ideal preventive maintenance time for the group. Setup cost-saving and penalty costs

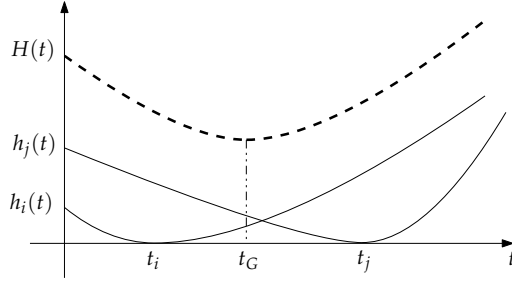


FIGURE 2.5: The incremental expected cost functions for single activities $h_i(t)$ and $h_j(t)$ as a function of the shift time t and the group expected cost of maintenance $H(t)$, which is the sum of $h_i(t)$ and $h_j(t)$ and presents a minimum in t_G .

payment is balanced by the algorithm during the search for the optimal grouping structure. When the earliest group of activities is carried out, new tasks are scheduled for the newly serviced components and the optimization is run again. Since the preventive maintenance schedule is repeatedly optimized leveraging new information, the approach is called the *rolling horizon* approach. In case of sudden failure, the penalty for shifting the corrective maintenance activity is arbitrarily set to a large value so that it would hardly be shifted, and the grouping structure is optimized. The *opportunistic* principle is seamlessly integrated with the grouping approach, thus allowing to optimize which components are worth preventive maintenance at any time.

Finding the optimal grouping for a set of components is a combinatorial problem, which was proved to be \mathcal{NP} -complete (Vu et al., 2014b; Vu et al., 2014a). Because of the problem complexity, exact methods provided few useful results and were often discarded in fa-vour of heuristic methods; for further insights about the heuristic techniques used in this research see Section 2.2.3.

Extensions and limitations of the dynamic grouping problem

The paper by Wildeman, Dekker, and Smit (1997) opened a new research avenue: The limitations left open by the original work were subsequently addressed and the model was integrated with new features.

The duration and type of maintenance activities were assumed to be negligible by Wildeman and colleagues, whereas it might be important in several real-world instances. Do Van et al. (2013) extended Wildeman et al.'s model considering multiple types of activities with different durations, which could be executed within time-limited opportunities.

Other extensions of Wildeman et al.'s model included modelling the degradation of components (Bouvard et al., 2011); the criticality of components and negative economic dependencies (Vu et al., 2014b; Vu et al., 2014a); multi-level condition-based maintenance on complex system structures (Nguyen, Do, and Grall, 2015); and the influence of humans on the quality of maintenance (Sheikhalishahi, Pintelon, and Azadeh, 2017). The applications in the real world of the above-mentioned models were very limited, with the only exception being Sheikhalishahi, Pintelon, and Azadeh (2017).

Finally, whereas the original model by Wildeman et al. exploited economic dependencies only, a maintenance model for complex systems might take advantage of other types of component dependencies. To extend Wildeman et al.'s model leveraging prognostic/predictive information, Van Horenbeek and Pintelon (2013) proposed a maintenance policy based on economic, stochastic, and structural dependencies that minimized the long-term mean cost per unit time. The proposed policy was compared to five preventive maintenance policies and showed the ability to produce significant cost savings. In addition, to consider economic and structural dependencies, Liang and Parlikad (2020) solved the dynamic grouping problem for multiple multi-component systems organized according to a networked structure. The previous models are two examples of efforts that have been made to integrate a grouping model with multiple component dependencies and complex system structures. We proposed a novel model to optimize maintenance leveraging both component dependencies and arbitrarily complex system structures. The approach proposed in Publication III addressed the dynamic grouping problem using multiple objectives, which helped to visualize the trade-off between the cost of maintenance and the performance of the system.

2.2.3 Tools and techniques used in this research

The publications Publication II, Publication III, and Publication IV rely on numerical tools that are widely used in operations research and computer science; these are briefly presented in the following.

Discrete event simulations and the Monte Carlo approach

In the context of maintenance, discrete event simulations (DES) can be used to replicate the behaviour of a system, thus predicting the performance of a maintenance policy.

According to Law (2014), DES “concerns the modelling of a system as it evolves by a representation in which the state variables change instantaneously at separate points in time”. There are three main hallmarks of DES models: These are *discrete*, that is the system-state changes at a countable number of points in time, and they can be described using rules that define how the system-state changes from one point in time to the next. Discrete models are alternative to continuous models, which capture the change of state variables at any point in time; an example of a continuous simulation tool is system dynamics (Forrester, 1994). Secondly, a DES model is *dynamic* because it can describe the change of a system with time. Thirdly, DES models are *stochastic*, because they are fed with random input components, which in turn produce a random output. In modelling a maintenance system, a stochastic representation is preferred to a deterministic one, because it better captures the probabilistic nature of failure phenomena, and it also provides a means to represent the uncertainty related to the adopted maintenance policy.

The Monte Carlo approach concerns the repetition of a large number of experiments with an uncertain outcome. The results of several trials, i.e., of an array of DES, can be analysed in an aggregated form to obtain a probabilistic representation of future scenarios. In their review, Alrabghi and Tiwari (2016) showed that simulations are largely used to improve and optimize maintenance in manufacturing systems, and that DES is the most widely used technique to model maintenance systems. The same authors proposed a novel approach to model maintenance of multi-component systems (Alrabghi and Tiwari, 2016), which was largely followed in developing the work in Publication II.

The main purpose of simulating a maintenance system is performance optimization. Alrabghi and Tiwari (2015) and Alrabghi and Tiwari (2016) found that manufacturing systems can be optimized using simulation-based approaches, which allow integrating maintenance models with production and spare parts models. DES also played an important role in Publication IV, where the goal was to maximize the profitability of the mining industry by leveraging the simulation of operations and maintenance. The experiment realized in the paper is carried out in a simulation-optimization setting (Fu, 2002), where a DES model is optimized using a heuristic algorithm.

The complexity of several optimization tasks in maintenance management, from the resolution of the grouping problem to simulation-optimization, yields very attractive ground for the use of heuristic approaches. Moreover, the ease of implementation and the flexibility of some heuristic algorithms make them usable in the rapid development of mock-up models.

Multi-objective optimization

The maintenance problem in complex technical systems typically requires the simultaneous consideration of multiple objectives. For instance, reliability, availability, and safety of a system are desirable characteristics that should be maximized, whereas risk and the LCC are to be minimized. However, the decisions that allow optimizing such characteristics are often conflicting, e.g., to carry out maintenance means a reduction in machine availability; in order to decrease risks, it might be necessary to invest in new equipment, or to implement solutions that require a reduction in productivity. In such a setting, it is not possible to optimize all the objectives simultaneously, because to increase some means to worsen at least one of the others. Moreover, the considered objectives might be incommensurable due to the different units of measurements that are used, and for this reason these problems cannot be transformed into a single-objective optimization problems. The maintenance of complex multi-component systems can be treated as a multi-objective optimization (MOO) problem (Zio, 2009). Some notions about MOO are summarized in the following.

Consider the following generic MOO problem

$$\begin{aligned} & \text{minimize} && (f_1(\mathbf{x}), \dots, f_p(\mathbf{x})) \\ & \text{subject to} && \mathbf{x} \in \mathcal{X} \end{aligned}$$

where $\mathbf{x} \in \mathcal{X}$ is the *decision variable* vector, and \mathcal{X} is the so-called decision variable space, which is a subset of \mathbb{R}^n . The functions $f_1(\mathbf{x}), \dots, f_p(\mathbf{x})$ are the $p \geq 2$ *objective functions* of the problem, also denoted by $\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_p(\mathbf{x}))^T$, and each $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$. The objectives are assumed to be minimized and, in order to avoid trivial solutions, it is also assumed that an \mathbf{x}^* that minimizes all objectives does not exist, i.e., that the objectives are (partly) conflicting. Due to the incommensurability and conflicts of problem criteria, it is possible to find not only one, but a set of equivalently good solutions, and the choice of the final solution is delegated to the decision-maker in any case.

A decision vector $\mathbf{x}^* \in \mathcal{X}$ is *Pareto optimal*, or *efficient*, if there is no other solution $\mathbf{x} \in \mathcal{X}$ such that $f_i(\mathbf{x}) \leq f_i(\mathbf{x}^*)$ for all $i = 1, \dots, p$ and $f_j(\mathbf{x}) < f_j(\mathbf{x}^*)$ for at least one objective f_j (Miettinen, 2012). Similarly, an *objective vector* $\mathbf{z}^* \in \mathcal{Z}$, where $\mathbf{z} = \mathbf{f}(\mathbf{x})$ and \mathcal{Z} is the *objective space*, “is Pareto optimal if there does not exist another objective vector $\mathbf{z} \in \mathcal{Z}$ such that $z_i \leq z_i^*$ for all $i = 1, \dots, p$ and $z_j < z_j^*$ for at least one index j ” (Miettinen, 2012). The set of Pareto optimal solutions is called *Pareto optimal set*, or *Pareto front*. The concept of Pareto optimality was largely used in Publication III, where a multi-objective genetic algorithm (GA) was employed to find an efficient frontier that was a close approximation of the true Pareto front. When the Pareto front (or a close

approximation to it) was found, the decision-maker is called to choose from the available options based on their experience.

Finding the Pareto optimal set requires the use of specifically designed techniques to deal with multiple objective problems. In practical MOO, it is not always possible to find the true Pareto optimal set, but a good enough approximation can usually be found. Multi-objective GAs have been shown to be an effective tool to find the efficient frontier for problems showing a non-linear, and non-con-vex, objective space (Konak, Coit, and Smith, 2006), similar to those tackled in this research. An example of the application of a multi-objective GA to maintenance policy optimization is provided by Hilber et al. (2007).

Exact and heuristic methods

Exact and heuristic methods were both used to tackle the optimization problems that were presented in Publications II, III, and IV. Several efforts were spent to design custom versions of heuristic algorithms that allowed complex combinatorial problems to be solved.

Specifically, the dynamic grouping problem was solved using a genetic algorithm (GA), both in its single- and multi-objective version. The GA belongs to the class of *evolutionary algorithms* (Goldberg and Holland, 1988) and it synthetically replicates an evolutionary process to select and improve a population of solutions. The information is stored in vectors x called *individuals*, which can be evaluated using a *fitness* function. The latter can be any kind of function that can be fed with x and returns a scalar value z in the case of a single-objective GA, or a vector z in case of a multi-objective GA. The main loop of a GA consists of 1) an evaluation step, during which a score is attributed to all individuals; then, 2) the most promising individuals are selected through a selection process, of which there exist of many kinds; and 3) the selected individuals are mutated to produce an offspring population. The algorithm iterates over the previous steps until a stopping criteria is not met.

The key difference between single- and multi-objective GAs regards the selection process: In the case of a single-objective, the fitness function returns a scalar value; therefore, the best individual is the one with the lowest score and the parent individuals can be selected accordingly. In the case of a multi-objective GA, the search procedure must ensure that the individuals improve both in terms of their criteria and that they are spread across the objective space. Multi-objective GAs have been shown to be effective in finding well-spread Pareto frontiers; a few examples of multi-objective GAs are those proposed by Fonseca and Fleming (1993) and Deb et al. (2002).

2.2.4 Beyond Reliability-Centred Maintenance

The RCM approach adopted in this research project is *time-based*, where-by the knowledge about the time to failure of an item is represented through a probability distribution. The probabilistic approach is justified by both the scarcity of failure data and the wide acceptance of such an approach in the scientific literature. An alternative to the time-based approach is *condition-based* maintenance (CBM), which tackles the maintenance problem by combining data-driven reliability models and sensor data gathered directly from the monitored system to design a PM policy (Alaswad and Xiang, 2017). The CBM approach is costlier than the time-based approach due to the need to continuously monitor the system, but it offers a high potential to avoid catastrophic failures, and the connected risks, and to intervene with PM only on-demand.

The CBM process can be summarized into three phases: data acquisition, feature extraction, and condition monitoring. Data acquisition relies on sensors for the acquisition of signals, which measure physical quantities such as temperature, humidity, speed, and pressure. The rise of the IoT contributed to speeding up the adoption of the monitoring solution and consequently spread CBM. However, the implementation of IoT devices alone is not enough to produce meaningful results and *feature extraction* is required to allow condition monitoring to be effective. Feature extraction deals with finding the set of signals that provide a correct system state representation under different working conditions. Data gathering and feature extraction are the starting point for carrying out the more complex activities involved in condition monitoring.

Condition monitoring in turn concerns two activities, i.e., fault detection and fault diagnosis. The former aims at identifying the presence of abnormal working conditions, which implies that the “normal” working condition is known; clearly, feature extraction plays a fundamental role in identifying the profiles that correspond to “normal”, “degraded”, and “faulty” states. Two examples of failure detection tools are the Auto-Associative Kernel Regression (AAKR) (Baraldi et al., 2015), which is used to reconstruct an observed signal according to the learned normal state, and the Principal Component Analysis (Baraldi et al., 2011). A limitation of CBM approaches is the need to determine a threshold over which the system is considered to be in an abnormal working condition; two statistical tests that may help with the threshold definition are Q Statistics and the Sequential Probability Ratio Test (Di Maio et al., 2013). Moreover, techniques akin to AAKR allow establishing which signal among those that are available is the likely cause of the abnormal behaviour; such a process is known as *fault diagnostics*. Fault diagnostics deals with identifying the nature, the extent, and the severity of a fault, that is, the identification of which one among the many possible components, or sensors, that are part of the system is the cause of the failure. Several statistical techniques can be used to carry out the task and

many of them are labelled as artificial intelligence (AI) techniques; several examples of AI applications for prognostics of rotating machines were reviewed by Liu et al. (2018). To summarize, failure detection works as an alarm, which is triggered in case there is an abnormal condition was detected; failure diagnostics is instead more like a doctor who drafts a diagnosis of the problem based on the state of the system. The reader may refer to Figure 2.6 for a conceptual map of all the concepts mentioned in Section 2.2 and to visualize how they are connected to Industry 4.0.

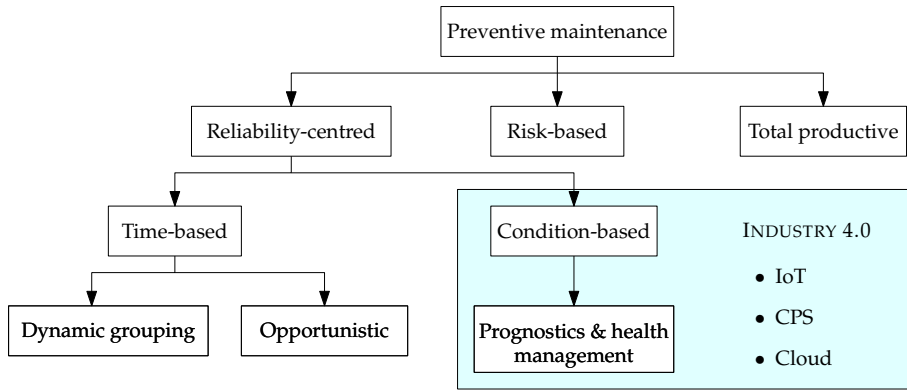


FIGURE 2.6: A schematic representation of preventive maintenance approaches.

The methodologies proposed so far are characterized by a retrospective approach to maintenance: Past data are exploited to create a benchmark for normal and abnormal conditions, and the observed signals are leveraged to determine whether or not there is a failure. This philosophy provides few insights into the remaining useful life of a component and the relative uncertainty. To this aim, the Prognostics & Health Management (PHM) (Kim, An, and Choi, 2017) approach was developed. The goal of PHM is to further reduce the cost of PM, and to increase safety and availability. The phases of the PHM process can be viewed as partly overlapped to the CBM process, to which the prognostics step is added. Prognostics is indeed a future-oriented PM approach, which is based on CBM and aims to predict how long it will take for a failure to reach a safety-critical state under the actual working conditions. ISO 13381-1 (2005) provides a more formal definition of prognostics, which is “an estimation of time to failure and risk for one or more existing and future failure modes”. The first industry to adopt the PHM approach was the aerospace industry, which presents a renowned need for safety and high maintenance-related costs. Subsequently, the PHM approach spread to several other industries, such as defence, civil infrastructure, electronics, manufacturing, and wind power; a review of PHM applications

up to 2012 was written by Sun et al. (2012). PHM presents a large potential to produce cost savings; as it was declared by five major companies that adopted this approach, the successful implementation of PHM solutions and advanced monitoring led to a saving of USD 885 million dollars (Kim, An, and Choi, 2017, p. 7).

The challenges linked to the implementation of PHM solutions are still several. Data management is one of these. Physical quantities are transformed into data by sensing devices, then data must be moved through the network, stored in high-capacity cloud infrastructures, processed and used to feed AI tools, and finally delivered to stakeholders. The creation of such a pipeline of data is a complex task, which requires involving experts with skills in specific fields, from software engineering to networking, and from business experts to maintenance managers. From a technological viewpoint, the advent of the fourth industrial revolution is a fruitful context within which to develop the full potential of advanced PM solutions.

2.3 Industry 4.0: new opportunities for maintenance optimization

The early 2010s were characterized by the occurrence of a major technological shift: The manufacturing industry started to experience a strong push towards the integration of physical and digital processes. The change of pace is taking place worldwide and it is known in Europe by the name “Industry 4.0”, which has been inherited from the German-born project INDUSTRIE 4.0 (GTAI, 2014). The technological shift that the developed countries are experiencing is acting as an enabler for the adoption of advanced maintenance management philosophies, such as PHM. In the following section, the main hallmarks of Industry 4.0 are outlined and the enabling technologies are presented.

2.3.1 On Industry 4.0

In the 18th century, the production of goods was still as slow and inefficient as it was in the Dark Ages due to the use of man- and animal-powered systems; craftsmanship was the key technology and industries could be compared to today’s craft workshops. At the end of the 18th century, the introduction of steam- and water-powered systems paved the way for the mechanization of manufacturing. The invention of the steam engine is recognized as one of the most important innovations in that the first industrial revolution ushered in the so-called Age of Steam. This technological shift enabled both an increase in productivity and an increasing size of machines.

At the end of the 19th century, the introduction of electricity led to the second industrial revolution. The mechanical energy generated by steam and water could be delivered to industries in the form of electrical energy, and then transformed into motion by the electrical engine. This technological shift made it possible to move factories far from the sources of water. The increase in productivity of manufacturing systems was such that it allowed a strong decrease in production-related costs, and the price of several products became affordable to large masses of the population; the mass consumption of goods characterized the so-called Age of Electricity.

Until the years following World War II, the number of electricity-based innovations grew at an ever-increasing pace and culminated in the introduction of microelectronics. The spread of microcontrollers for industrial applications enabled the rise of automation in machines, which became capable of performing several tasks without the supervision of human operators. The rise of automation was closely related to information and communication technologies (ICT), the advancement of which enabled the development of technologies like computer numerical control (CNC) machines and robots, computer-aided design (CAD), and computer-aided manufacturing (CAM) technologies. Due to the increased reliance of the industry on ICT, the third industrial revolution ushered in the Age of Information.

In Europe, the fourth industrial revolution began with the INDUSTRIE 4.0 project (GTAI, 2014) supported by the German government in 2010. The ideas contained in the German-born project found fertile ground in the rest of Europe as well, where they spread under the name Industry 4.0. However, the phenomenon took off almost simultaneously on a global scale, and although it was referred to by different names, the underlying ideas and concepts were the same. Factories should be transformed into integrated environments where the physical and the virtual world communicate seamlessly; the resulting ecosystem is the so-called “smart factory”—that is, a physical-virtual ecosystem where machines are interconnected and exchange information among each other and with human stakeholders. The adjective “smart” stems from blending different aspects of science, engineering, and business: The knowledge in electrical engineering, business administration, computer science, mechanical engineering, and business and information systems engineering are put together to enhance efficiency, competitiveness, and flexibility of companies (Xu, Xu, and Li, 2018). The model of industry proposed by Industry 4.0 hinges on new technologies, since it would not be possible to realize the physical-virtual integration without machines connected in a network, and data being accessible from everywhere.

In the Age of Information, mechanical systems were equipped with embedded electronics, which enabled automated production systems, and industrial processes benefited from

the development of ICT, which empowered machines with software. However, the elements of the industrial system were interacting weakly with each other and there was little integration with business processes. The aim of Industry 4.0 is to connect the machines in the workshop to each other and to the people, thus allowing real-time control and optimization of processes. Physical facilities are equipped with devices that can sense the surrounding environment and stream the data to the cloud. Machines are connected to the internet through wireless sensor networks, which allow data to travel to and from the workshop. The control decision can be performed remotely with the aid of a decision support system and sent to the machine after a simulation-optimization process showed the decision-maker what are the most likely failure scenarios. Such a connection of embedded system production technologies with the so-called *smart* production processes enables the change of paradigm that is brought forward by Industry 4.0 (GTAI, 2014).

A systemic analysis of the proposed model highlights a transformation from a centralized production logic to a decentralized setting, where processes are distributed between the workshop and the cloud. This is going to produce an augmented reality, richer in information, more efficient, and more engaging. A meaningful example is the mass customization of products, which is already a reality for several goods, from shoes to cars. Customers are involved in the design process through online platforms, where they can personalize the final product, and customers' preferences become part of the production process as they are key information for operations. From an operational viewpoint, processes are adapted to reduce a production batch to the size of one object without losing economic profitability. Another meaningful task of the production process runs in the cloud, where customer preferences are used in an aggregated form to optimize the production process, e.g., by improving the supply chain through forecasting of future needs, by optimizing the production schedules, and by offering insights on the product features to be further developed.

Data can also be used the other way around—i.e., from the production process to the end-user—to increase awareness about an ongoing process. In the transportation sector, real-time information is already exploited to inform users about delays or changes of schedule; similarly, the logistics industry relies on tracking technologies to inform its customers about the delivery status of parcels. By the same token, information about the advancement of processes can be delivered to different company departments, which can harness timely information to optimize inefficiencies. Focusing on performance is typical of service-oriented businesses, which sell products on a per-use basis. The Industry 4.0 philosophy fosters this approach, which is made possible by the new technological advancements.

2.3.2 Industry 4.0-enabling technologies

Three fundamental technologies enable the realization of Industry 4.0: Cyber-physical systems (CPS), the Internet of Things (IoT), and cloud computing.

The IoT concerns a dynamic ecosystem of interconnected devices that embed different sensing, radio communication, networking, and information processing technologies. The heterogeneity of the elements composing the IoT poses a challenge to the realization of an IoT infrastructure, and hence to the creation of the smart factory. To achieve the interoperability of devices, a Service-Oriented Architecture (SOA) is considered a good design approach (Xu, He, and Li, 2014). The most common models of SOA separate IoT devices into layers according to their concern; for instance, the International Telecommunication Union proposes a five-layer architecture that is composed of sensing, accessing, networking, middleware, and application layers (Xu, He, and Li, 2014).

On the sensing layer of an IoT ecosystem, there are the so-called “things”, i.e., the devices that transform physical quantities, such as temperature, speed, or humidity, into data, and that translate signals into actions. Subsequently, “things” send data to data warehouses through a Wireless Sensor Network (WSN) and then through the internet. A recently established trend foresees that sensors and actuators provide also computing power locally, thus allowing to implement AI algorithms “on the edge”.

Several SOA models present a service, or application, layer as the last layer of the architecture (Xu, He, and Li, 2014); here is where services exchange information, and where data are accessed and processed. The operations of the application layer are designed to be carried out in a highly distributed infrastructure, which is usually called the *cloud*. The cloud is made of both powerful computers for High-Performance Computing (HPC), and large Network Attached Storages (NAS) for data warehousing.

Several challenges must be addressed to both build the physical-virtual infrastructure and to protect IoT systems from external threats. From a strategic viewpoint, a standard protocol to enable interoperability, compatibility, and reliability of IoT systems on a global scale is still missing. The networking of devices poses some technical challenges too. For instance, the impossibility to test the software under all devices’ possible working conditions limits the reliability of the system, whereas the difficulty to realize predictable timing in a networked environment may make it difficult to rely on the IoT for safety-critical applications. Scalability is also a challenge due to authentication management, bandwidth requirements, and data processing and management. Last but not least, information security and privacy protection are difficult to guaranteed since constant research and development is required to combat an increasing number of threats.

CPS are a key technology for the realization of the smart factory—i.e., of an environment designed to achieve adaptability and learning characteristics, fault tolerance, and risk management of production processes (Xu, Xu, and Li, 2018). Moreover, CPS are expected to improve resource productivity and efficiency and to enable more flexible models of work organization. In comparison to a classic production system, a smart factory should present the ability to respond almost in real-time to quality issues and to optimize the use of resources.

On the other hand, the implementation of CPS presents some challenges. Derler, Lee, and Sangiovanni Vincentelli (2012) highlighted that designing a CPS “requires understanding the joint dynamics of computers, software, networks, and physical processes”. Since CPS are expected to react to multiple signals from sensors and to control multiple actuators concurrently and in real-time, they are required to show *concurrency* and *predictability* characteristics. According to Lee (2008), achieving such abilities is the main hurdle due to the philosophy of traditional programming languages. The current computing and networking abstractions are designed to match object-oriented and service-oriented architectures, rather than the physical dynamics, where “time cannot be abstracted away” (Lee, 2008). Furthermore, since the physical world itself is not entirely predictable, CPS should show robustness to unexpected conditions. System designers should strive for reliability and predictability of the elements that compose a system as long as this is technically possible and cost-effective; where this would be impossible, it is required to act at the higher level of abstraction.

2.3.3 Digital twins

Digital twins (DT) are an innovative technology that enables the realization of cyber-physical products or systems. The goal of a DT is to realize a seamless real-time connection between a physical object and its digital counterpart, a virtual model that exists in a digital environment, which can be used to deliver optimization and predictive abilities. The DT idea gained momentum and started to spread after the publication of the Modelling, Simulation, Information Technology & Processing Roadmap (Shafto et al., 2010) by the National Aeronautics and Space Administration (NASA) at the beginning of the 2010s. Shafto et al. defined a DT as “*an integrated multi-physics, multi-scale, probabilistic simulation of a vehicle or system that uses the best available physical models, sensor updates, fleet history, etc. to mirror the life of its flying twin. It is ultra-realistic and may consider one or more important and independent vehicle systems*” (Shafto et al., 2010). In their seminal review article, Negri, Fumagalli, and Macchi (2017) list several definitions of DT that were proposed in the literature and they show that researchers do not yet agree on a unique DT definition.

The early applications of the DT concept were in the field of aero-space, e.g., (Tuegel et al., 2011) and the main objectives included monitoring the health of an asset, performing predictive maintenance, checking mission requirements, and providing a more transparent life cycle-view of the asset (Cimino, Negri, and Fumagalli, 2019). Although the initial idea of DT was limited to single products, the DT model was soon extended to entire manufacturing systems, which are also referred to as Cyber-Physical Production Systems (CPPS) (Uhlemann, Lehmann, and Steinhilper, 2017).

According to Kritzinger et al., 2018, all DT are made of a physical and a digital part, and of a bi-directional communication process between the two. The actual physical asset is equipped with sensors and actuators, which allow data to be gathered about the state of machines and to actuate decisions respectively. The digital part plays the role of a decentralized brain for the physical one: Historical data about the working conditions of machines can be exploited together with real-time observations to run predictive maintenance algorithms, and to simulate forward the state of the equipment is a faster-than-real-time way to predict future events. Within the digital part, an ultra-high fidelity copy of the physical system usually exists and can be harnessed for control purposes (Wright and Davidson, 2020).

The advantages of adopting a DT approach for complex systems are several. Data collection issues can be improved through the use of IoT devices, predictive abilities can be achieved, and operations can be optimized by using predictive analytics. Furthermore, the possibility to predict future system downtimes helps improve resource management, production planning, and procurement. Although most of the attention in this thesis is dedicated to the operational phase of an engineering system, DTs can help to improve the whole product life cycle (Fei et al., 2018), and business innovation (Lim, Zheng, and Chen, 2019). For instance, maintainability issues can be resolved on the drawing board both by improving the design and by embedding IoT devices for prognostics and diagnostics.

2.3.4 Additive manufacturing for maintenance

Additive manufacturing (AM) technology consists of joining materials to make parts from three-dimensional (3D) Computer-Aided Design (CAD) systems, usually layer upon layer, as opposed to subtractive and formative manufacturing methodologies (ISO 52900, 2015). The production of an object through AM starts from a CAD model of the end product. Then, a virtual model is enriched with additional structures that allow parts to be printed without support, and the object is further decomposed into cross-sectional layers, which corresponds to one cycle of material deposition. After the printing process

ends, the support structures are removed, either mechanically or through a chemical process. The result is an object born of fewer restrictions to the *freedom of design* that can be delivered only by AM technology. An example of where AM differs from traditional manufacturing is that with AM, cavities can be built within traditionally solid shapes, which helps to save material and create lighter shapes. According to ISO 52900 (2015), there are seven groups of technologies in AM, which differ due to the used material, how the layers are created, and how the layers are bonded to each other. These technologies are binder jetting, direct energy deposition, material extrusion, material jetting, powder bed fusion, sheet lamination, and VAT photo-polymerization. The chosen AM technology influences the accuracy of the final part and the properties of the material: For instance, the thinner the layers, the higher the quality of the end product, and the longer the printing process. Moreover, the throughput of material and the processing and post-processing times vary according to the technology; such factors contribute to the profitability of the process, which should be carefully estimated and considered together with the cost of the machine.

The value of AM resides in the ability to realize shapes that could not be obtained through traditional forming and subtractive technologies. For instance, the subtractive and forming technologies cannot shape the internal regions of an object; this is instead possible with AM, which can produce objects with variable density. By the same token, different materials can be mixed to produce parts with variable material grades.

The combination of the additive and subtractive technologies into so-called *hybrid* machines enables the realization of efficient maintenance processes. For example, Publication V describe how surface cracks that occur on metallic dies can be refurbished thanks to hybrid AM machines. Metallic dies for metal shaping present peculiar shapes and are made of hard alloys, which are expensive and should be used sparingly. AM allows the use of lower quantities of raw materials and, thanks to the combination with advanced software tools (Perini, Bosetti, and Balci, 2020), it can automatize the repair of unique mechanical parts with no loss of functionality and in a way that is also economically viable. The logistics of maintenance could also benefit from the adoption of AM. The supply chain of repair parts can be transformed into the supply chain of materials to print repair parts; the benefits from this shift include the reduction of lead times and an increased flexibility of the refurbishment process, which can be applied to objects showing different shapes and volumes.

The integration of AM in maintenance processes also has serious implications for the business models for service provision (Weller, Kleer, and Piller, 2015). A company may have to choose between outsourcing and internalizing the maintenance process: The latter option may require investments for the adoption of the AM technology, which is expensive

both in terms of equipment and expertise. On the other hand, outsourcing maintenance provides flexibility and does not require the employment of costly resources necessary for AM. There is a plethora of available options for companies willing to adopt AM in their processes (Thomas, 2016). Savolainen and Collan (2020) made the effort to review the literature about business models involving AM by directly answering the question “*How does additive manufacturing technology change business models?*”

Chapter 3

Publications and contribution

The present chapter provides a summary of the research objectives and the main contributions of the publications resulting from this research project.

3.1 Publication I

The publication “*Maintenance-management in light of Manufacturing 4.0*” (Publication I) is a chapter in the book “*Technical, Economic and Societal Effects of Manufacturing 4.0*” edited by Collan and Michelsen (2020), about automation, adaptation, and manufacturing in Finland and beyond.

Manufacturing 4.0 is a Finnish Strategic Research Council (SRC) project regarding the economic, societal, and technological aspects of the future of manufacturing. The goal of the project is to understand, study, and draw conclusions about the impact of the fourth industrial revolution on Finnish society. The future challenges and requirements posed by the fourth industrial revolution are identified and addressed through seven work packages (WP), which aim to study technical and business aspects, educational and societal policy aspects, and the future of manufacturing. Manufacturing 4.0 intends to provide a good understanding of the future of manufacturing so that Finland and Finnish companies can identify their strengths and weaknesses, and prepare for upcoming changes. Although the project was born and fulfilled in Finland, the results of Publication I are wide-ranging and can be extended to all countries that are addressing the challenges of Industry 4.0.

To disseminate the results of the project, stakeholders from the seven WPs collaborated in the realization of the book edited by Collan and Michelsen (2020). The maintenance of industrial assets is one of the key factors that provide competitiveness to physical assets-based companies, and it was the subject of study of one of the three technological WPs. Publication I provides an introduction to the basics of maintenance and an overview of its evolution from the early time-based approaches to the use of cutting-edge technologies

such as digital twins for predictive maintenance. The objectives of the research include identifying future research avenues and directions of development for maintenance in the industry.

The chapter summarizes how maintenance methodologies evolved and when they are worth being adopted. Whereas maintenance of non-safety-critical assets concerns scheduling inspections and repair actions, for safety-critical applications *condition-based* approaches may be preferred. Where PM approaches fail to guarantee the desired safety and reliability standards, the adoption of a costly CBM system is justified; these proved to be effective in managing safety-critical applications. However, CBM for safety-critical applications may soon be superseded by Prognostics and Health Management, which empowers CBM with a predictive module aimed at producing component Remaining Useful Life estimations. The implementation of PHM is enabled by the recent technological advancements, both in the field of IoT devices and cloud infrastructures and in the application of AI techniques for problems such as feature extraction and failure predictions. Finally, CPS are identified as the technology that could foster the implementation of the PHM approach. Digital Twin technology, i.e., an instance of CPS with specific properties, is presented as a potential solution to the problem of integrated control, production scheduling, and maintenance management. DTs start to fit best practices together in engineering design and process control, and they can help to cover all phases of the life cycle of a system. Among the drawbacks of advanced maintenance methodologies, the reliance on complex technological systems is a hurdle both to the implementation and the maintenance of the monitoring system itself.

Urbani is the primary author. Collan proposed the research topic, and Petri provided the material and the knowledge to write the contents. Urbani contributed to the design and general writing of the chapter supervised by Petri. Urbani carried out the literature study that provided adequate references for the topics addressed in the chapter. Collan carried out the editing of the content, and Brunelli supervised the final revision of the artefact.

3.2 Publication II

Publication II is titled “*A comparison of maintenance policies for multi-component systems through discrete event simulation of faults*” (Publication II) and the goal of the research was to propose a methodology for maintenance policy selection in complex multi-unit systems. Finding a good maintenance policy increases the competitiveness of companies with strict production and reliability requirements, and allows us to estimate the connected uncertainty. Although a wide range of maintenance policies has been proposed

in the body of literature, few efforts were documented in comparing these policies under operative conditions.

The testing of maintenance policies under operational conditions should be the duty of maintenance managers. The problem is particularly relevant when failures occur randomly and the system under analysis is made of multiple non-identical components connected in series, i.e., when each component is critical to the function of the system. The objective of the research is to develop a decision-making methodology that allows such a comparison to be carried out systematically. Moreover, the existence of economies of scale makes it possible to discard maintenance policies for single components in favour of system-level policies. The latter can exploit the grouping of maintenance activities to save maintenance costs and to identify which components can be effectively maintained when an unforeseen failure occurs. To define which policy is the most effective among those that are considered, the maintenance cost was initially considered as the only relevant criteria. The proposed methodology should be flexible and adaptable to different series systems, and the choice of Discrete Event Simulation allows the policies to be tested using different values of the setup cost S of maintenance. Each of the proposed policies can be used to make practical decisions and yields a maintenance schedule.

The proposed analysis is limited to system downtime, which is translated into a monetary value, and system operations have only a side role. The maintenance problem so formulated is relevant in a continuous production setting, namely when the cost of missed production is high and any interruption should be avoided. To validate the proposed methodology, a numerical study is designed and the results of three different policies are compared with the limit that only varies.

In drafting maintenance policies for multi-component systems, the model proposed by Wildeman, Dekker, and Smit (1997) is taken as a starting point. Wildeman et al.'s model has been extended by several other authors to implement aspects that were not addressed in the original paper; however, these models were seldom assessed through simulation of operative conditions, i.e., occurrences of sudden faults. Under the assumptions made in the paper, the setup cost and the overall cost of maintenance are linearly related to all the policies. The higher the setup cost, the more the policies tend to group maintenance activities; consequently, the number of PM tasks tend to decrease as shown by Figure 3.1 for all the policies. Figure 3.1 shows the number of interventions that were performed on average at each S ; the average number of activities per intervention is indeed higher for PM-triggered interventions than for CM ones. The average cost of the adaptive grouping and the opportunistic grouping policies was found to be similar under the assumptions made in the paper. Using the cost as the only criterion was found to be a poor choice because the observed distributions of costs overlapped. Discerning which policy was more

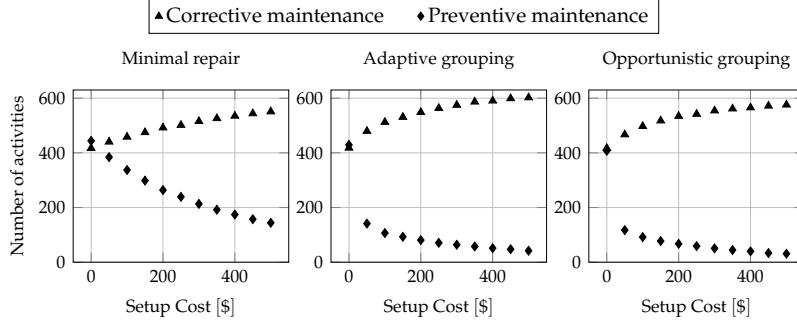


FIGURE 3.1: The number of maintenance interventions caused by PM and CM per each policy. Reprinted from Publication II.

effective than the others required considering at least a second criterion; the choice of the authors fell on availability. Figure 3.2 shows a comparison of policies' performance in terms of cost and availability. When the policies are compared using multiple crite-

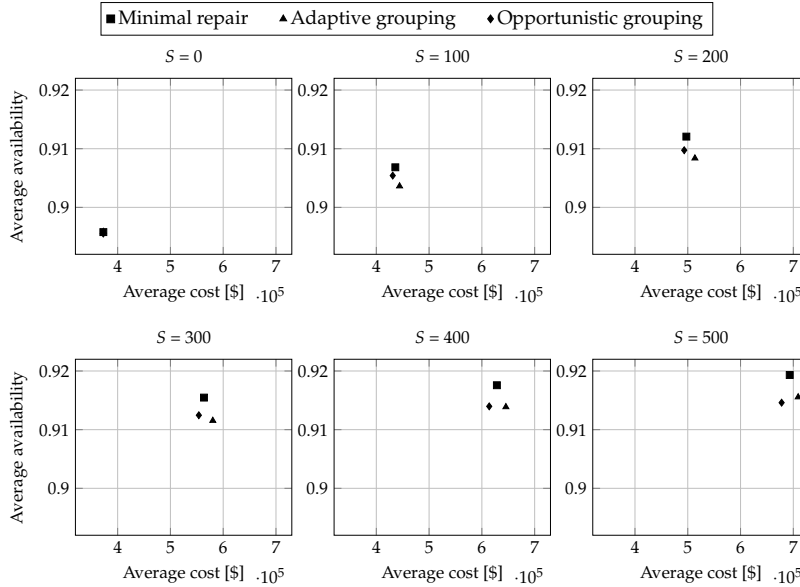


FIGURE 3.2: A bi-criteria comparison of the analysed policies at variable setup costs S . Reprinted from Publication II.

ria, there is no dominant one, and the final choice is a matter of trade-offs. This result

highlights the fact that the studied policies aim at minimizing the cost of maintenance, but they overlook the availability of the system. The opportunistic grouping policy was expected to be the most effective policy due to its ability to optimize maintenance even in the case of unforeseen failures. However, the tendency of the opportunistic grouping policy to anticipate some maintenance activities for purely economic reasons favoured the adaptive grouping policy in terms of availability. At a low value of S , even the minimal repair policy, which presented no mechanism to exploit grouping, presented a higher availability than the opportunistic grouping policy.

Urbani is the primary author. Urbani proposed the research questions and carried out the literature research. Urbani autonomously designed and coded the numerical simulation experiments to test the maintenance policies. The design and general writing of the paper were conducted by Urbani with the supervision of Brunelli. Collan contributed to the general supervision and final editing of the manuscript.

3.3 Publication III (under review)

Maintenance of complex multi-unit systems is often a multi-objective optimization problem. In the manuscript “*An approach for bi-objective maintenance scheduling on a networked system with limited resources*” (Publication III), the preventive maintenance problem is solved for a multi-unit system presenting an arbitrary structure in a multi-objective setting.

The goal of the project was to develop a model for PM scheduling that could be used in practice. In complex systems with machines connected arbitrarily, PM can be organized by exploiting synergies among components, that is, servicing multiple components simultaneously may yield a more efficient use of resources. Grouping PM of multiple components allows to exploit the missing flow of work from downstream machines to perform PM on upstream machines and vice versa; conversely, stopping certain combinations of machines may lead to severe performance reductions. The goal is to exploit such synergies to maximize the flow of jobs through the system during a given period. On the other hand, shifting PM activities from their optimal date triggers the payment of an expected CM cost, which accounts for the risk of delaying maintenance and for the waste of resources due to advancing PM. The scope of the problem also includes constraining the number of available maintenance staff to fulfil PM; model validation is carried out through a sensitivity analysis of maintenance staff, and the relative cost and flow of jobs of the found solutions.

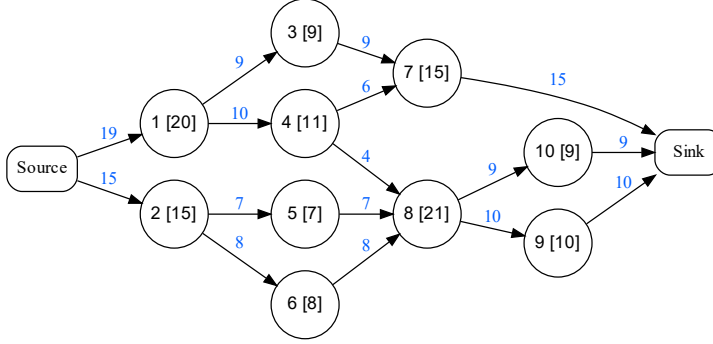


FIGURE 3.3: A graph representing the system analyzed in the paper. The circles represent physical assets and arcs represent the connections between assets. The numbers above the arcs (in blue) describe how the maximum flow of 34 is distributed to the nodes, which are represented by their IDs and capacities (in brackets). Reprinted from Publication III.

The multi-unit system of non-identical components was modelled using a directed graph, which described the connections among machines and the flows of jobs that are transferred among the components; Figure 3.3 depicts the system model used in the paper. The numbers above the arcs are the number of jobs per unit time that flow between the connected assets when all the assets are fully available. This choice turned out to be successful, since it was possible to harness already available algorithms for solving flow optimization problems. The grouping of PM maintenance activities was formally demonstrated to be optimal for system components showing submodularity of the loss of throughput, which is a measure of the reduction of the flow of jobs through the system. Such properties inspired the development of a heuristic for maintenance scheduling optimization: The NSGA-II (Deb et al., 2002) multi-objective genetic algorithm (GA) was adapted to the problem and showed to be an effective tool for finding a set of non-dominated solutions. Figure 3.4 shows four sets of points corresponding to a given number of available maintenance staff r . The solutions with a high number of available maintenance staff enabled the flow loss of jobs FL to stay at a low level, but they were more expensive in terms of the expected cost of maintenance IC. Since the multi-objective GA does not provide a unique solution, the decision-maker is in charge of choosing the desired trade-off between FL and IC using experience. To help the decision maker, the Gantt chart of a PM schedule was coupled to the flow level diagram, thus making it possible to analyze the system performance as a function of time. Figure 3.5 shows the Gantt chart

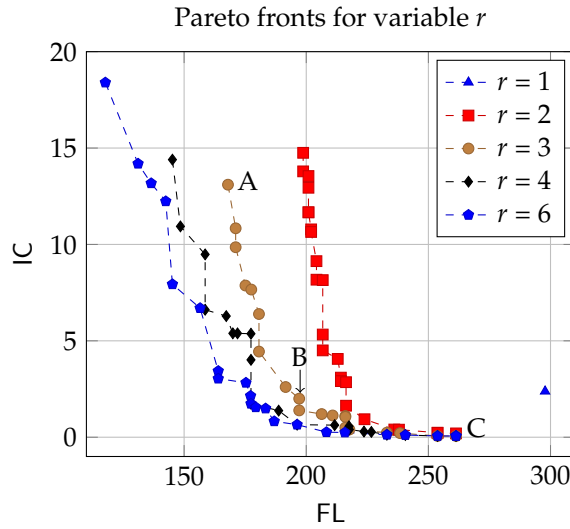


FIGURE 3.4: The frontiers of non-dominated solutions obtained using a variable number of resources. Reprinted from Publication III.

of a schedule of activities and the corresponding flow level; the leftmost solution corresponds to a high degree of grouping of PM activities, whereas the rightmost one shows a schedule with lower performance in terms of FL but a very low IC. The solution in Figure 3.5b can be thought of as an intermediate solution between the previous two.

A practical contribution linked to Publication III regards the creation of a web-based dashboard that allowed a user to interact with the model. The tool was developed with only the purpose of testing the proposed model and producing graphical insights into the results; despite the dashboard being far from deployed in a production setting, it is reasonable to think that the obtained results can be used in practice. Figure 3.6 shows the system design page, whilst parameters of the system components and the system configuration are declared.

Urbani is the primary author. Urbani proposed the research topic and carried out the literature research to motivate the development of the proposed model. The proposition that motivates the grouping approach was developed and proved by Brunelli. Urbani carried out the development of the algorithmic procedure to solve the model under the guidance of Brunelli. Urbani performed the implementation of the algorithm and numerical analysis. Urbani, Brunelli, and Punkka contributed to the design of the manuscript. Urbani and

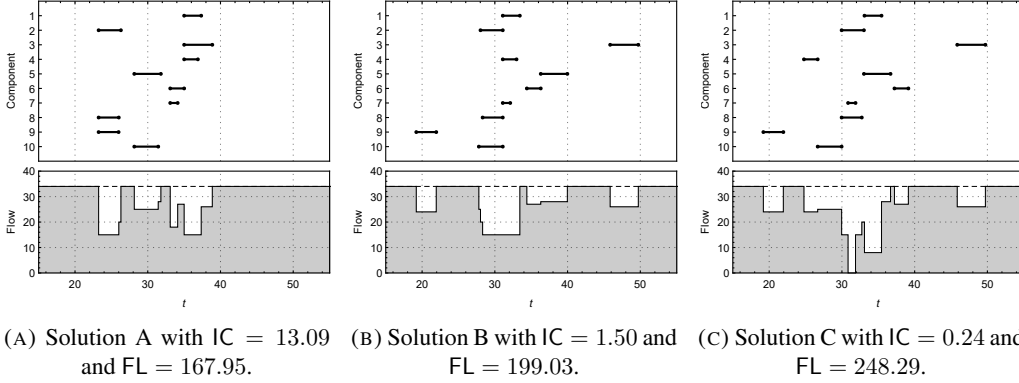


FIGURE 3.5: Gantt charts and flow levels of the solutions A, B, C in Figure 3.4. The figure is reproduced from Publication III.

Brunelli wrote the manuscript. Comments to the results and conclusions are the outcome of the joint effort of Urbani, Brunelli, and Punkka.

3.4 Publication IV

Publication IV is titled “*Maintenance optimization for a multi-unit system with digital twin simulation*” (Publication IV) and concerns maintenance optimization of a complex industrial system through simulation of a digital twin model.

The goal of the work is to study the influence that macro-economic variables play on the profitability of a mining industry. The scope of the work was to build a multi-domain simulation model that integrates an Operations and Maintenance (O&M) simulation-optimization module and a cash-flow module for profitability analysis of the mining industry. Such a model resembles a digital twin (DT), which exploits information about the spot price of the ore to steer decisions about O&M. The scope of the paper includes modelling the uncertainty inherent to macro-economic variables, such as the spot price of ore, and the cost of maintenance. In addition to showing how the model can optimize the long-term O&M policy, the limitations of such an approach are also presented.

The choice of the DT technology to replicate mining operations and cash-flow generation was shown to be an element of novelty in a literature study. Moreover, maintenance optimization was found to be a major goal of the early DT models, which were designed to control and optimize physical systems utilizing a digital counterpart. The study also highlighted that 1) a clear framework for the use of co-simulation models—i.e., of software

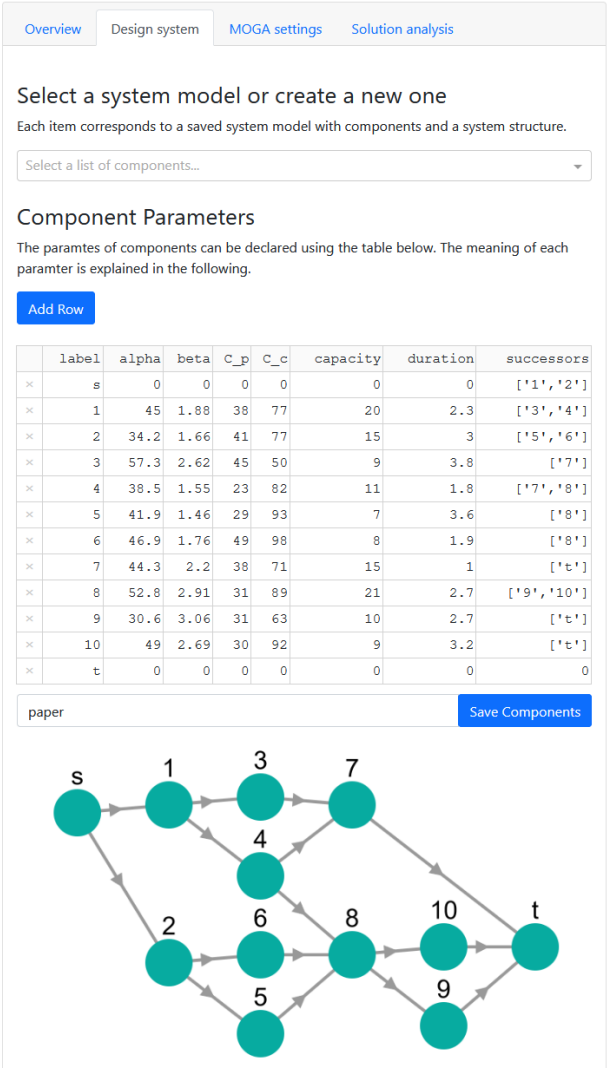


FIGURE 3.6: A screenshot of the “System design” page of the web app realized during the development of Publication III.

systems integrating multi-domain simulation models—does not exist, and 2) the efforts of the scientific community were mainly focused on co-simulation models of technical systems. Only a few papers addressed the modelling of techno-economi-cal systems through

the use of DT; the works produced by the Centre for Digital Built Britain at the University of Cambridge through the National Digital Twin programme were a remarkable example (Centre for Digital Built Britain, 2021).

Figure 3.7 depicts the architecture of the proposed model. Similarly to CPS, DT connects

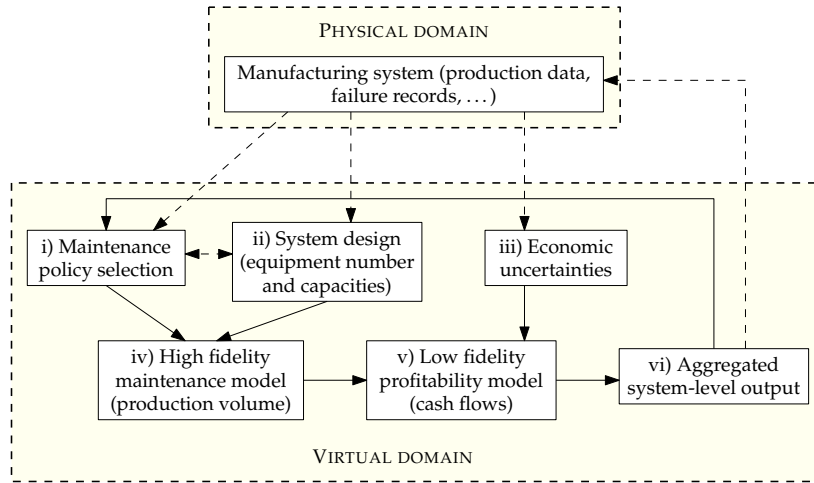


FIGURE 3.7: A schematic representation of the flow of information within the simulated digital twin model. Reprinted from Publication IV.

the physical and the virtual domains. Since a real virtual domain was not available, it was reproduced through a simulation model, which was used both to test the effect of decisions and to evaluate the uncertainty connected with future scenarios. The virtual domain included a high-fidelity maintenance model, represented by box iv) in Figure 3.7, and a low-fidelity profitability model v), which made use of aggregated data from iv) and economic uncertainty from iii) to produce the output vi) that was used to make decisions. The proposed model was tested through two numerical experiments. The first experiment aimed at verifying and validating the overall DT approach. Within the limits of the proposed model, we observed that the spot price of ore had little effect on the maintenance policy, and it was rather preferred to maximize utilization of the facilities and production throughput. The results highlighted the role of maintenance as a “necessary evil” with only little potential on the upside. The second experiment added uncertainty about the cost of maintenance and optimized the system under these new assumptions. Simulations were confirmed to be an effective tool for optimizing operations considering the uncertainty of the parameters used and it provided meaningful insights for operational and investment decisions—e.g., despite the large availability of equipment, the configuration

that maximized revenues involved a low number of trucks and shovels, due to the strong dependence on the limited available maintenance resources.

The limitations connected to the use of multi-domain models in managerial decision-making are discussed. First of all, a large number of degrees of freedom in modelling the system allowed high dimensional data to be produced. The high dimensionality and heterogeneity of data did not allow them to be fully exploited and often it was necessary to aggregate data to transfer them from one module to the other.

The validation and verification of results were limited by the availability of real data and by the absence of a real physical system to study. On the other hand, studies like this allow rapid development and testing of new ideas.

Urbani is the secondary author. The research questions were formulated by Savolainen, whereas Urbani carried out the literature study to show the relevance of the questions and the novelty of the research. Savolainen provided expertise in the mining industry. Urbani designed and coded the simulation-optimization experiment, to which the SD module written by Savolainen was connected. The design and general writing of the paper, except the results regarding the SD module, was conducted by Urbani, whereas Savolainen edited the contents.

3.5 Publication V

The book chapter “*Additive manufacturing cases and a vision for predictive analytics and additive manufacturing-based maintenance business model*” (Publication IV) is also part of the book “*Technical, Economic and Societal Effects of Manufacturing 4.0*” edited by Collan and Michelsen (2020), about automation, adaption and manufacturing in Finland and beyond.

To investigate the potential benefits offered by additive manufacturing (AM) to maintenance optimization, two innovative applications of AM in manufacturing and health-care are presented and the in-use AM-based maintenance business models are outlined. Finally, the authors envisage the potential business models that could be developed by exploiting AM and predictive maintenance; the implications of these systems on the surrounding manufacturing ecosystem are mentioned.

In contrast to the other publications, the followed research method was purely qualitative. The data required to describe the two application cases were gathered by interviewing subject experts and through literature research.

The application of AM to the refurbishment of metallic dies, otherwise called saddles, showed to be a successful application of AM to maintenance. The shaping of metallic parts requires the use of metallic dies made of hard alloys, which show high resistance to wear. Hard alloys are expensive and their use is limited to the minimum that is necessary. Due to regular wear and tear phenomena, dies lose their original shape and they can finally cause quality issues to the end product; refurbishment of the die is therefore compulsory to restoring the desired quality level. Machines that integrate additive and subtractive manufacturing, also referred to as *hybrid* machines, can perform refurbishment as a unique process. This allows them to i) keep the piece in the same venue and hence to avoid calibration issues, and ii) reduce the use of expensive materials through 3D printing. On the other hand, the manual process would require the use of separate machines for subtraction and addition of material, and calibration issues may yield a poor quality of the final result. What makes it worth using AM is its ability, thanks also to the software, to optimize the refurbishment of specific failure instances, which can hardly be automatized in other ways; this is an example of the *mass customization* of a process. The upside potential of using AM is not only technological. New business models can be developed around sharing both the physical facilities and the expert knowledge that are required for using AM. The refurbishment of mechanical parts is already offered as a service for several manufacturing applications, therefore AM could be introduced incrementally to the range of products. The adoption of hybrid AM technology is indeed limited by its cost and by the work volume, which already allows processing “fairly” large parts and is bound to improve.

New business models revolving around the use of AM can be envisaged. There are two items that contribute to making the adoption of AM expensive: Firstly, the cost of equipment, which in the case of hybrid AM machines could grow up to a million euro, and secondly, the cost of trained technicians. The latter varies depending on the AM technology. The expertise required to set up, maintain, and use a polymer printer is rather low if compared to metal printers. Due to the high initial costs paid for the adoption of AM, the equipment utilization must be high to reach the profitability of the investment; often, Machine as a Service (MaaS) business models may be a more effective alternative to buying.

In the authors’ opinion, the possibility to integrate predictive maintenance and AM may deliver a high competitive advantage to companies. The envisioned model is depicted in Figure 3.8, where starting from “Instrumented equipment” it is possible to build “Predictive maintenance optimization systems” that can be harnessed for different purposes. The ability to forecast the remaining useful life of components can be exploited to coordinate maintenance personnel, to optimize the management of spare parts inventory, and to refine the design of components. For instance, inventory levels can be lowered by printing

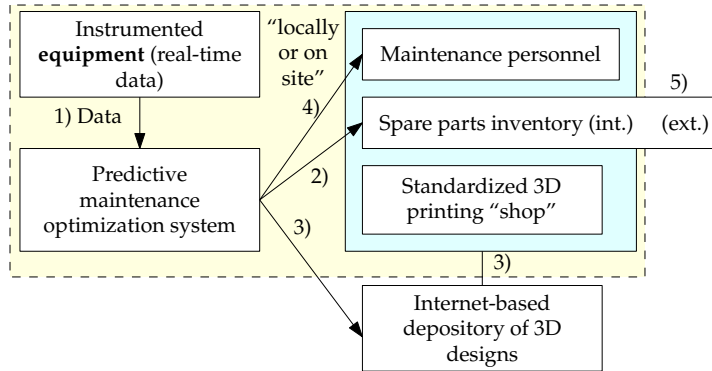


FIGURE 3.8: A blueprint model for PM-oriented maintenance ecosystem.
Reprinted from Publication IV.

some components on-demand or by refurbishing components on the fly instead of holding a spare in the inventory. Storage of parts would be reduced to those parts that could not be produced just in time. The refurbishment venue would be potentially moved closer to the workshop and delivery issues could be eliminated. The logistics of spare parts would become the logistics of materials to produce spares and the digital logistics of models.

Urbani is the primary author. Collan proposed the research topic. Urbani interviewed the subject expert, Prof. Paolo Bosetti from the University of Trento, and gathered the information about both case studies. Urbani contributed to the design and wrote the first two sections of the chapter. Collan wrote the third section of the chapter and performed the editing and supervision of the whole manuscript.

3.6 Positioning of the research

In light of the methodological framework presented in Section 2.1, the positioning of the research activity within such a framework is discussed in the following.

Publications I and V are based on a qualitative research method (Bell, Bryman, and Harley, 2018, p. 383) and as such, they do not involve the activities connected to the "Scientific model" in Figure 2.1. This approach to scientific investigation moves back and forth between points I), II), and IV) in Figure 2.1 with limited connections between the solution IV) and reality I). The main concern during this research activity is the conceptualization of real-world maintenance problems, which are subsequently categorized and located in the scientific literature. The solution to different maintenance problems is

presented, and in case this has been implemented a description of the case study is proposed. The goal of Publications I and V is to provide a broad range of readers with an overview of what benefits are delivered by the fourth industrial revolution to maintenance; a qualitative research method has been regarded as more efficient to the dissemination of scientific knowledge towards a non-technical audience.

Publications II, III, and IV deal with the resolution of specific maintenance problems, which were addressed using a quantitative research approach (Bell, Bryman, and Harley, 2018, p. 147). The scientific investigation concerned the loop among parts II), III), IV), and II) again in Figure 2.1, and the goal was to further develop existing conceptual models of the preventive maintenance problem. Publications II and III addressed the dynamic grouping problem and proposed an investigation into its effectiveness and an extension to complex systems respectively. Both the activities started from the “Conceptual model”, II) in Figure 2.1, and expanded a previously existing idea before to formalize it in a “Scientific model”. The scientific model was then solved, i.e., an algorithmic procedure was developed that can produce feasible solutions to the problem. The quality of the solution was evaluated by feeding it back into the conceptual model and looking at some performance measures of the studied system, such as availability and productivity. Activity number 5 in Figure 2.1 was particularly intense, because it triggered several changes to the algorithmic procedures developed during model solving. In Publication IV, the conceptual model was the main result of the research, and the goal was to validate such a model and study its limitations. The scientific activity resembles that of Publications II and III. However, most of the activity concerned observing how different solutions to the same problem returned a feedback to the conceptual model, e.g., how the conceptual model could eventually be improved. Although Publication IV does not directly address a real-world problematic issue, the experience in the mining industry of one of the authors played an important role in drafting the research question, and in guiding the development of the conceptual model. As mentioned in Section 2.1, the scientific investigation process can be schematized according to the view proposed by Mitroff et al. (1974), but it is in fact a complex process involving several parts of Figure 2.1 concurrently.

3.7 Summary of publications

Table 3.1 summarizes the objectives, research method, research data, and the contribution of each publication.

TABLE 3.1: A summary of the research results.

Publication	Objective	Research method(s)	Data	Contribution
Publication I	To provide insight into the evolution of maintenance research in the context of Industry 4.0.	Qualitative research	Literature review	To identify the key points of RCM maintenance approaches.
Publication II	To develop a methodology for comparison of maintenance policies in multi-component systems through DES.	Quantitative research	Data used in simulations are generated synthetically from arbitrary TTF distributions.	A multi-criteria analysis of three maintenance grouping policies is presented. The simulation workflow used in the numerical experiments is summarized and offered as a reference.
Publication III	To develop a model for PM of multi-unit systems with non-identical components connected according to an arbitrary structure. The models must be able to account for limited maintenance staff.	Quantitative research	Data were synthetically generated from TTF distributions and maintenance cost values that were arbitrarily chosen.	Grouping of maintenance tasks was proved to be optimal for specific instances of the problem. The bi-objective approach to PM of complex systems proved to be effective and it allowed finding sets of equivalent solutions among which to choose. The developed tools allowed the visualization of the PM schedule and the relative system performance
Publication IV	To study the influence of macro-economic variables on O&M through simulation of a digital twin.	Qualitative and quantitative research	Qualitative data from a structured literature study, and numerical data generated synthetically through simulation mining industry environment.	A model to simulate a DT is drafted; numerical experiments validate the proposed DT model and show the potential use; the limitations of using heterogeneous simulation modules and data handling issues are highlighted.

Continued on next page

Table 3.1 – Continued from previous page

Publication	Objective	Research method(s)	Data	Contribution
Publication V	To show the potential of AM for maintenance, and to offer a vision for predictive maintenance-oriented ecosystems.	Qualitative research	Information gathered from interview data of subject experts.	A vision on the use of AM for maintenance is offered; AM-based business models are identified; and a blueprint model for the use of PdM is drafted.

Chapter 4

Discussion and conclusions

The research questions posed in Section 1.2 are answered in Section 4.1 and the implications and limitations of the research outcomes are discussed in Sections 4.1.2 and 4.1.3 respectively. Future research avenues are outlined in Section 4.2 and conclusions are drafted in Section 4.3.

4.1 Discussion

This research focused on developing novel OR models to optimize maintenance policies, and on studying the implications of Industry 4.0 technologies on maintenance management. The research outcomes support the research questions posed in Section 1.2, and they also have practical and theoretical implications.

4.1.1 Answering the research questions

Answer to Question 1 Question 1 asked, “*How is maintenance optimization evolving in light of the fourth industrial revolution?*” Publication I provides a non-exhaustive overview of maintenance optimization from the early models to the use of cutting-edge technologies such as digital twins. The level of sophistication in policy design increased with the need to manage ever more complex systems and to guarantee a high safety level. Maintenance approaches are changing in order to embed IoT devices, which enable real-time monitoring of the system state, and maintenance policies are developed to embed novel dynamic models. Maintenance optimization is evolving along three main directions: First, the formalization of maintenance problems allows harnessing traditional mathematical optimization techniques from the field of OR, hence solving larger problem instances. The availability of large datasets of signals enables the use of artificial intelligence techniques, which have been specifically developed for feature selection, signal reconstruction, and fault detection and diagnosis. The third direction of development

concerns the creation of “smart” environments, which allow the seamless integration of business objectives with operations. The trend seems to be towards that of creating holistic models, which can account not only for equipment monitoring and PM scheduling, but also for resource management and coordination of O&M.

Answer to Question 2 Question 2 regarded balancing system performance and the need for the maintenance of physical assets. The answer is certainly system-specific, but some general tenets can be exploited when a maintenance policy is designed. First, multiple criteria should be considered to evaluate the performance of the system, hence the maintenance policy. Reliability, availability, maintainability, and safety (RAMS) are four examples of commonly used evaluation criteria. The utilitarian assumption whereby system performance can be measured attributing a monetary value to all criteria may be unsuitable, and RAMS criteria might better be accounted for separately because of their incommensurability. Secondly, network effects may exist in complex systems. These can be of different types—i.e., economic, stochastic, structural, and resource dependencies—and they can be exploited both to improve system performance and to avoid unfavourable combinations of events. The creation of ad-hoc models for maintenance management is the preferred way to identify cost-effective combinations of maintenance events.

The question is partly answered by Publication III, where a model for maintenance optimization of complex systems is proposed. By leveraging on positive network effects, the proposed model allows optimizing system performance and maintenance costs, which are the two criteria used to evaluate the resulting maintenance schedules. Moreover, Publication II proposes a numerical study about maintenance policies comparison and it highlights the importance of considering multiple criteria for the evaluation of policies. Finding the optimal maintenance policy hardly concerns only the cost of maintenance; rather it involves other criteria, e.g., the availability of the system.

Answer to Question 3 Question 3 concerns the integration of maintenance management systems into Cyber-Physical Systems (CPS). The question is directly addressed in Publication IV, where the goal is to study the effect of macro-economic variables on operations and maintenance of the mining industry through simulating a digital twin (DT). According to the literature study conducted in Publication IV, maintenance optimization was one of the main goals in realizing DTs, which were identified as an effective tool for integrating heterogeneous models and fulfilling CBM. The main hallmark of a DT is indeed the possibility to exchange data between the physical and the virtual domain in almost real-time. The creation of a virtual counterpart for physical systems improves the ability to read the future through simulation of the possible scenarios, which is a key factor to maintenance optimization. The challenges to realizing such a physical/virtual

space are several: The absence of a common framework to integrate multi-domain models and the difficulty to exchange data between models are two examples. In Publication IV, descriptive statistics concepts, such as the mean and the standard deviation of data, were used to allow the communication between models; however, it is simple to envisage that whole time-series can be used to provide complete information. The limitations connected to the implementation of CPS regard the validation of the proposed model, and the reliability of hardware and software parts, when dealing with safety-critical systems.

Answer to Question 4 There are several ways in which AM can improve maintenance processes. The creation of mock-ups of body parts and the refurbishment of worn mechanical parts are only two examples, which are described in Publication V. These examples are representative of two general use-cases of AM, i.e., the ability to create mock-ups of newly designed objects, and the ability to repair worn mechanical parts. In addition, the ability of AM to produce single-item batches without losing profitability is relevant to the deployment of refurbishment processes. Thanks to the integration of additive and subtractive manufacturing into *hybrid* machines, the refurbishment of mechanical parts can be carried out as a unique automated process. Compared to the manual one, the new process yields more precise results in shorter times, and it uses a lower amount of material. On the other hand, advanced software solutions and skilled personnel are required to deploy the above-mentioned refurbishment process. From a technological perspective, this is only one example of what can be realized using AM.

New business models are also enabled by the adoption of AM. These can rely on more efficient logistics and on the advantages delivered by smart environments. For instance, if the repair of parts can be fulfilled in-house, repair times could be shortened and delivery issues could be avoided. The logistics of repair parts could potentially be transformed into the logistics of materials for repair parts, and the need for spare parts would be reduced to those items that could not be fixed by AM. Moreover, in a smart factory, predictive abilities could be exploited to optimize the timing of PM intervention, hence printing and refurbishment operations.

4.1.2 Theoretical and practical implications

The theoretical contribution of this thesis progressively increased with the advancement of the doctoral studies. Publication II was the first work in chronological order and it was carried out with limited knowledge of reliability theory and the relative literature; the early idea behind Publication II was the development of the author's master's thesis. Publication IV, which was the second in chronological order, witnesses the achievement

of a methodologically structured research project, where a clear literature gap is identified and addressed through the proposal of a framework for the use of digital twins for techno-economical analysis. The theory proposed in Publication IV is validated through a numerical experiment and some practical implications are identified; the latter regard the coexistence of heterogeneous modules in a co-simulation setting and the exchange of data between such modules. In Publication III, the highest level of theoretical contribution of the author's doctoral project was achieved: Starting from formally proving a proposition about the maintenance system under study, a novel model for preventive maintenance scheduling was proposed. The model optimizes the highly non-linear performance behaviour of a system showing an arbitrary structure to smartly schedule PM activities. The resolution procedure found a set of non-dominated solutions by leveraging the optimality of grouping maintenance activities. Finally, Publication V provides a limited theoretical contribution by laying out the components of future potential business models that could be built around the use of AM for maintenance.

On the other hand, Publications I and V present several practical implications. The former introduces the reader to the subject of reliability-centred maintenance. Starting from the basic concepts of time-based maintenance and run to failure policies, the reader is guided through the most recent approaches and cutting-edge technologies for maintenance. The practical contribution lies in identifying the advantages and drawbacks of each maintenance approach, and in providing an overview of the maintenance processes connected to Industry 4.0. By the same token, Publication V describes two applications of AM in maintenance, one in the healthcare sector and the other in the manufacturing sector, and presents the relative business models. Identifying such a connection is of practical interest because it allows figuring out how the AM technology can be profitable.

4.1.3 Limitations of the research

The major limitation of this research regards the validation of the proposed theory. According to the methodology presented in Section 2.1 and Section 3.6, the scientific inquiry carried out in this research rarely touched the “Reality” node in Figure 2.1. Due to the lack of an industrial partner, both the validation through an empirical application of the proposed theory and through discussion with a potential end-user could not be carried out. The author acknowledges that this limitation weakens the credibility and relevance of the proposed work, but, on the other hand, it allowed for the quick development and testing of new ideas. The proposed models are indeed “production-ready” in the sense that they have already been formulated, and they could be transferred to reality by taking advantage of the insights provided in the research outcomes of this thesis. In the case of Publication III, the proposed model is usable through a web app and closing the gap with the real world would take a little effort.

A second limitation of this research concerns the aspects relative to Industry 4.0. The proposed theory is mainly focused on OR methods to find and exploit network effects in complex systems and, in particular, to organize maintenance activities in groups. Sections 2.2 and 2.3 showed how maintenance management is relying more and more on condition-based maintenance and on artificial intelligence techniques for fault detection and diagnosis. Despite their central role in modern maintenance, there was no room to carry out research specifically devoted to developing novel contributions in the fields of AI, IoT, or signal processing.

A third and final limitation relates to the framework used to study maintenance policies. Publication II relied on the implementation of deterministic rules and on the use of simulations to study the performance of the proposed policies on a specific system. The approach is effective and can be considered a traditional approach in OR at large, and in maintenance management more specifically. A similar approach was adopted in Publication IV, where one step forward was made. The combined use of simulations and optimization—i.e., of the so-called *simulation-optimization* approach—allowed learning the preventive maintenance thresholds that maximized the profitability of the mine. Publication III optimized the organization of preventive maintenance on systems showing network effects, but it was limited to a static version of the problem. Two interesting research questions regarding Publication III are *how would the proposed policy behave in the context of sequential decision-making?* And, *how would it be possible to optimize a rule for sequential decision-making under uncertainty in such a complex system?* A proposal to continue the present research in a more general decision-making framework is formulated in the following section.

4.2 Prospective future research questions

The future research avenues identified in Publications II, III, and IV are expanded in the following. Moreover, new research trends and active research fields within the area of maintenance optimization are presented.

In Publication IV, research on the use of co-simulation for digital twins embedding techno-economical aspects of a mine was carried out. The focus was on the interaction of heterogeneous simulation modules. The technology used for the simulation-optimization of operations and maintenance (O&M) could be improved by integrating condition-based maintenance (CBM) of the mining equipment. The system would then present a three-level structure, i.e., an equipment-level layer, which mimics the CBM system, a PM layer that schedules PM activities, and a module for the management of financial aspects. The novelty of the proposed research resides in the hierarchical structure of the system, which

would allow studying how to optimize the interaction of heterogeneous and multi-level simulation modules.

Publication III tackled the PM maintenance problem of a complex system showing an arbitrary structure. The proposed model is *static* in the sense that it considers a single optimization time step, whereas in a real-world instance multiple time steps and multiple PM activities per component should be considered; such limitations are justified in the manuscript. However, the model could be further developed under the lens of maintenance *policy optimization*, where a policy is intended as a rule for subsequent decisions under uncertainty. A suitable framework for optimization of maintenance policies has recently been proposed by Barlow et al. (2020), who considered a complex multi-unit system with limited maintenance resources and a single optimization objective. According to Zio (2009), maintenance of complex systems could benefit from considering multiple objectives, therefore an interesting research avenue could be to consider the sequential decision-making problem with multiple objectives in the context of maintenance. In such a case, a clear hurdle to overcome would be the curse of dimensionality, which soon arises when combinatorial problems, e.g., the dynamic grouping problem, are tackled. In this case, exact models could be limited, whereas approximated methods could be effective (Bengio, Lodi, and Prouvost, 2021). By the same token, the multi-objective sequential decision-making problem could be tackled effectively by adopting a multi-objective reinforcement learning approach (Liu, Xu, and Hu, 2015).

The application of approximated methods to solve combinatorial problems in industrial management showed to be effective; for instance, Huang, Chang, and Arinez (2020) applied a deep reinforcement learning algorithm to solve the PM problem on a series production line. Huang et al.'s algorithm found the grouping/opportunistic strategy to be the most effective, thus confirming the results previously obtained in the literature, but with the exception that they were obtained in a fraction of the time required by exact and heuristic algorithms. In the context of decision-making problems that can be formulated as Markov Decision Processes, the ability of approximated algorithms to overcome the curse of dimensionality proved promising. A few examples are the PM problem with limited resources (Barlow et al., 2020), repair parts management (Compare et al., 2020), production order dispatching (Stricker et al., 2018), and operations management (Rocchetta et al., 2019). A further step towards the application of approximated algorithms to real-world cases concerns the use of simulation models together with artificial intelligence agents (Pinciroli et al., 2020a; Pinciroli et al., 2020b). The expected advantages are learning and adaptability characteristics, and a rapid decision-making process; on the other hand, the interpretability of results, the long training time, and a high demand of computational resources are some of the challenges still to be addressed to obtain efficient decision support systems.

4.3 Conclusions

The scientific activity that was carried out during this research presented a theoretical approach, and consequently, its impact was limited to the academic world. The lack of time and real-world problematic issues to be studied weakened the impact of the research on society. However, the results presented a certain degree of novelty to the actual body of the literature and contributed to expanding the knowledge of the grouping of maintenance activities, which is a relevant and long-debated problem in maintenance planning. Despite the limited practical contribution of the research results, the ability of the author to produce practical tools for decision-making grew steadily during the doctoral studies, and it will be relevant to making an impact on society. To improve the societal impact of research remains author's main objective.

The research conducted in Publication IV can also be seen as an attempt to bring the attention of the scientific community to a less developed aspect of digital twins, i.e., to the implementation of macro-economic variables in technical decision-making processes. As shown in the paper, the novelty of the topic is clear, but the credibility of the proposed solution is weakened by the lack of a real-world validation case.

The contribution of the present research to the practise of maintenance consistently advanced thanks to the results achieved in Publication III. The development of a ready-to-use graphical tool of the proposed model made its validation on a real-world application a concrete possibility. The knowledge of ICT tools and the modelling abilities that were acquired by the author during the doctoral studies are fundamental to the spread of advanced scientific models in the industrial practice.

References

- Alaswad, Suzan and Yisha Xiang (2017). “A review on condition-based maintenance optimization models for stochastically deteriorating system”. In: *Reliability Engineering & System Safety* 157, pp. 54–63. DOI: 10.1016/j.ress.2016.08.009.
- Alrabghi, A. and A. Tiwari (2016). “A novel approach for modelling complex maintenance systems using discrete event simulation”. In: *Reliability Engineering & System Safety* 154, pp. 160–170. DOI: 10.1016/j.ress.2016.06.003.
- Alrabghi, Abdullah and Ashutosh Tiwari (2015). “State of the art in simulation-based optimisation for maintenance systems”. In: *Computers & Industrial Engineering* 82, pp. 167–182. DOI: 10.1016/j.cie.2014.12.022.
- Aven, Terje (2012). *Foundations of Risk Analysis*. John Wiley & Sons.
- Baraldi, P. et al. (2011). “A randomized model ensemble approach for reconstructing signals from faulty sensors”. In: *Expert Systems with Applications* 38.8, pp. 9211–9224. DOI: 10.1016/j.eswa.2011.01.121.
- Baraldi, Piero et al. (2015). “Robust signal reconstruction for condition monitoring of industrial components via a modified Auto Associative Kernel Regression method”. In: *Mechanical Systems and Signal Processing* 60–61, pp. 29–44. DOI: 10.1016/j.ymssp.2014.09.013.
- Barlow, E. et al. (2020). “A performance-centred approach to optimising maintenance of complex systems”. In: *European Journal of Operational Research*. DOI: 10.1016/j.ejor.2020.11.005.
- Bedford, Tim, Roger Cooke, et al. (2001). *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press.
- Bell, Emma, Alan Bryman, and Bill Harley (2018). *Business Research Methods*. Oxford university press.
- Ben-Daya, Mohamed, Uday Kumar, and DN Prabhakar Murthy (2016). *Introduction to Maintenance Engineering: Modelling, Optimization and Management*. John Wiley & Sons. DOI: 10.1002/9781118926581.
- Bengio, Yoshua, Andrea Lodi, and Antoine Prouvost (2021). “Machine learning for combinatorial optimization: A methodological tour d’ horizon”. In: *European Journal of Operational Research* 290.2, pp. 405–421. DOI: <https://doi.org/10.1016/j.ejor.2020.07.063>.

- Bouvard, Keomany et al. (2011). "Condition-based dynamic maintenance operations planning & grouping. Application to commercial heavy vehicles". In: *Reliability Engineering & System Safety* 96.6, pp. 601–610. DOI: 10.1016/j.ress.2010.11.009.
- Brans, Jean-Pierre and Pierre L. Kunsch (2010). "Ethics in operations research and sustainable development". In: *International Transactions in Operational Research* 17.4, pp. 427–444. DOI: <https://doi.org/10.1111/j.1475-3995.2009.00728.x>.
- BS 4778 (1991). *Glossary of Terms Used in Quality Assurance Including Reliability and Maintainability Terms*. British Standards Institution, London.
- Centre for Digital Built Britain (2021). *National Digital Twin Programme*. <https://www.cdabb.cam.ac.uk/what-we-do/national-digital-twin-programme>. Accessed on 29/06/2021.
- Cho, Danny I. and Mahmut Parlar (1991). "A survey of maintenance models for multi-unit systems". In: *European Journal of Operational Research* 51.1, pp. 1–23. DOI: 10.1016/0377-2217(91)90141-H.
- Cimino, Chiara, Elisa Negri, and Luca Fumagalli (2019). "Review of digital twin applications in manufacturing". In: *Computers in Industry* 113, p. 103130. DOI: <https://doi.org/10.1016/j.compind.2019.103130>.
- Collan, Mikael and Karl-Erik Michelsen (2020). *Technical, Economic and Societal Effects of Manufacturing 4.0: Automation, Adaption and Manufacturing in Finland and Beyond*. Springer Nature. DOI: 10.1007/978-3-030-46103-4.
- Compare, Michele et al. (2020). "A reinforcement learning approach to optimal part flow management for gas turbine maintenance". In: *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability* 234.1, pp. 52–62. DOI: 10.1177/1748006X19869750. eprint: 10.1177/1748006X19869750.
- Dao, Cuong D. and Ming J. Zuo (2017). "Selective maintenance of multi-state systems with structural dependence". In: *Reliability Engineering & System Safety* 159, pp. 184–195. DOI: <https://doi.org/10.1016/j.ress.2016.11.013>.
- De Jonge, B. and P. A. Scarf (2020). "A review on maintenance optimization". In: *European Journal of Operational Research* 285.3, pp. 805–824. DOI: 10.1016/j.ejor.2019.09.047.
- Deb, K. et al. (2002). "A fast and elitist multiobjective genetic algorithm: NSGA-II". In: *IEEE Transactions on Evolutionary Computation* 6.2, pp. 182–197. DOI: 10.1109/4235.996017.
- Dekker, Rommert, Ralph E Wildeman, and Frank A Van der Duyn Schouten (1997). "A review of multi-component maintenance models with economic dependence". In: *Mathematical Methods of Operations Research* 45.3, pp. 411–435. DOI: 10.1007/BF01194788.

- Derler, Patricia, Edward A. Lee, and Alberto Sangiovanni Vincentelli (2012). "Modeling Cyber-Physical Systems". In: *Proceedings of the IEEE* 100.1, pp. 13–28. DOI: 10.1109/JPROC.2011.2160929.
- Di Maio, Francesco et al. (2013). "Fault detection in nuclear power plants components by a combination of statistical methods". In: *IEEE Transactions on Reliability* 62.4, pp. 833–845. DOI: 10.1109/TR.2013.2285033.
- Dinh, Duc-Hanh, Phuc Do, and Benoit Iung (2020). "Degradation modeling and reliability assessment for a multi-component system with structural dependence". In: *Computers & Industrial Engineering* 144, p. 106443. DOI: <https://doi.org/10.1016/j.cie.2020.106443>.
- Do, Phuc, Phil Scarf, and Benoi Iung (2015). "Condition-based maintenance for a two-component system with dependencies". In: *IFAC-PapersOnLine* 48.21. 9th IFAC Symposium on Fault Detection, Supervision and Safety for Technical Processes SAFEPROCESS 2015, pp. 946–951. DOI: 10.1016/j.ifacol.2015.09.648.
- Do Van, Phuc et al. (2013). "Dynamic grouping maintenance with time limited opportunities". In: *Reliability Engineering & System Safety* 120, pp. 51–59. DOI: 10.1016/j.res.2013.03.016.
- Ersdal, Gerhard and Terje Aven (2008). "Risk informed decision-making and its ethical basis". In: *Reliability Engineering & System Safety* 93.2, pp. 197–205. DOI: <https://doi.org/10.1016/j.res.2006.12.018>.
- Fei, Tao et al. (2018). "Digital twin-driven product design, manufacturing and service with big data". In: *The International Journal of Advanced Manufacturing Technology* 94.9-12, pp. 3563–3576. DOI: <https://doi.org/10.1007/s00170-017-0233-1>.
- Fonseca, C.M. and P.J. Fleming (1993). "Multiobjective genetic algorithms". In: *IEE Colloquium on Genetic Algorithms for Control Systems Engineering*, pp. 6/1–6/5.
- Forrester, Jay W. (1994). "System dynamics, systems thinking, and soft OR". In: *System Dynamics Review* 10.2-3, pp. 245–256. DOI: 10.1002/sdr.4260100211.
- Fu, Michael C. (2002). "Optimization for simulation: Theory vs. practice". In: *INFORMS Journal on Computing* 14.3, pp. 192–215. DOI: 10.1287/ijoc.14.3.192.113. eprint: 10.1287/ijoc.14.3.192.113.
- Gallo, Giorgio (2004). "Operations research and ethics: Responsibility, sharing and co-operation". In: *European Journal of Operational Research* 153.2. Management of the Future MCDA: Dynamic and Ethical Contributions, pp. 468–476. DOI: [https://doi.org/10.1016/S0377-2217\(03\)00167-X](https://doi.org/10.1016/S0377-2217(03)00167-X).
- Geng, Junbao, Michael Azarian, and Michael Pecht (2015). "Opportunistic maintenance for multi-component systems considering structural dependence and economic dependence". In: *Journal of Systems Engineering and Electronics* 26.3, pp. 493–501. DOI: 10.1109/JSEE.2015.00057.

- Goldberg, David E and John Henry Holland (1988). *Genetic Algorithms and Machine Learning*. Kluwer Academic Publishers-Plenum Publisher.
- GTAI (2014). *Industrie 4.0 Smart Manufacturing for the Future*. Germany Trade & Invest.
- Hilber, Patrik et al. (2007). “Multiobjective optimization applied to maintenance policy for electrical networks”. In: *IEEE Transactions on Power Systems* 22.4, pp. 1675–1682. DOI: 10.1109/TPWRS.2007.908468.
- Huang, Jing, Qing Chang, and Jorge Arinez (2020). “Deep reinforcement learning based preventive maintenance policy for serial production lines”. In: *Expert Systems with Applications* 160, p. 113701. DOI: 10.1016/j.eswa.2020.113701.
- IEC 30600 (1992). *Dependability Management*. International Electrotechnical Commission, Geneva.
- ISO 13381-1 (2005). *Condition Monitoring and Diagnostics of Machines – Prognostics – Part 1: General Guidelines*. International Standards Organization.
- ISO 52900 (2015). *Additive Manufacturing – General Principles – Terminology*. International Standard Organization, Geneva.
- ISO 55000 (2014). *Asset management – Overview, principles and terminology*. International Standard Organization, Geneva.
- ISO 8402 (1986). *Quality Vocabulary*. International Standard Organization, Geneva.
- Kim, Nam-Ho, Dawn An, and Joo-Ho Choi (2017). *Prognostics and Health Management of Engineering Systems*. Switzerland: Springer International Publishing.
- Kleindorfer, George B., Liam O’Neill, and Ram Ganeshan (1998). “Validation in simulation: Various positions in the philosophy of science”. In: *Management Science* 44.8, pp. 1087–1099. DOI: 10.1287/mnsc.44.8.1087.
- Konak, Abdullah, David W. Coit, and Alice E. Smith (2006). “Multi-objective optimization using genetic algorithms: A tutorial”. In: *Reliability Engineering & System Safety* 91.9, pp. 992–1007. DOI: 10.1016/j.res.2005.11.018.
- Kritzinger, Werner et al. (2018). “Digital Twin in manufacturing: A categorical literature review and classification”. In: *IFAC-PapersOnLine* 51.11. 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018, pp. 1016–1022. DOI: <https://doi.org/10.1016/j.ifacol.2018.08.474>.
- Kuhn, Thomas S. (2012). *The Structure of Scientific Revolutions*. University of Chicago press.
- Law, A. M. (2014). *Simulation Modeling and Analysis. Vol. 5th International Edition*. New York, New York: McGraw-Hill, Inc.
- Lee, E. A. (2008). “Cyber physical systems: Design challenges”. In: *2008 11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing (ISORC)*, pp. 363–369. DOI: 10.1109/ISORC.2008.25.

- Liang, Zhenglin and Ajith Kumar Parlikad (2020). “Predictive group maintenance for multi-system multi-component networks”. In: *Reliability Engineering & System Safety* 195, p. 106704. DOI: 10.1016/j.ress.2019.106704.
- Lim, Kendrik Yan Hong, Pai Zheng, and Chun-Hsien Chen (2019). “A state-of-the-art survey of Digital Twin: techniques, engineering product lifecycle management and business innovation perspectives”. In: *Journal of Intelligent Manufacturing*, pp. 1–25. DOI: 10.1007/s10845-019-01512-w.
- Liu, C., X. Xu, and D. Hu (2015). “Multiobjective reinforcement learning: A comprehensive overview”. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45.3, pp. 385–398.
- Liu, Ruonan et al. (2018). “Artificial intelligence for fault diagnosis of rotating machinery: A review”. In: *Mechanical Systems and Signal Processing* 108, pp. 33–47. DOI: 10.1016/j.ymssp.2018.02.016.
- Miettinen, Kaisa (2012). *Nonlinear Multiobjective Optimization*. Vol. 12. Springer Science & Business Media. DOI: 10.1007/978-1-4615-5563-6.
- Mild, Pekka and Ahti Salo (2009). “Combining a multiattribute value function with an optimization model: An application to dynamic resource allocation for infrastructure maintenance”. In: *Decision Analysis* 6.3, pp. 139–152. DOI: 10.1287/deca.1090.0143.
- Mill, John Stuart (1998). *The Cambridge Companion to Mill*. Cambridge University Press.
- Mitroff, Ian I et al. (1974). “On managing science in the systems age: two schemas for the study of science as a whole systems phenomenon”. In: *Interfaces* 4.3, pp. 46–58. DOI: 10.1287/inte.4.3.46.
- Nakajima, Seiichi (1988). “Introduction to TPM: Total Productive Maintenance.” In: *Productivity Press, Inc., 1988*, p. 129.
- Negri, Elisa, Luca Fumagalli, and Marco Macchi (2017). “A Review of the Roles of Digital Twin in CPS-based Production Systems”. In: *Procedia Manufacturing* 11. 27th International Conference on Flexible Automation and Intelligent Manufacturing, FAIM2017, 27-30 June 2017, Modena, Italy, pp. 939–948. DOI: <https://doi.org/10.1016/j.promfg.2017.07.198>.
- Nguyen, Kim-Anh, Phuc Do, and Antoine Grall (2015). “Multi-level predictive maintenance for multi-component systems”. In: *Reliability Engineering & System Safety* 144, pp. 83–94. DOI: 10.1016/j.ress.2015.07.017.
- (2017). “Joint predictive maintenance and inventory strategy for multi-component systems using Birnbaum’s structural importance”. In: *Reliability Engineering & System Safety* 168, pp. 249–261. DOI: 10.1016/j.ress.2017.05.034.
- Nicolai, Robin P and Rommert Dekker (2008). “Optimal maintenance of multi-component systems: a review”. In: *Complex system maintenance handbook*, pp. 263–286. DOI: 10.1007/978-1-84800-011-7_11.

- Olde Keizer, Minou C.A., Simme Douwe P. Flapper, and Ruud H. Teunter (2017). “Condition-based maintenance policies for systems with multiple dependent components: A review”. In: *European Journal of Operational Research* 261.2, pp. 405–420. DOI: 10.1016/j.ejor.2017.02.044.
- Papadakis, Ioannis S and Paul R Kleindorfer (2005). “Optimizing infrastructure network maintenance when benefits are interdependent”. In: *OR Spectrum* 27.1, pp. 63–84. DOI: 10.1007/s00291-004-0167-5.
- Pargar, Farzad, Osmo Kauppila, and Jaakko Kujala (2017). “Integrated scheduling of preventive maintenance and renewal projects for multi-unit systems with grouping and balancing”. In: *Computers & Industrial Engineering* 110, pp. 43–58. DOI: 10.1016/j.cie.2017.05.024.
- Peng, Ying, Ming Dong, and Ming Jian Zuo (2010). “Current status of machine prognostics in condition-based maintenance: a review”. In: *The International Journal of Advanced Manufacturing Technology* 50.1-4, pp. 297–313. DOI: 10.1007/s00170-009-2482-0.
- Perini, Matteo, Paolo Bosetti, and Nicolae Balci (2020). “Additive manufacturing for repairing: From damage identification and modeling to DLD”. In: *Rapid Prototyping Journal* 26.5, pp. 929–940. DOI: 10.1108/RPJ-03-2019-0090.
- Petchrompo, Sanyapong and Ajith Kumar Parlikad (2019). “A review of asset management literature on multi-asset systems”. In: *Reliability Engineering & System Safety* 181, pp. 181–201. DOI: 10.1016/j.res.2018.09.009.
- Pinciroli, Luca et al. (2020a). “Agent-Based Modeling and Reinforcement Learning for Optimizing Energy Systems Operation and Maintenance: The Pathmind Solution”. In: *Proceedings of the 30th European Safety and Reliability Conference and the 15th Probabilistic Safety Assessment and Management Conference*. Ed. by Piero Baraldi, Francesco Di Maio, and Enrico Zio.
- Pinciroli, Luca et al. (2020b). “Deep Reinforcement Learning for Optimizing Operation and Maintenance of Energy Systems Equipped with PHM Capabilities”. In: *Proceedings of the 30th European Safety and Reliability Conference and the 15th Probabilistic Safety Assessment and Management Conference*. Ed. by Piero Baraldi, Francesco Di Maio, and Enrico Zio.
- Rausand, Marvin and Arnljot Høyland (2003). *System Reliability Theory: Models, Statistical Methods, and Applications*. Vol. 396. John Wiley & Sons. DOI: 10.1002/9780470316900.
- Rausand, Marvin and Jørn Vatn (2008). “Reliability centred maintenance”. In: *Complex System Maintenance Handbook*. Springer, pp. 79–108. DOI: 10.1007/978-1-84800-011-7_4.

- Rocchetta, R. et al. (2019). “A reinforcement learning framework for optimal operation and maintenance of power grids”. In: *Applied Energy* 241, pp. 291–301. DOI: 10.1016/j.apenergy.2019.03.027.
- Saaksvuori, Antti and Anselmi Immonen (2008). *Product Lifecycle Management*. Springer Science & Business Media. DOI: 10.1007/b138258.
- Savolainen, Jyrki and Mikael Collan (2020). “How Additive Manufacturing Technology Changes Business Models? – Review of Literature”. In: *Additive Manufacturing* 32, p. 101070. DOI: <https://doi.org/10.1016/j.addma.2020.101070>.
- Shafto, Mike et al. (2010). *Modeling, Simulation, Information Technology & Processing Roadmap*. URL: https://www.nasa.gov/pdf/501321main_TA11-MSITP-DRAFT-Nov2010-A1.pdf. (Accessed: 29-12-2020).
- Sheikhalishahi, Mohammad, Liliane Pintelon, and A Azadeh (2017). “An integrated approach for maintenance planning by considering human factors: Application to a petrochemical plant”. In: *Process Safety and Environmental Protection* 109, pp. 400–409. DOI: 10.1016/j.psep.2017.04.016.
- Shi, Hui and Jianchao Zeng (2016). “Real-time prediction of remaining useful life and preventive opportunistic maintenance strategy for multi-component systems considering stochastic dependence”. In: *Computers & Industrial Engineering* 93, pp. 192–204. DOI: 10.1016/j.cie.2015.12.016.
- Smith, Peter Godfrey (2003). *Theory and Reality: An Introduction to the Philosophy of Science*.
- Stahel, Walter R (2016). “The circular economy”. In: *Nature News* 531.7595, p. 435. DOI: 10.1038/531435a.
- Stricker, Nicole et al. (2018). “Reinforcement learning for adaptive order dispatching in the semiconductor industry”. In: *CIRP Annals* 67.1, pp. 511–514. DOI: 10.1016/j.cirp.2018.04.041.
- Sun, Bo et al. (2012). “Benefits and Challenges of System Prognostics”. In: *IEEE Transactions on Reliability* 61.2, pp. 323–335. DOI: 10.1109/TR.2012.2194173.
- Thomas, Douglas (2016). “Costs, benefits, and adoption of additive manufacturing: a supply chain perspective”. In: *The International Journal of Advanced Manufacturing Technology* 85.5, pp. 1857–1876. DOI: 10.1007/s00170-015-7973-6.
- Tuegel, Eric J et al. (2011). “Reengineering aircraft structural life prediction using a digital twin”. In: *International Journal of Aerospace Engineering* 2011. DOI: 10.1155/2011/154798.
- Uhlemann, Thomas H.-J., Christian Lehmann, and Rolf Steinhilper (2017). “The Digital Twin: Realizing the Cyber-Physical Production System for Industry 4.0”. In: *Procedia CIRP* 61. The 24th CIRP Conference on Life Cycle Engineering, pp. 335–340. DOI: 10.1016/j.procir.2016.11.152.

- Van Horenbeek, Adriaan and Liliane Pintelon (2013). "A dynamic predictive maintenance policy for complex multi-component systems". In: *Reliability Engineering & System Safety* 120, pp. 39–50. DOI: 10.1016/j.ress.2013.02.029.
- Vu, Hai Canh et al. (2014a). "Maintenance grouping strategy for multi-component systems with dynamic contexts". In: *Reliability Engineering & System Safety* 132, pp. 233–249. DOI: 10.1016/j.ress.2014.08.002.
- (2014b). "Maintenance planning and dynamic grouping for multi-component systems with positive and negative economic dependencies". In: *IMA Journal of Management Mathematics* 26.2, pp. 145–170. DOI: 10.1016/j.ress.2014.08.002.
- Wang, Hongzhou (2002). "A survey of maintenance policies of deteriorating systems". In: *European Journal of Operational Research* 139.3, pp. 469–489. DOI: 10.1016/S0377-2217(01)00197-7.
- Weber, Ron (Sept. 2003). "Editor's comment: Theoretically speaking". In: *MIS Q.* 27.3, iii–xii.
- Weller, Christian, Robin Kleer, and Frank T. Piller (2015). "Economic implications of 3D printing: Market structure models in light of additive manufacturing revisited". In: *International Journal of Production Economics* 164, pp. 43–56. DOI: <https://doi.org/10.1016/j.ijpe.2015.02.020>.
- Wenstøp, Fred (2010). "Operations research and ethics: development trends 1966–2009". In: *International Transactions in Operational Research* 17.4, pp. 413–426. DOI: 10.1111/j.1475-3995.2009.00730.x.
- White, Leroy (2009). "Challenge of research ethics committees to the nature of operations research". In: *Omega* 37.6. Ethics and Operations Research, pp. 1083–1088. DOI: <https://doi.org/10.1016/j.omega.2008.12.003>.
- Wildeman, R.E., R. Dekker, and A.C.J.M. Smit (1997). "A dynamic policy for grouping maintenance activities". In: *European Journal of Operational Research* 99.3. Eleventh EURO Summer Institute: Operational Research Models in Maintenance, pp. 530–551. DOI: 10.1016/S0377-2217(97)00319-6.
- Wireman, Terry (2004). *Total Productive Maintenance*. Industrial Press Inc.
- Wright, Louise and Stuart Davidson (2020). "How to tell the difference between a model and a digital twin". In: *Advanced Modeling and Simulation in Engineering Sciences* 7.1, pp. 1–13. DOI: 10.1186/s40323-020-00147-4.
- Xu, Li Da, Wu He, and Shancang Li (2014). "Internet of Things in industries: A survey". In: *IEEE Transactions on Industrial Informatics* 10.4, pp. 2233–2243. DOI: 10.1109/TII.2014.2300753.
- Xu, Li Da, Eric L. Xu, and Ling Li (2018). "Industry 4.0: state of the art and future trends". In: *International Journal of Production Research* 56.8, pp. 2941–2962. DOI: 10.1080/00207543.2018.1444806.

- Zhou, Xiaojun et al. (2015). “Preventive maintenance modeling for multi-component systems with considering stochastic failures and disassembly sequence”. In: *Reliability Engineering & System Safety* 142, pp. 231–237. DOI: <https://doi.org/10.1016/j.ress.2015.05.005>.
- Zio, E. (2009). “Reliability engineering: Old problems and new challenges”. In: *Reliability Engineering & System Safety* 94.2, pp. 125 –141. DOI: [10.1016/j.ress.2008.06.002](https://doi.org/10.1016/j.ress.2008.06.002).
- Zio, Enrico (2007). *An Introduction to the Basics of Reliability and Risk Analysis*. Vol. 13. World scientific.

Publication I

Urbani, M., Petri, D., Collan, M., and Brunelli, M.

Maintenance-management in light of manufacturing 4.0

In: *Technical, Economic and Societal Effects of Manufacturing 4.0: Automation, Adaption and Manufacturing in Finland and Beyond*

Ed. by Mikael Collan and Karl-Erik Michelsen.

Cham: Springer International Publishing, pp. 97–111

Reprinted with permission from

©2020, Springer



Maintenance-Management in Light of Manufacturing 4.0

*Michele Urbani, Dario Petri, Matteo Brunelli,
and Mikael Collan*

I INTRODUCTION

During the last decade, the manufacturing industry has gone through a deep transformation with the digitalization of processes, the arrival of the Internet of Things, the spread of artificial intelligence (AI) in daily practices, and the ubiquitous presence of data—thanks to the cloud technologies lifting the efficiency of manufacturing systems to a new level. Notwithstanding these radical changes, the manufacturing industry still has a strong dependence on maintenance, a field that is still considered to be a *necessary evil* by most managers, but without which plants and equipment will not remain safe and reliable. The importance of

M. Urbani (✉) • D. Petri • M. Brunelli

Department of Industrial Engineering, University of Trento, Trento, Italy

e-mail: michele.urbani@unitn.it; dario.petri@unitn.it; matteo.brunelli@unitn.it

M. Collan

School of Business and Management, LUT University, Lappeenranta, Finland

e-mail: mikael.collan@lut.fi

© The Author(s) 2020

M. Collan, K.-E. Michelsen (eds.), *Technical, Economic and Societal
Effects of Manufacturing 4.0*,

https://doi.org/10.1007/978-3-030-46103-4_5

maintenance-management as part of tangible asset-management is clearly inscribed within modern international industry standards [1], where asset-management is defined as “the coordinated activity of an organization to realize value from assets”. Maintenance-management takes care of physical assets with the aim of minimizing their life-cycle cost and achieving stated business objectives. Depending on the specific sector of industry, maintenance takes different forms—its most elementary form involves simple operations and inspections of and on machines, while the most cutting-edge applications include intelligent maintenance control-systems capable of predicting the remaining useful life (RUL) of components and triggering maintenance activities automatically when needed. Moreover, some companies are adopting more holistic approaches to maintenance, aimed at improving the efficiency of the whole productive unit. Such approaches are called *total productive maintenance* (TPM) [2] and they aim at improving the quality of products, developing corporate culture, and enhancing the attention to safety and environment.

The popularization of Industry 4.0 paradigm around the year 2011 represented a new starting point for the manufacturing industry after the financial crisis of 2008. Asset-management and maintenance-management of physical equipment underwent a transformation: real-time monitoring of working conditions became very common due to decreasing cost of sensor technology (IoT devices), thus making possible the development of new technologies such as Virtual Factories and Digital Twins (DTs) of machines and processes. The digital replication of the physical environment allows the optimization of processes already during the design phase and the optimization of running processes during the production phase. Real-time monitoring of assets and the direct control of processes remotely has become a part of the new paradigm of manufacturing; with respect to maintenance, diagnostics and prognostics of equipment are spreading into daily practices and a new stream of research is contributing to the development of these technologies.

In this chapter we illustrate some of the connections between modern manufacturing (Manufacturing 4.0) and maintenance-management, present shortly the evolution of maintenance-methodologies starting from early models until today and summarizing the most important concepts relevant to the field including a discussion of how the digital twin concept may become an important issue for maintenance-management.

2 MAINTENANCE-MANAGEMENT: AN OVERVIEW

Maintenance-management is nowadays a fundamental function in most industry. In its traditional form, maintenance is aimed at ensuring that a system performs its function in a safe and efficient manner. Due to information technology (IT) development, maintenance-management has seen a significant evolution within its best practices: the classical methods for maintenance-planning and scheduling have been integrated and improved by technologies such as the Internet of Things, cloud computing, and artificial intelligence.

Engineering systems often have a complex structure, with a limited number of dedicated resources and strict requirements on safety and on performance—under these circumstances maintenance is an issue that needs to be handled in a systematic way. A clear strategy for maintenance must be defined, where components of a system to be maintained should be documented and listed according to priority, then a set of rules for the daily management of operations must be drafted. The set of rules that are used to coordinate maintenance tasks are typically called a *maintenance-policy*. As basic example, maintenance-policies for lifts and elevators that typically depend on country-wise regulations and that state that maintenance must be carried out on regular intervals, such as “every twelve months”, which is then the rule that triggers a maintenance intervention that is aimed at avoiding sudden failures of the system. The above discussed types of interventions that are carried out before a failure has taken place are called *preventive* and they may range from simple inspections to the replacement of broken components. Maintenance actions undertaken after a failure are called *corrective* and they typically consist of the replacement and/or the repair of failed components. Usually corrective actions are more expensive than preventive, but when this is not the case it is sometimes possible to let a system *run to failure* that is, a system is left un-serviced until it fails, or until it fails and its failure is detected. Non-critical system components with a steady failure rate are often let run to failure.

Implementing preventive maintenance-policy typically requires more in terms of analysis, than a corrective policy—it requires information about the state of the maintained system such as information about the degradation level of system components. Depending on the information available, preventive maintenance-policies can be time-, or condition-based.

Time-Based Maintenance

Time-based or predetermined, as they are also called, maintenance-policies were the first approach adopted to effectively manage maintenance. In these types of policies maintenance actions are scheduled to take place on predefined times, according to set intervals of duration t_M , or upon failure (whichever occurs first). The aim of the policies is to preventively maintain the asset through shorter, but planned downtimes and by doing so avoiding longer and more expensive corrective maintenance actions. In this way the asset availability increases and consequences of failure can most often be avoided.

Scheduling of activities can be organized according to block-based- or age-based approaches. Block-based approaches schedule maintenance actions at constant time intervals, regardless of the asset operating time. The block-based approach is commonly used, when several assets of the same class (a block) are in (constant) use simultaneously. Age-based, or runtime, models are applied, when asset degradation and failures depend on the cumulative load exposure. Since the active age of a mechanical component has a strong correlation with the physical wear, or fatigue, of a component the maintenance of mechanical systems is often managed according to the age of system components. Asset age can be measured by using the working time of a machine as proxy, or in other ways, such as by observing the number of kilometres travelled or by the number of take-offs or landings, as can be done with aircraft. Approaches that combine more than one proxy for component states are also possible. Literature is ripe with research on time-based approaches for maintenance-optimization, we refer the interested reader to see the review by Wang [3]. It is worth to mention that time-based maintenance-policies carry a risk of over-maintenance, as some of the performed actions may not be necessary, on the other hand, time-based policies cannot weed-out failures, when component-deterioration happens at a non-standard pace—these are clear handicaps, when compared to condition-based policies. In fact, when the cost-risks of a time-based policy, or the costs of over-maintenance, are too high, condition-based maintenance may represent a feasible alternative.

Condition-Based Maintenance

Experience shows that failures can occur independently of the asset age, but at the same time most of these undesired events give some sort of

warning about the fact that they are about to occurring—thanks to the presence of such symptoms an early detection of fault occurrence is possible. This means that preventive actions can be taken, if the signals and symptoms of impending failures are understood, this is the fundamental concept that underpins condition-based maintenance. According to condition-based maintenance-policies maintenance actions are initiated by performance of a system reaching a trigger-level, typically determined by monitoring one or more indicators (sensors) of the maintained system. This means that maintenance is not done based on a predetermined schedule, but actions are taken based on observed, evidence-based deterioration of system performance that signals impending (component) failure and as such on only-when-needed basis.

A prerequisite for condition-based maintenance-policies (CBM) is that there is objective *monitoring* of the system state in place—the monitoring should be carried out in a non-invasive way and it is typically achieved by using sensors. Monitoring can be scheduled or continuous and the output from monitoring is a set of observations (indicators, failure precursors) that describe the capacity of a system to perform its function. A typical example of a failure precursor is the vibration frequency of a rotating machine—shift in the frequency is a clear indication of a change in the working conditions. As a rule of thumb used in CBM, once enough data has been gathered, thresholds on the monitored feature-values are established to more reliably identify degraded asset performance—a comparison between the system-state and the thresholds is used to track the system health. With knowledge about the system health and history-based thresholds a decision about maintenance-scheduling can be made in a way that actions are performed only when needed and as a result both the probability of failure and the overall cost of maintenance can be optimized.

3 MORE ABOUT CONDITION-BASED MAINTENANCE

Setting up condition-based maintenance is a process and it can be divided roughly into three main steps. Condition-based maintenance assumes that objective monitoring of the system is possible, which means that *acquisition of data* about the system state is in place. Sensors that measure issues such as material cracking, corrosion, vibration, and change in electrical resistance are the types of information that are usable from the point of view of understanding the system state—one must also remember that these issues depend on the operating and the environmental conditions,

such as the frequency of use, ambient temperature, and humidity. It is typical that a monitored system must be equipped with sensors, signal conditioning and digitizing components that are typically already embedded in new modern machines. We emphasize the importance of sensors, because they are a core technology needed for the implementation of the Manufacturing 4.0 paradigm in maintenance—they are the bond that connects machines into networks and they allow the realization of the Internet of Things.

Based on the data collected the features that explain and describe the state of the system and allow determining whether maintenance is necessary must be estimated. Features can be difficult to observe directly (by observing the system), but by exploiting data and a priori knowledge of the system *feature extraction* can be made easier. The quality of a feature is determined by its capacity to represent the system state, in order to achieve a better state representation, usually a set of features is used—the more clearly different system states can be distinguished from each other the better the condition of the system can be described. In practice finding the correct features or sets of features that allow high failure detection capability and a low false alarm probability are problems that can be solved by specific methods created for feature-selection and for information fusion. Improvement in feature-selection methods has been fuelled by the great interest analytics and AI have received in recent times. One must remember that sudden changes in the operative and environmental conditions may render features that work well under normal conditions imprecise—this is why the best modern systems may use different sets of features for different operating conditions and are able to change the feature sets used “on the fly”, when conditions change.

Once the data acquisition and feature extraction processes are ready *condition monitoring* can be effectively performed. Monitoring is the last step prior to the definition of the maintenance-strategy that is forming the set of rules that aids managers in taking maintenance decisions.

The main goal of condition monitoring is to provide fault-recognition, which typically foresees three sub-goals: (1) fault detection, aimed at identifying if a fault or the degradation of a component occurred; (2) fault isolation, that identifies the damaged component among many others; and (3) fault identification, aimed at determining the nature, extent, and severity of the isolated fault. In the following we look at these issues in more detail.

Fault Detection

The task of fault detection is to identify the presence of abnormal working conditions in a system by leveraging the information from the system history and information that can be learned from actual data. Typically a benchmark that defines the “normal” working conditions of the system is needed—the normal conditions depend on the task that the system is carrying out and on the environment surrounding the system. Because of different environments a system may have several normals—each normal will have a “profile” that is a set of features that defines it. Another thing is the extraction of profiles for different fault-states, such as “healthy”, “degraded”, and “faulty”. The state of the system can be compared to the different profiles and this allows one to understand the state of the system and to predict the failure. Typically one will want to see several system states that precede the “failed” state, because the more states there are the finer is the information about the system state and better one can predict what will happen next. The comparison of the observed system state and the normal state can be done by different means, two examples of usable modelling techniques for this purpose are the auto associative kernel regression (AAKR) [4] and principal component analysis (PCA) [5] for the identification of the state and subsequently a statistical test is applied to identify the extent to which the state of the system differs from a normal condition. Typically used tests include the threshold based approach, Q statistics, and the Sequential Probability Ratio Test (SPRT) [6]. When the state of the system is known an action is taken (not taken) depending on the recommendations described for each state—the recommendations are drafted by using fault diagnosis techniques.

In order to clarify how fault detection works, we provide a simple example of condition monitoring. We assume that the state of a system is represented by a single feature $x(t)$. We define two thresholds considered important for the component. In Figure 1, the first threshold x_w identifies a warning-level, while threshold x_f identifies the failure of the component. When the value of $x(t)$ surpasses level x_w , an alarm is triggered, and a preventive action can be undertaken to prolong the life of the component, or to change it, to avoid incurring a sudden failure. The curve representing the behaviour of $x(t)$ is known as the Performance/Failure curve and it expresses the evolution of the system-feature as a function of either calendar time or system age time. A realistic mathematical model of $x(t)$ will also include the uncertainty related to the estimated quantity, which in

Fig. 1 A performance/failure curve for a generic system

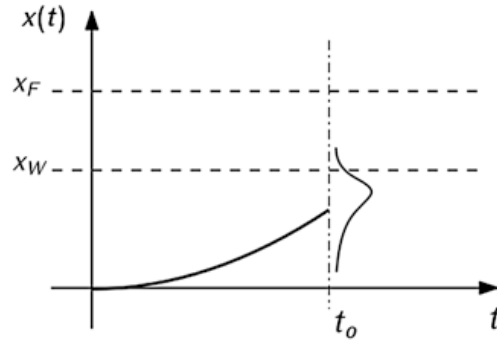


Fig. 1 is represented by a generic probability distribution. The importance of modelling the degradation of a component using a random variable is represented by the possibility to express the result using a probability that is, a degree of belief about the triggering event.

Fault Diagnosis

Fault diagnosis is isolating and identifying the fault and typically means identifying the cause, this means identifying which component in a system is degrading among many possible components and to determine the nature, the extent, and the severity of the fault. Isolating and identifying the fault are sometimes overlapping and not always clearly separable. Fault diagnostics means most often solving a classification problem—any given set of measurements from the system can be matched to a single component if sufficient data is available for training a machine learning classification algorithm. In cases where data is abundant algorithms can even spot specific conditions within components and provide a credible probability of a failure event. Many techniques are good for this task, the interested reader may find an extensive review about modern fault diagnostics techniques applied to rotatory machines in [7], where the authors describe both the fundamental principles behind adopted AI algorithms and present numerous application examples. As a caveat about AI-based techniques one must observe that where there is no data, or data is very incomplete, machine learning algorithms cannot be used—in such cases suitable data must first be collected. In the cases of very rare faults diagnosis is difficult and diagnostics performance for them is typically poor.

The performance of condition-based maintenance systems is only as good as the system in place and there is uncertainty associated with the

outputs (alarms) from these systems. Uncertainty is caused by a number of things, some were already mentioned above such as the operating conditions and the environment, but others like production tolerances also affect the reliability of CBM system—because of tolerances two nominally identical machines may have a different wear. Due to this inherent inaccuracy the output from CBM systems is most often expressed as a probability or an interval. We refer the reader interested in deepening their knowledge in maintenance and maintenance optimization to read the review by De Jonge and Scarf [8].

4 PROGNOSTICS AND HEALTH MANAGEMENT—TOWARDS INDUSTRY 4.0

Thanks to the availability of cheap networked sensors the monitoring and maintenance of systems is undergoing a fast and deep change. In the past, manual collection of maintenance-relevant data made the processing slow and unreliable—today technology allows abundant collection of data often in real-time. This profound change has caused the attention of maintenance systems development to move towards maintenance process-optimization. The new generation of production systems that are “smart” and networked has been labelled as Cyber Physical Production Systems (CPPS)—important to maintenance, they offer the possibility to perform real-time monitoring and accurate analysis of the degradation of critical components. This means that the long stream of research carried out on condition-based maintenance can now be exploited for its full potential—this change has given rise to the term Prognostics and Health Management (PHM), which can be said to be the cutting-edge approaches to predictive maintenance born within the last two decades. Keeping in mind that PHM is part of the same continuum with CBM and that the two cannot be sharply separated, it can be said that PHM aims higher than the “traditional CBM” and uses more advanced tools to get there.

The higher goals of PHM include, for example, optimization of maintenance-planning, reduction of downtimes, just in time spare parts provision, energy consumption optimization, minimization of raw material use and of pollution—all in all the focus is on increasing profitability through “better maintenance”. PHM means effectively the same thing that is meant when the term Predictive Maintenance is used in common parlance. A fundamental prerequisite for a well-functioning predictive

maintenance system is the high quality of information that is used as an input into the system. This is true for both the real-time operation of the system as it is true for the information that is needed to construct or teach the system to be able to operate reliably—the information needed typically includes operating and maintenance histories, prior knowledge about system failure modes, resource constraints, and mission requirements. The information is used in tuning complex models the architecture of which may include numerous machine learning sub-systems and that require top of the line know-how. This means that these systems are expensive and they can be constructed only for systems that either merit such costs from the point of view of safety or that are business-critical and can economically justify the expenses.

In prognostics and health-management systems the system status received as input from condition monitoring is used to create an estimate of the system degradation state, which is used together with the P/F curve, or by using a classification-based architecture, to determine the distance between the current degradation level and a failure threshold (health-margin). The idea of the modern systems is to not only identify the cause of the fault but also to predict any secondary failures that may occur and to forecast the system health evolution as reliably as possible. Prognostics is considered the “holy grail” of PHM systems [9], because diagnostics has a retrospective approach to failure that consists of identifying and quantifying failures that have already occurred, while prognostics is about forecasting and as such, if successful means that the remaining useful life (RUL) of components can be accurately predicted. This will happen simply by being able to accurately estimate the end of life of a component and calculating the time to the end of life—the more accurate this ability is, the more precise can any optimizations performed based on it, including just in time deliveries of spare parts and maintenance scheduling become. The difference between high accuracy and medium accuracy can mean great savings in cases, where multiple systems are maintained and costs associated with maintenance are high. Another important issue is to know how much in advance a prognostics system can (accurately) predict the failure time—in fact, the relative RUL estimation accuracy and the prognostic horizon are key performance parameters of PHM systems.

In the literature, three types of approaches to prognostics have been identified, namely (1) experience-based approaches, which exploit historical information of a similar components; (2) model-based approaches, which make use of a physical fault model, and; (3) data-driven approaches,

which are mainly based on AI-techniques. We propose the interested reader to explore model-based and data-driven approaches by reading the book by Kim et al. [10].

Digital Twins and Their Connection to Maintenance

According to recent literature on maintenance and industrial management [11, 12] prognostics and health management systems be viewed as an examples of cyber-physical systems (CPS). The idea of CPS started to spread in the beginning of the 2010's, when NASA published their Modelling, Simulation, Information Technology & Processing Roadmap [13]—the document delineated the intention to integrating all the available physical and virtual technologies, the context back then was aeronautics. In essence the idea is that of a digital replica of a physical asset and it was called a Digital Twin (DT) and defined as “*an integrated multi-physics, multi-scale, probabilistic simulation of a vehicle or system that uses the best available physical models, sensor updates, fleet history, etc. to mirror the life of its flying (sic) twin. It is ultra-realistic and may consider one or more important and independent vehicle systems*”. What makes this interesting from the point of maintenance is that predictive maintenance was one of the first fields of application of the DT concept, together with the check of mission requirements and a more transparent life-cycle view. The DT concept was subsequently extended to the manufacturing industry and the term Cyber-Physical Production System (CPPS) was coined to indicate the specific application area. A CPPS is composed of a physical part, a virtual part (the DT), and a stream of data between the two [14]. The DT strives to hold a perfect real-time synchronization between the physical and the virtual worlds, the physical part sends data to the virtual model, and the virtual part reproduces the physical system with ultra-high fidelity. As this is the case, historical data stored can be used together with real-time sensory information from the physical system in order to run, e.g., simulations and to optimize the production process virtually and then transmit “orders” to the physical system in order to optimize the way it functions. Theoretically the CPPS can harness the interaction between the virtual and the physical parts in order to create a continuously improving system. Digital twins are a clear way to remedy the typical problems of data collection, organization, and exploitation widespread in the context of production systems.

In fact, digital twins start to look like the key to reaching solutions for the problems of fitting together the best practices in engineering design and in process control. The advantages of adopting the DT concept seem cover the whole of product lifecycle that is, production design, manufacturing, and service providing are all immersed in the realm of DT [14]. In the design phase, if realized with a sophisticated digital model, issues that have to do with the maintainability of the production system can perhaps be addressed already on the drawing board—this may include the instrumentation of the system for best possible diagnostics and prognostics. During the production life of the production system the DT can perhaps assist in production planning, resource management, and procurement that can be optimized also with regards to predicted downtimes due to maintenance. The DT may run failure prediction algorithms in real-time so that users can be notified when the system state changes and in cases of imminent failure. It seems feasible to say that there is clear potential for maintenance systems development based on the digital twin concept.

5 CONCLUSION

Maintenance has always been a part of the management of production systems and it has become a craft of its own, the early mathematical models for maintenance management were based on the notion of optimizing the interval between maintenance activities in order to minimize downtime and the maintenance related costs. This type of maintenance management systems may still exist in cases, where preventive maintenance is the norm and the systems maintained are “old school” and not instrumented with sensors.

The modern approach of maintenance management is based on condition-based maintenance, which in the early days was more expensive than time-based maintenance management and thus reserved to high-risk and high-cost applications. Today the price of sensors and instrumentation is considerably low, which has made condition-based maintenance the leading way of handling maintenance management. Improvement of maintenance policies has created competitive advantages for companies that have been able to adopt them successfully and therefore a shift to modern maintenance management approaches is occurring in many companies. Automation of industrial facilities, such as the increasing use of robotics, improves productivity and safety, but it also increases the technological complexity of industrial assets and means a higher dependence on

production systems—this accentuates the role of effective and efficient maintenance.

Key Industry 4.0 technologies, such as artificial intelligence and Internet-of-Things, enable the implementation of very effective maintenance policies at an affordable cost and have paved the way for better diagnostic and prognostic systems, which can be said to be the backbone of what is typically called predictive maintenance. These systems are able to make fault-prediction even more accurate than what is possible with traditional condition-based maintenance methods and therefore offer a possibility for even further savings through better optimization. Predictive maintenance most importantly is a forward looking approach to maintenance, where traditionally the policies have been based on after-the-fact optimization.

The concept of digital twin is interesting from the point of view of maintenance management, as it is based on the idea of having a highly accurate real-time virtual model of a physical system that are “conversing” with one another. In effect, this is a concept that is not very far away from the ideal maintenance management system in terms of the information exchange between a production system and the maintenance management system. The digital twin, as it is used in the lifecycle management of products today is already opening avenues for many issues that are relevant to making maintenance better—looking forward there is potential for much more, specifically in terms of using digital twins in a maintenance focused way.

Getting back to the real-world, one must observe that the choice of maintenance management systems and policies is always constrained by the economic and technical realities surrounding the maintained systems. In this respect, predictive maintenance is at the start of a road that may lead at some point to something that resembles a digital twin—one thing is for sure, the Industry 4.0 paradigm and what we already can see beyond it will change maintenance management.

REFERENCES

1. ISO 55000, “Asset Management—Overview, Principles, and Terminology.” 2014.
2. S. Nakajima, *Introduction to TPM: Total Productive Maintenance*. Productivity Press, Inc., 1988.

3. H. Wang, "A survey of maintenance policies of deteriorating systems," *Eur. J. Oper. Res.*, vol. 139, no. 3, pp. 469–489, 2002.
4. P. Baraldi, F. Di Maio, P. Turati, and E. Zio, "Robust signal reconstruction for condition monitoring of industrial components via a modified Auto Associative Kernel Regression method," *Mech. Syst. Signal Process.*, vol. 60, pp. 29–44, 2015.
5. P. Baraldi, G. Gola, E. Zio, D. Roverso, and M. Hoffmann, "Expert Systems with Applications A randomized model ensemble approach for reconstructing signals from faulty sensors," *Expert Syst. Appl.*, vol. 38, no. 8, pp. 9211–9224, 2011.
6. F. Di Maio, P. Baraldi, E. Zio, and R. Seraoui, "Fault detection in nuclear power plants components by a combination of statistical methods," *IEEE Trans. Reliab.*, vol. 62, no. 4, pp. 833–845, 2013.
7. R. Liu, B. Yang, E. Zio, and X. Chen, "Artificial intelligence for fault diagnosis of rotating machinery: A review," *Mech. Syst. Signal Process.*, vol. 108, pp. 33–47, 2018.
8. B. de Jonge and P. A. Scarf, "A review on maintenance optimization," *Eur. J. Oper. Res.*, 2019.
9. J. B. Coble and J. W. Hines, "Prognostic algorithm categorization with PHM challenge application," *2008 Int. Conf. Progn. Heal. Manag. PHM 2008*, 2008.
10. N.-H. Kim, D. An, and J.-H. Choi, *Prognostics and Health Management of Engineering Systems*. Springer, 2017.
11. E. Negri, L. Fumagalli, and M. Macchi, "A Review of the Roles of Digital Twin in CPS-based Production Systems," *Procedia Manuf.*, vol. 11, no. June, pp. 939–948, 2017.
12. W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, "Digital Twin in manufacturing: A categorical literature review and classification," *IFAC-PapersOnLine*, vol. 51, no. 11, pp. 1016–1022, 2018.
13. NASA, "Modeling, Simulation, information Technology & Processing Roadmap—Technology Area 11," *Natl. Aeronaut. Sp. Adm.*, p. 27, 2010.
14. F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, and F. Sui, "Digital twin-driven product design, manufacturing and service with big data," *Int. J. Adv. Manuf. Technol.*, vol. 94, no. 9–12, pp. 3563–3576, 2018.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



Publication II

Urbani, M., Brunelli, M., and Collan, M.

**A comparison of maintenance policies for multi-component systems through
discrete event simulation of faults**

Reprinted with permission from

IEEE Access

Vol. 8, pp. 143654–143664

©2020, IEEE

A Comparison of Maintenance Policies for Multi-Component Systems Through Discrete Event Simulation of Faults

MICHELE URBANI^{1,2}, MATTEO BRUNELLI¹, AND MIKAEL COLLAN², (Associate Member, IEEE)

¹Department of Industrial Engineering, University of Trento, 38122 Trento, Italy

²School of Business and Management, Lappeenranta-Lahti University of Technology (LUT) University, 53850 Lappeenranta, Finland

Corresponding author: Michele Urbani (michele.urbani@unitn.it)

This work was supported in part by the Italian Ministry of Education, University and Research (MIUR) through the Departments of Excellence Program, and in part by the Finnish Strategic Research Council under Grant 313396 MFG40.

ABSTRACT Finding optimal maintenance policies for complex multi-component systems is a real-world challenge in the industry. This paper compares three maintenance policies for complex systems with non-identical components and economic dependencies in case of fault. Discrete event and Monte Carlo simulation are used to replicate fault occurrences, while a genetic algorithm is used to minimize the cost of maintenance by finding optimal groups of maintenance activities. Low total average maintenance cost and high average availability of the system are considered as desirable objectives and the capacity of the studied policies to achieve these goals is analyzed. None of the policies dominates the others (in a Pareto efficiency sense), thus making the policy choice context dependent and subject to decision makers' preferences.

INDEX TERMS Maintenance policies, simulations, genetic algorithm, opportunistic maintenance.

I. INTRODUCTION

This paper studies a set of maintenance policy alternatives available for managing the maintenance of complex systems—policies with and without grouping of maintenance activities are considered. More specifically, three possible real-world maintenance policies are studied and (dynamically) optimized for simulated maintenance schedules that include simulated occurrences of fault events; subsequently, the resulting cost of maintenance and the observed reliability of the system are recorded for each simulation and for each policy, and used to compare the policies.

Besides presenting results, our goal is to show that this approach is sufficiently holistic and general to be used in aiding industry decision making on maintenance policy selection and to illustrate the real-world applicability of the methods. The possibility to test a maintenance policy for suitability in advance is a substantial improvement to the decision making process connected to choosing a maintenance policy, a task typically carried out by a maintenance department in an industrial company.

The associate editor coordinating the review of this manuscript and approving it for publication was Chaitanya U Kshirsagar.

The need to choose a suitable maintenance policy is especially pressing in the context of industries with high setup and downtime costs due to maintenance operations. A few examples of high downtime cost industries are the oil and gas industry, where reaching offshore extraction facilities is costly; the steel making industry, where the shutdown of blast furnaces requires long times and causes large losses of material; and the production of pulp and paper, where the cost of missed production is high enough to justify a 24/7 operation. Other examples are electrical networks [18] and manufacturing systems [19].

The result of choosing a maintenance policy is also practical: each policy produces a set of actions, or a *maintenance schedule*, that can be implemented in practice.

A. MOTIVATION

Finding a good maintenance policy is a challenge of primary importance for many production-based industries running physical production assets. Some authors [5], [27] claimed that maintenance costs in the industry could range between 15 and 70 % of total production costs. It is then clearly in the best interest of these organizations to try to minimize maintenance related costs and to maximize the reliability

and availability of their machines. Previous research in the field of maintenance policy research has produced a rather wide range of policies for the management of maintenance of complex multi-component systems [8], [9], [30], [45]. These theoretical policy-models typically optimize the maintenance schedule with regards to several objectives and are able to integrate short-term information on system status. A sudden fault is a typical short-term occurrence (information) which negatively influences the performance of a machine and of a maintenance-system. It may be impossible to fully insulate a system against sudden faults, but having a maintenance policy in place that is able to minimize the costs from a sudden fault can work to decrease the associated costs. The question is then about how to choose a good maintenance policy. By knowing the lifetime distributions of system components, the failure process can be quite accurately replicated, and the performance of different maintenance models can be estimated in advance. Using a simulation-based approach allows the effective comparison of different maintenance policies and is a suitable tool for analyzing them [4]. Simulation models are also able to consider the dependencies between components. Modeling of economic dependencies, such as (high) setup costs of activities, is of fundamental importance for systems with series of components.

The ability to intelligently group maintenance activities, when failures take place, increases the ability of an organization to minimize maintenance related costs. In addition to the optimization results any further exploitation of statistics from simulation-based analyses may help managers to obtain additional insight in the reliability of a system and on the robustness of a maintenance policy.

B. STATE OF THE ART

Manufacturing systems are increasingly complex and their effective maintenance is a challenge for maintenance managers and researchers. The complexity of the resulting models is often on such a level that analytical solutions for optimal maintenance schemes are seldom available [30] and simulation-based approaches are often used.

The literature on the topic proposes a great number of different models to study and create policies for the maintenance of single- and multi- component systems. Reviews and a classification of existing models were done by Cho and Parlar [8], and by Dekker *et al.* [9]. Both reviews agree on categorizing maintenance models into five groups; out of these five, four are of interest for this research: *group/ block/ cannibalization/ opportunistic models* aim at identifying the components that may be changed during preventive, or corrective, maintenance.

Component dependencies can be exploited in multi-unit systems to reduce maintenance costs. There are three types of dependence: structural, which identifies the possibility to maintain components independently [14]; stochastic, where failure of one component may influence the lifetime of other components [10], [22], [35], [40]; and economic dependence. Economic dependence is typically investigated to

establish whether it is possible to save on maintenance costs by contemporarily executing multiple activities, or if, instead, the execution of activities separately is economically more feasible. Although there is potential in considering the three dependencies together, in the literature they are usually considered separately [9]; only Van Horenbeek and Pintelon [42] presented an all-encompassing approach to model all types of dependencies. Maintenance models for multi-component systems with economic dependencies were exhaustively reviewed by Nikolai and Dekker [30].

Maintenance-models can be further classified as *static* or *dynamic*, depending on their ability to include (pieces of) short-term information about the status of the system. Static models are usually based on an infinite length planning horizon and they are devoted to optimizing the maintenance frequency of a component. A clear limit of static models is their inability to consider new information about unforeseen events. Other authors [9], and [45] presented reviews on static maintenance models.

Dynamic models are more flexible than static models: they exploit short- and long-term information together in order to combine corrective maintenance (CM) with preventive maintenance (PM) interventions. Dynamic models that combine interventions on different components of the same system are also known as opportunistic dynamic grouping models. They exploit component dependencies to defer activities from their initially scheduled dates and thus try to make savings on setup cost of activities. A cornerstone in dynamic grouping of maintenance activities is the work by Wildeman *et al.* [48], which was subsequently extended by several authors. Meaningful improvements of the model regarded the inclusion of health status and failure occurrence of components [6], criticality of components [43], [44], and multi-level condition based maintenance (CBM) [29]. One further improvement of the model consists of the addition of activities duration, which had previously been considered to be zero: Do Van *et al.* [12] added multiple maintenance activities with different durations, Pargar *et al.* [32] proposed grouping and balancing of activities, Sheikhalishahi *et al.* [39] accounted for human influence on quality of maintenance and illustrated it with a case study on an offshore oil plant in Iran.

Given the complexity of the dynamic grouping models under analysis, a simulation-based approach seems to be the most suitable approach to solve these types of problems [1]. Simulation-based models have been shown to be effective in bringing results in many industries, including semiconductor manufacturing, plastic industry, transportation infrastructure, and train maintenance facilities [3]. Alrabghi and Tiwari [2], [3] reviewed the literature on maintenance-system simulations and provided several examples, where Discrete Event Simulation (DES) was proficient in modeling fault occurrences.

The principle behind DES is easy to understand: each time the state of the system changes, the simulator applies the required changes to the system in accordance with the adopted maintenance policy (e.g., CM, PM, or CBM) and the

result with statistics is registered. DESs can be used to reach many kinds of results, e.g., to find the optimal capacity of inter-operational buffers which minimize the cost of a plant downtime [28], to optimize spare parts availability [20], [47], to estimate reliability of systems made of rotatable parts [13], to optimize thresholds in CBM policies [7], [16], to optimize maintenance intervals [31], and to develop knowledge for maintenance management [33].

The aforementioned DES models analyze a single policy at a time. Finding a suitable (optimal) maintenance-policy for a complex multi-component systems requires the comparison of policy alternatives—in the past, only limited efforts have been made to compare different maintenance policies under operative conditions [3]. There are few exceptions: Hani *et al.* [17] compared policies for train maintenance, Van der Duyn Schouten and Vanneste [41] for management of buffers in production systems subjected to maintenance, and Van Horenbeek and Pintelon [42] for multi-component systems dependencies on the components. Only [38] tested maintenance policies for flexible manufacturing systems, i.e., systems where wear out risk is higher than in standard systems, operating under different failure rates. One can say that the literature on comparing maintenance-policies is not complete.

An improvement in DES for maintenance, was provided by the framework of Alrabghi and Tiwari [4]. They designed a general procedure for DES with different policies (including CM, PM, and CBM), and this work will partly follow their footsteps. In order to deduce meaningful insights about the policies under analysis, our study combines DES in Monte Carlo experiment. Rao and Bhadury [36] showed how the comparison of opportunistic maintenance policies is possible by using the Monte Carlo technique. In addition, the complexity of the combinatorial problem pushed us to use a genetic algorithm (GA) to obtain satisfactory solutions in a reasonable time. While our research uses GA, we acknowledge that also other optimization methods can be used. Although the Monte Carlo method and a GA can provide useful insights on suitable (optimal) maintenance-policy identification, they have been the subject of only few publications in the past [23].

C. CONTRIBUTION

The body of literature is populated by several complex maintenance models, as shown in Section I-B. A multitude of optimization problems were tackled regarding cost minimization, or availability and reliability maximization. Only few simulation studies [17], [36], [41], [42] and a simulation methodology [4] are available to compare maintenance policies. This study differs from those already presented in the literature, thanks to the combination of tools that is used and by the methodology that is followed. In this research, a hypothetical industrial system is modeled taking into account the presence of multiple non-identical components connected in a series. Activities duration is also considered when the maintenance schedule is drafted, by using a genetic algorithm to optimize

the grouping structure. Cost minimization is the only objective, while the reliability of the system is considered for policy evaluation. An opportunistic maintenance policy similar to that of Wildeman *et al.* [48] against other heuristic policies is an element of novelty of this study, which, to best of our knowledge, has never been done before. The policies analyzed here can be considered realistic approximations of real-world maintenance needs. The numerical results are used to compare different maintenance policies, when the setup cost of the activities varies. The setup cost is the key factor which pushes the algorithm to group maintenance activities whenever possible. Descriptive statistics like the expected cost of each policy, the distribution of the variance of a policy's costs, and the average reliability of the system are calculated to compare the effectiveness and robustness of the studied policies.

The rest of this paper is organized as follows. Section II describes the model and the optimization technique used in the analysis and the determination of the lowest cost solution for each policy. Section III describes the simulation procedure for comparison of policies and summarizes the obtained results. Discussion, conclusions, and suggestions for further research are presented in Section IV.

II. THE MODEL

The model presented below is developed according to the five phases, rolling horizon approach, proposed by Wildeman *et al.* [48], with the addition of activities duration, which can be summarized in the following steps:

- 1) *Decomposition*: determine the optimal frequency for maintenance of each component separately; the planning horizon is considered to be of infinite length during this step.
- 2) *Penalty functions*: a penalty function is determined for each activity, and it is used to quantify the cost for deferring the activity from its ideal execution date. Activities can be shifted backward or forward in time.
- 3) *Tentative planning*: the duration of the plan is now considered finite and multiple maintenance activities are possible for each component.
- 4) *Grouping maintenance activities*: the maintenance activities are allowed to be moved within the planning horizon. The aim of this step is to maximize the save on set-up cost due to grouping of activities, and to minimize the cost due to shifting of activities.
- 5) *Rolling-horizon step*: once new short-term information is available, it can be supplied to the model and the model can be executed again to obtain an optimized maintenance schedule.

We consider a multi-component system with N components connected in series, which means that all the components are considered critical; Namely, a fault of one component compromises the whole system. The choice to analyze a series system reflects the approach of previous studies on the opportunistic maintenance policy [24], [49], [50], where

the criticality of each component makes the opportunistic approach particularly effective.

A. THE COST STRUCTURE

We assume that only two types of maintenance activities are possible: (i) preventive maintenance activities and (ii) corrective maintenance activities; the first are considered to be planned activities, whereas the second are unplanned. Both these activities, in terms of costs, are treated as the sum of three factors: a set-up cost, a cost for replacement of the component, and a cost for missed production, this is similar to what was used by [44]. The cost of a *preventive* maintenance activity i is:

$$C_i^p = S + c_i^p + C_{sys}^p d_i, \quad (1)$$

where S is the set-up cost, c_i^p is the cost for replacement of the component, and $C_{sys}^p d_i$ is the cost of missed production, which is calculated as the product of the coefficient C_{sys}^p [\$/time] and the duration of the i -th activity d_i . On the other hand, the cost of an *unplanned* maintenance activity i is:

$$C_i^c = S + c_i^c + C_{sys}^u d_i, \quad (2)$$

where S is the set-up cost, c_i^c is the cost for replacement of the component, and $C_{sys}^u d_i$ is the coefficient for unplanned missed production.

B. DECOMPOSITION

The model we use is the dynamic grouping maintenance model with the opportunistic approach proposed by Wildeman *et al.* [48], with some changes. The time to failure of each component is considered as a random variable X_i , and its probability of occurrence before time t is characterized by a two parameter Weibull distribution with the cumulative density function (CDF):

$$\Pr\{X_i \leq t\} = F_i(t) = 1 - \exp\left(-\frac{t}{\lambda_i}\right)^{\beta_i}, \quad (3)$$

with scale parameter $\lambda_i > 0$, shape parameter $\beta_i > 1$, and probability density function (PDF):

$$f_i(t) = \left(\frac{\beta_i}{\lambda_i}\right) \left(\frac{t}{\lambda_i}\right)^{\beta_i-1} \exp\left(-\frac{t}{\lambda_i}\right)^{\beta_i}, \quad (4)$$

As recalled in the literature [26], Weibull distributions are sufficiently general to fit a wide range of empirical distributions, and the lower bound imposed on the shape parameter ($\beta > 1$) implying increasing failure rate is not restrictive for our analysis. In fact, with $\beta_i \leq 1$ implies a non-increasing failure rate which, in turn, makes preventive maintenance activities on single components unreasonable. Moreover, with today's high quality standards, infant mortality of components is often a negligible phenomenon and, as claimed by Love and Guo [26], "most often a rising force of mortality is assumed".

At this point, the mathematical treatment to obtain the optimal interval length for preventive maintenance x_i^* is contained in [48] and the full presentation is thus omitted here.

However, it has been shown that, if the duration of a PM activity is $d_i \ll x_i^*$, then x_i^* can be approximated as:

$$x_i^* = \lambda_i \sqrt[\beta_i]{\frac{(C_i^p + S)}{C_i^c (\beta_i - 1)}}. \quad (5)$$

Eq. (5) allows us to compute the value of the minimal long-run average cost of maintenance per unit time:

$$\phi_i^* = \phi_i(x_i^*) = \frac{(C_i^p S) \beta_i}{x_i^* (\beta_i - 1)}. \quad (6)$$

The minimal long-run average cost of maintenance per unit of time for the whole system can be calculated as the sum of all these costs for all the components:

$$\phi_{sys}^* = \sum_{i=1}^N \phi_i^*. \quad (7)$$

C. TENTATIVE MAINTENANCE

A finite length planning horizon is now considered in order to realize the grouping of activities. The initial time of the plan is t_{begin} , whereas the date of the last maintenance action on component i is t_i^e ($\leq t_{begin}$). The cumulative duration D_i^{\sum} of all the replacement activities between t_i^e and the first activity on i is used jointly with the length of the preventive maintenance cycle x_i^* in order to determine the date t_{ij}^* of the first repair action $j = 1$ on component i . The date t_{i1} can be calculated by using the following equation:

$$t_{i1} = t_{begin} - t_i^e + d_i + D_i^{\sum} + x_i^* \quad (8)$$

where d_i is the duration of the maintenance activity on component i . Instead, the end date of the planning horizon t_{end} is equal to a multiple of:

$$t_{end} = \max_{i=1, \dots, N} (t_{i1}) + d_i. \quad (9)$$

It is important to note that the maintenance activity of each component i might be executed more than once within the interval $[t_{begin}, t_{end}]$, therefore the maintenance dates of activities with $j \geq 2$ are calculated as follows:

$$t_{ij} = t_{ij-1}^* + d_i + D_i^{\sum} + x_i^* \quad \forall t_{ij} \leq t_{end}, \quad (10)$$

where t_{ij-1}^* is the optimal execution date of the previous maintenance activity on component i , D_i^{\sum} is the cumulative duration of the preventive maintenance (PM) activities within the interval $[t_{ij-1}^*, t_{ij}]$. The process is represented in a simplified manner in Fig.1.

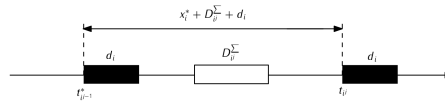


FIGURE 1. Representation of how the dates of the activities are calculated.

D. GROUPING MAINTENANCE ACTIVITIES

The economic dependence among the components of the system is a key variable of the optimization model and it is used to produce savings on maintenance costs. Savings on setup costs are generated when maintenance activities are executed in the same moment; Namely, one activity is subsequent or simultaneous to the other, and it is equal to:

$$U_{G^k} = [|G^k| - 1] S, \quad (11)$$

where $|G^k|$ is the cardinality of the group, namely the number of activities simultaneously executed. The higher the number of activities in G^k the greater the savings. The shifting of maintenance activities from their ideal date leads to costs. Suppose that the activity i^j is shifted from its ideal date t_{ij} , to the date t_{G^k} of the group it belongs to: the new date of execution is equal to $t_{G^k} = t_{ij} + \Delta t_{ij}$, where the temporary shift Δt_{ij} can be “positive” or “negative”, i.e., the activity can be anticipated or postponed. In order to avoid infeasible shifts of activities, the following constraint is imposed: $\Delta t_{ij} > -x_i^*$. In order to quantify the cost of activities shifting penalty functions are introduced. The change of date of an activity i^j has effect on the following activities on component i , which are moved using a long term shift; namely, the interval between the first two maintenance activities becomes $x_i^* + \Delta t_{ij}$, whereas the remaining intervals remain x_i^* . The process is graphically represented in Fig. 2. Once this choice has been made, the penalty function for each activity is composed of two parts:

- 1) an increase of the expected cost with regards to the j -th renewal cycle, which is given by $E_i(x_i^* + \Delta t_{ij}) - E_i(x_i^*)$,
- 2) and a changing cost due to the deferments of the future activities executed after t_{ij} , which is calculated as $\Delta t_{ij} \phi_i^*$.

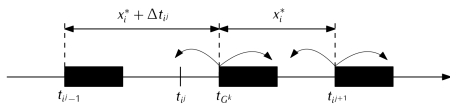


FIGURE 2. The rectangles represent activities on the time axis; with a long term shift all the activities after t_{ij} are moved accordingly to t_{ij} .

Therefore, a penalty function $h_i(\Delta t_{ij})$ for the shifting of activity i^j on component i can be expressed as:

$$\begin{aligned} h_i(\Delta t_{ij}) &= E_i(x_i^* + \Delta t_{ij}) - E_i(x_i^*) - \Delta t_{ij} \phi_i^* \\ &= \underbrace{\left[C_i^p + C_i^c \left(\frac{x_i^* + \Delta t_{ij}}{\lambda_i} \right)^{\beta_i} \right]}_{E_i(x_i^* + \Delta t_{ij})} \\ &\quad - \underbrace{\left[C_i^p + C_i^c \left(\frac{x_i^*}{\lambda_i} \right)^{\beta_i} \right]}_{E_i(x_i^*)} - \Delta t_{ij} \phi_i^*. \end{aligned} \quad (12)$$

Details of this formula were presented by Wildeman et al. [48]. The cost $\Delta H_{G^k}(t_{G^k})$ of shifting all the activities i^j within a group G^k to a new execution date t_{G^k} can be expressed as:

$$\Delta H_{G^k}(t_{G^k}) = \sum_{i^j \in G^k} h_i(t_{G^k} - t_{ij}) = \sum_{i^j \in G^k} h_i(\Delta t_{ij}), \quad (13)$$

which in turn is strictly convex ($\Delta H_{G^k}''(\cdot) > 0$). The optimization of the ideal execution date of the group $t_{G^k}^*$ is represented graphically in Fig. 3. The economic profit $EP(G^k)$ generated by a group is then calculated as:

$$EP(G^k) = U_{G^k} - \Delta H_{G^k}^*; \quad (14)$$

a negative value of EP means that the grouping of the activities within a group G^k is not convenient, and that it is possible to split the group in two or more subgroups which lead to higher savings. The set of all groups is called *grouping structure*, is identified with SGM , and represents a partition of the set of preventive maintenance activities.

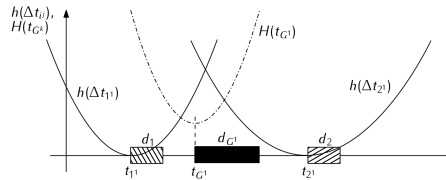


FIGURE 3. The solid curves in the plot represent the value of the penalty function of each activity as function of the deferment Δt_{ij} . The dot-dashed curve is the penalty function of group G^1 , within which we suppose to group activities 1^1 and 1^2 . The minimum of the curve corresponds to the ideal execution date $t_{G^k}^*$ of the group.

The economic profit of a grouping structure $EPS(SGM)$ is defined as:

$$\begin{aligned} EPS(SGM) &= \sum_{G^k \in SGM} EP(G^k) \\ &= \sum_{G^k \in SGM} (U_{G^k} - \Delta H_{G^k}^*). \end{aligned} \quad (15)$$

The goal of the problem is to maximize the profit given by the grouping structure. More formally, we search the optimal grouping structure SGM_{opt} :

$$SGM_{opt} = \arg \max_{SGM} EPS(SGM). \quad (16)$$

E. OPPORTUNISTIC APPROACH

The model for dynamic grouping maintenance presented above can be modified to include special needs of maintenance managers. With special needs we mean the occurrence of a sudden fault, or any planned activity that must be performed at a certain time. An opportunity to perform maintenance at a time t_{opp} on a component i is modeled with the following penalty function:

$$h_i(t) = \begin{cases} 0, & \text{if } t = t_{opp} \\ +\infty, & \text{if } t \neq t_{opp}. \end{cases} \quad (17)$$

Eq. (17) reads that if the maintenance activity on the faulty component i is executed at time t , with $t \neq t_{opp}$, the cost for shifting the activity is extremely high. Therefore, the activity on the faulty component will most likely be executed at t_{opp} and other activities will possibly be anticipated and grouped with it.

III. SIMULATIONS

A. METHODOLOGY

The goal of the simulation approach is to analyze the cost of different maintenance policies in a multi-component system environment with randomly generated faults and variable set-up costs. We have designed a set of discrete event simulation (DES) procedures that mimic different maintenance policies.

Three preventive maintenance policies are tested for 8670 hours, namely one year of simulated time, in a Monte Carlo experiment using different setup costs $S \in \{0, 50, 100, 150, \dots, 600\}$. Each combination of policy and setup cost is tested 1,000 times in order to obtain information about the average cost of maintenance, the average availability of the system, and maintenance frequency. Independently of the policy, the system is shut down every time that a corrective or preventive maintenance policy is performed, and, when maintenance is executed, there is no ageing of components. After a component has been repaired, its degradation state is considered as-good-as-new. The experiment mimics the dynamic environment that we have in the real world by simulating the occurrence of faults according to the time to failure distribution of each component.

Maintenance activities are scheduled and managed according to the following three maintenance policies:

- **Minimal repair policy (MRP):** A preventive maintenance activity is scheduled for all the components at time intervals of x_i^* (> 0) working hours based on the age. According to the rolling horizon approach, the event with the earliest date is processed and it can be either a corrective or a preventive maintenance intervention; after a component has been processed, a new PM is scheduled.
- **Adaptive grouping policy (GPa):** According to this policy, maintenance activities are initially planned at time intervals x_i^* for all the components. The grouping structure is optimized using the GA and the first group of activities, i.e. group G_1 in Fig. 4, is added to the maintenance plan. According to the rolling horizon approach, the system starts to process the first event, which could be either a group of PM activities, or a CM activity. If the upcoming event is a group of activities it is regularly executed, then new PM activities are planned and a new grouping structure is found using the GA. The process loops until a failure event occurs, in which case the system executes maintenance only on the faulty component; indeed, in case a fault occurs, no grouping

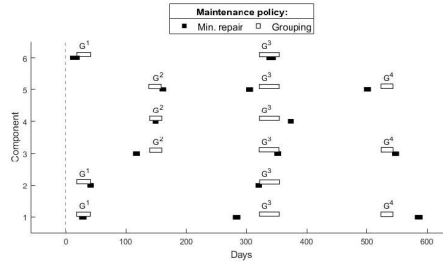


FIGURE 4. The black rectangles represent the activities planned at their optimal date, whereas the white rectangles represent the groups of activities; each group is identified with a label G^k . The example is realized with data from Table 1, and $S = 100$.

is performed, but the faulty component is immediately repaired. A new grouping structure needs to be found taking into account that a new PM activity has been added to the plan at $t_{opp} + x_i^*$, where i is the faulty component.

- **Opportunistic grouping policy (OGP):** According to the rolling horizon approach, a preventive maintenance intervention is planned for all the components, thus making available a temporary schedule. Based on the information contained in the schedule, the grouping structure is optimized using a genetic algorithm (GA) and the system produces a new maintenance schedule made of groups of activities, which resembles that of Fig. 4. If the upcoming event in the simulation is a preventive maintenance intervention, i.e. group G_1 in Fig. 4, this is regularly executed and new PM activities are planned for each component. The grouping structure is then optimized again using the GA and the new maintenance dates of the components just maintained; subsequently, the group of PM activities with the earliest date is executed, unless a fault event occurs. When a fault occurs, the GA is called down to optimize the grouping structure implementing (17), which has the aim to lock down the CM activity at the time of fault t_{opp} . The effect produced on the schedule is represented in Fig. 5, where it is possible to see that the activity on the faulty component

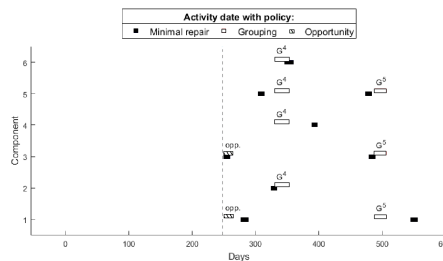


FIGURE 5. OGP: When component 3 fails the maintenance on component 1 is anticipated at t_{opp} .

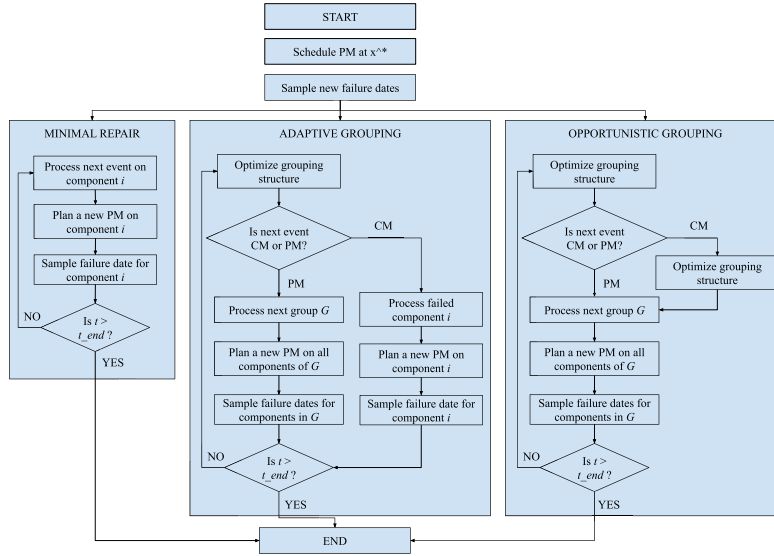


FIGURE 6. Simulation procedures for the different policies analyzed.

(component number 3) is grouped together with component 1 at time t_{opp} , thus saving one setup cost S . The simulation restarts by planning new PM activities for the components that were maintained in the last group and subsequently a new grouping structure is optimized based on this schedule.

In all the policies, the algorithm continues to process events according to the previous rules until the end of the simulation horizon t_{end} is reached.

After a group of activities has been executed, its cost C is calculated according to the following equation for all the policies:

$$C = \sum_{i \in G \setminus f} C_i^P + C_f^C + S \quad (18)$$

The equation is valid for groups with one or more components, among which at most one can be failed. The set of components involved in the group maintenance is indicated by G , C_i^P is the cost of preventive maintenance of a component, and C_f^C is the cost of corrective maintenance on the failed component f (if there is a faulty component). The last term represents the setup cost S which, as stipulated in (11), is paid only once instead of $|G|$ times.

The policies studied in this paper are summarized in the flowchart presentation in Fig. 6, which can also be considered a small, but original, contribution to the field of maintenance-systems simulation.

TABLE 1. Characteristics of the components of the system.

Comp.	λ_i	β_i	C_i^P	C_i^C	d_i
1	450	2.1	300	500	5.61
2	550	1.4	500	700	5.28
3	340	2.2	300	500	6.12
4	440	1.3	300	490	5.85
5	450	2.0	500	1010	4.50
6	340	1.9	300	450	4.38

1) DATA

The data used to simulate the system are similar to those of previous studies [11], [12], [43], [48]. Table 1 lists the data about the six components used to simulate the system.

Durations of maintenance activities, d_i 's in Table 1, are assumed to be deterministic in the experiment. Testing of the model assuming a stochastic duration d_i of maintenance activities was carried out and no meaningful effects on results were observed; therefore, we report the results for the deterministic case in order to avoid overparametrization of the experiment.

2) GROUPING STRUCTURE OPTIMIZATION WITH A GENETIC ALGORITHM

Grouping structure optimization is a complex combinatorial problem. It has been demonstrated that similar optimization problems [43], [44] are \mathcal{NP} -complete. The reason why we decided to use a genetic algorithm (GA) in the optimization

of the grouping structure is the known ability of GAs to find (near-optimal) solutions for combinatorial problems, as confirmed by many authors [11], [17], [31], [43], [44]. In the context of this research, the GA was written for this specific purpose and includes a feasibility check of the individual solutions for the initial population in the hope of speeding up the optimization. Further details on the GA implementation can be found in the appendix.

The simulation procedure was implemented using the object oriented programming approach in Python 3.7. The realization of the discrete event simulation is based on the Python library SimPy distributed freely under MIT license.

The total set of simulations required roughly 30 hours to run on a desktop computer with the following characteristics: 64-bit Windows Server 2016, Intel® Xeon® Platinum 8160 CPU 2.10 GHz, and 768 GB of RAM. The time to run a single optimization of the grouping structure required few seconds using the *stall generations* stopping criteria; that is, the algorithm stopped after the best value had not changed in the last 15 iterations.

B. SIMULATION RESULTS

The results of the simulations offer insight on the costs associated with different maintenance policies. As shown in Fig. 7, the minimal repair policy (MRP) is the least efficient policy with the highest cost in all cases, while the opportunistic grouping policy (OGP) is the lowest cost policy. This result is not completely unexpected since the OGP is the most sophisticated policy. However, given the complexity of the problem, it was not obvious, at least for us, to observe a linear relation between setup costs and total maintenance costs.

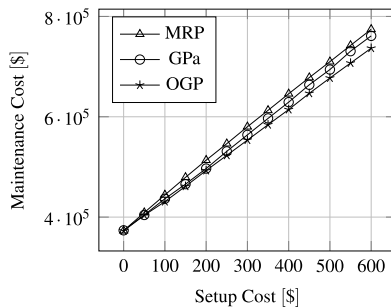


FIGURE 7. The average cost of maintenance with different policies with respect to different set-up costs.

The analysis also returned the number and the type of maintenance activities performed within the simulation horizon. The number of corrective interventions was always higher on average than the number of preventive interventions for all the policies, as it is possible to see in Fig. 8.

colorblue It is possible to notice that the number of corrective interventions increases as a function of the setup costs. In fact, exploiting the grouping strategy can induce the

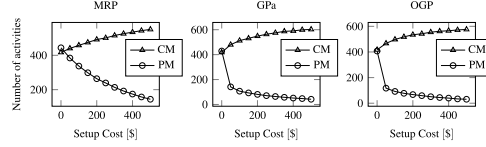


FIGURE 8. The average number of maintenance activities divided by type. The execution of PM on a group of components is counted as a single maintenance activity.

algorithm to increase the risk of component failures. Also, already with a small setup cost, e.g. $S = 50$, there is a substantial reduction of preventive maintenance activities, due to their grouping.

The values in Fig. 8 count the amount of groups, but they provide no information on the duration of interventions and on the number of components maintained within a group. According to the MRP and GPa policies, CM activities are carried out singularly, whereas according to OGP, a CM activity might involve multiple components; this difference significantly affects the availability of the system. The availability of the system is a relevant and practical metric of system effectiveness. Therefore the average availability produced by each policy for each setup cost was measured. In this experiment, the availability of a system at time t (in the past) corresponds to its state and is assumed to be binary, as follow:

$$A(t) = \begin{cases} 1, & \text{if the system was working at time } t \\ 0, & \text{if the system was not working at time } t. \end{cases}$$

In particular, we are interested in measuring the average availability of the system over the simulation horizon. This is defined as follows [37]:

$$\bar{A} = \frac{1}{t_{end} - t_{begin}} \int_{t_{begin}}^{t_{end}} A(t) dt. \quad (19)$$

where t_{begin} and t_{end} are the beginning and the end of the simulation horizon, respectively.

The average cost of maintenance and the average availability of the system were compared using a bi-objective analysis: in Fig. 9, each policy is represented by a point in the cost-reliability space at a given value of S . The coordinates of each point are the average value of maintenance cost and availability produced by the relative policy. Uncertainty about cost and availability are not represented in Fig. 9 since the standard deviations of the underlying distributions are too small to provide clear information.

To achieve the best operating performance, a policy should maximize availability and minimize the expected cost of maintenance.

The results in Fig. 9 show that there is not a dominating policy, and the final choice is a matter of trade-offs. This result highlights the fact that the maintenance policies studied so far aim at minimizing the maintenance cost, but they overlook the availability of the systems.

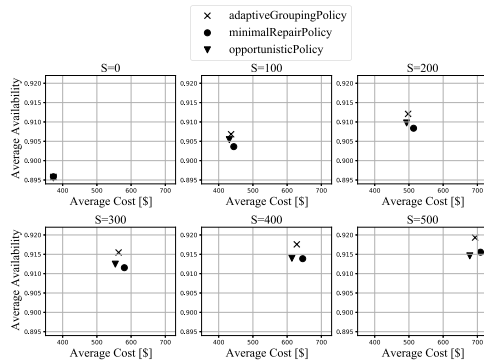


FIGURE 9. Bi-objective (cost vs. availability) comparison of policies for different setup costs ($N = 6$). Costs are expressed in 1,000 units.

More specifically, the greater availability associated to the GPa policy compared to the OGP may be due to the fact that OGP tends to anticipate some maintenance activities on the ground of purely economic reasons, but by doing so it results in a larger number of maintenance activities which ultimately leads to a worse availability of the system. Note that, with large setup costs even the MRP beats the OGP in terms of availability.

We also explored the case with $N = 10$ where four additional components — with characteristics similar to the already existing six — were added in the system. The results are shown in Fig. 10 with three setup costs and strengthen the results obtained with $N = 6$: in this case, the loss of availability associated with the opportunistic policy is even more evident.

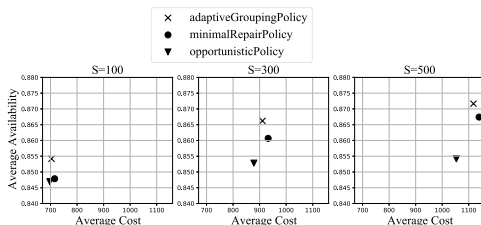


FIGURE 10. Bi-objective (cost vs. reliability) comparisons of policies for different setup costs ($N = 10$). Costs are expressed in 1,000 units.

IV. DISCUSSION AND CONCLUSION

As recalled by George-Williams and Patelli [15]: “identifying the optimal maintenance strategy is a challenge”. In the hope of helping to solve this challenge, we presented a comparative study of selected maintenance policies and a framework for analyzing them. The policies were tested in an operational environment with randomly generated faults by means of discrete event simulations. The results clearly

indicate that there are tangible cost savings that can be reached by using maintenance policies based on taking an opportunistic approach for grouping of maintenance activities. In this sense, this study brings new numerical evidence to support the importance of grouping activities to save on maintenance costs. On the other hand, simulations showed that the majority of maintenance intervention was corrective, as confirmed by Fig. 8. This means that scheduling PM activities according to (10) leads to the execution of a few groups of PM activities along the simulation history; the opportunistic approach is thus particularly useful in practice, when there is uncertainty about which components to maintain in case of a sudden failure. On the other hand, the implementation of a monitoring system and a so-called condition based maintenance (CBM) approach would help optimize the PM schedule by detecting a state of imminent failure of a component. The importance of anticipating a near-to-failure condition is corroborated by the non-negligible amount of CM activities shown in Fig. 8. We conjecture that, in a real-world application, the opportunistic policy would benefit from CBM either by lowering the cost due to unplanned shutdowns, and by relieving their technical consequences. Thus, our results can be also interpreted as additional evidence pushing towards the adoption of condition based maintenance systems.

A further step to approach maintenance to reality is the implementation of imperfect maintenance interventions. The as-good-as-new assumption for repair of components could be relaxed through the addition of a significant number of new parameters: random variables to describe the degradation level of a component [42], new TTF distributions for imperfectly repaired components [25], and additional repairing costs. Moreover, imperfect maintenance models have already been extensively addressed in the literature [34], [46], and, in the context of this research, a maintenance policy implementing imperfect maintenance interventions would be of little help to clarify our contribution on the comparison of maintenance policies.

Nevertheless, by extending the analysis to consider also the average availability of the system we were unable to find a dominating policy (in a Pareto efficiency sense). This corroborates the importance of a careful *a priori* selection of the preferred policy considering the preferences of a decision maker in terms of cost vs. availability trade-off. Hence, in practice, as argued in the introduction, the discrete event simulation methodology employed in this paper can be seen as a valuable support to choose the most suitable maintenance policy to any given context.

Besides its use to obtain the results analyzed in the previous section, it is possible to imagine that the presented methodology has at least two more uses. Firstly, it can be employed, for budgetary purposes, as a *predictive analytics* tool to forecast the expected maintenance costs for a given period of time. Secondly, it can be used in *prescriptive analytics* to optimize maintenance schedules. In fact, despite the long time required by our simulations (about 30 hours), a single optimization of the maintenance schedule for the near

future — i.e. an instance containing the next few PMs on all the components — required few seconds both with $N = 6$ and $N = 10$ components.

Let us remark that, in spite of our simplifying assumption that all groups of activities are feasible and the setup cost is the same for all of them, our framework is flexible and can encompass more specific cases. In fact, different setup costs can be defined for different subset of activity, i.e. instead of a single S we may have $S_A \forall A \subseteq \mathcal{N}$, where \mathcal{N} is defined as the set of all components. This could be useful to model technical dependencies between components. One use could be that of assigning an extremely high setup costs to technically unfeasible groups to make them non-optimal and therefore never appear in the optimal maintenance schedule.

Further work is required to optimize the running time so that more complex model environments become tractable. This may include testing various optimization routines to check which kind of optimization methods perform best in the maintenance policy optimization environment.

Other topics for further research include making modifications to the maximum number of activities grouped together, in order to be compatible with real-world shift duration, and to be in sync with real-world availability of repairmen. A more complex model could be built by including a connection to a spare parts management model or a workforce management model. One important avenue for further research is taking the model to the real world and testing it with real data, further enhancements could then, for example, also include a prognostic learning model for the estimation of the useful life of the components used.

APPENDIX: GENETIC ALGORITHM

The use of an heuristic method becomes necessary due to the computational complexity of the problem, especially for $N = 10$. We chose to represent the grouping structure *SGM* using a vector of integer numbers. The list of activities (i.e. all the activities on all the components) was sorted by ideal execution date t_{ij} and each element of the *SGM* vector encoded the group to which one activity belongs. In the case of $n = 6$, each group contains at most 6 activities, therefore a feasible *SGM* consisted in a partition of the set of activities, where each partition contains at most six activities. Each partition was identified with an integer number, thus the resulting vector looked like the following:

$$SGM = (1, 1, 1, 2, 2, 2, 2, \dots, 5, 5, 6, 6, 6).$$

The constraint on partitioning can be exploited to generate new feasible *SGMs*.

The selection of parents for the next generation was performed according to the *wheel of fortune method*, i.e. the *SGMs* showing the highest scores were more likely to be selected as parents for the next generation.

Both mutation and crossover operations were carried out with respect to the structure of the solution. The mutation operation required to choose a mutation point, which could be each of the elements of the *SGM* vector. A single mutation

occurred with a probability of 1 for all the selected individuals. The crossover operation was performed at a single point of the *SGM* vector on a selected pool of individuals in order to produce the desired number of modified individuals.

An elitist strategy was adopted, therefore the individuals with the best fitness score were copied to the next generation.

Finally, the adopted stopping criterion was *generation limit*. That is, the algorithm stopped if the average relative change in the best fitness value did not change for more than 15 generations.

REFERENCES

- [1] A. A. Alabdulkarim, P. D. Ball, and A. Tiwari, "Applications of simulation in maintenance research," *World J. Model. Simul.*, vol. 9, no. 1, pp. 14–37, 2013.
- [2] A. Alrabghi and A. Tiwari, "A review of simulation-based optimisation in maintenance operations," in *Proc. 15th Int. Conf. Comput. Modelling Simulation (UKSim)*, Apr. 2013, pp. 353–358.
- [3] A. Alrabghi and A. Tiwari, "State of the art in simulation-based optimisation for maintenance systems," *Comput. Ind. Eng.*, vol. 82, pp. 167–182, Apr. 2015.
- [4] A. Alrabghi and A. Tiwari, "A novel approach for modelling complex maintenance systems using discrete event simulation," *Rel. Eng. Syst. Saf.*, vol. 154, pp. 160–170, Oct. 2016.
- [5] M. Bevilacqua and M. Braglia, "The analytic hierarchy process applied to maintenance strategy selection," *Rel. Eng. Syst. Saf.*, vol. 70, no. 1, pp. 71–83, Oct. 2000.
- [6] K. Bouvard, S. Artus, C. Bérenguer, and V. Cocquemot, "Condition-based dynamic maintenance operations planning & grouping. Application to commercial heavy vehicles," *Rel. Eng. Syst. Saf.*, vol. 96, no. 6, pp. 601–610, Jun. 2011.
- [7] Z. Cheng, Z. Yang, and B. Guo, "Optimal opportunistic maintenance model of multi-unit systems," *J. Syst. Eng. Electron.*, vol. 24, no. 5, pp. 811–817, Oct. 2013.
- [8] D. I. Cho and M. Parlar, "A survey of maintenance models for multi-unit systems," *Eur. J. Oper. Res.*, vol. 51, no. 1, pp. 1–23, Mar. 1991.
- [9] R. Dekker, R. E. Wildeman, and F. A. van der Duyn Schouten, "A review of multi-component maintenance models with economic dependence," *Math. Methods Oper. Res.*, vol. 45, no. 3, pp. 411–435, Oct. 1997.
- [10] P. Do, P. Scarf, and B. Iung, "Condition-based maintenance for a two-component system with dependencies," *IFAC-PapersOnLine*, vol. 48, no. 21, pp. 946–951, 2015.
- [11] P. Do, H. C. Vu, A. Barros, and C. Bérenguer, "Maintenance grouping for multi-component systems with availability constraints and limited maintenance teams," *Rel. Eng. Syst. Saf.*, vol. 142, pp. 56–67, Oct. 2015.
- [12] P. Do Van, A. Barros, C. Bérenguer, K. Bouvard, and F. Brissaud, "Dynamic grouping maintenance with time limited opportunities," *Rel. Eng. Syst. Saf.*, vol. 120, pp. 51–59, Dec. 2013.
- [13] M. El Hayek, E. van Voorthuysen, and D. W. Kelly, "Optimizing life cycle cost of complex machinery with rotatable modules using simulation," *J. Qual. Maintenance Eng.*, vol. 11, no. 4, pp. 333–347, Dec. 2005.
- [14] J. Geng, M. Azarian, and M. Pecht, "Opportunistic maintenance for multi-component systems considering structural dependence and economic dependence," *J. Syst. Eng. Electron.*, vol. 26, no. 3, pp. 493–501, Jun. 2015.
- [15] H. George-Williams and E. Patelli, "Maintenance strategy optimization for complex power systems susceptible to maintenance delays and operational dynamics," *IEEE Trans. Rel.*, vol. 66, no. 4, pp. 1309–1330, Dec. 2017.
- [16] G. Guizzi, M. Gallo, and P. Zoppoli, "Condition based maintenance: Simulation and optimization," in *Proc. 8th WSEAS Int. Conf. Syst. Sci. Simulation Eng. (ICOSSSE)*, 2009, pp. 319–325.
- [17] Y. Hani, L. Amodeo, F. Yalaoui, and H. Chen, "Simulation based optimization of a train maintenance facility," *J. Intell. Manuf.*, vol. 19, no. 3, pp. 293–300, Jun. 2008.
- [18] P. Hilber, V. Miranda, M. A. Matos, and L. Bertling, "Multiobjective optimization applied to maintenance policy for electrical networks," *IEEE Trans. Power Syst.*, vol. 22, no. 4, pp. 1675–1682, Nov. 2007.
- [19] J. Huang, Q. Chang, J. Zou, and J. Arinez, "A real-time maintenance policy for multi-stage manufacturing systems considering imperfect maintenance effects," *IEEE Access*, vol. 6, pp. 62174–62183, 2018.

- [20] Y. Jiang, M. Chen, and D. Zhou, "Joint optimization of preventive maintenance and inventory policies for multi-unit systems subject to deteriorating spare part inventory," *J. Manuf. Syst.*, vol. 35, pp. 191–205, Apr. 2015.
- [21] D. Kececioglu and F.-B. Sun, "A general discrete-time dynamic programming model for the opportunistic replacement policy and its application to ball-bearing systems," *Rel. Eng. Syst. Saf.*, vol. 47, no. 3, pp. 175–185, Jan. 1995.
- [22] N. H. Kim, D. An, and J. H. Choi, *Prognostics and Health Management of Engineering Systems: An Introduction*. Cham, Switzerland: Springer, 2017.
- [23] K. A. H. Kobbacy, "Application of artificial intelligence in maintenance modelling and management," *IFAC Proc. Volumes*, vol. 45, no. 31, pp. 54–59, 2012.
- [24] J. Koochaki, J. A. C. Bokhorst, H. Wortmann, and W. Klingenberg, "Condition based maintenance in the context of opportunistic maintenance," *Int. J. Prod. Res.*, vol. 50, no. 23, pp. 6918–6929, Dec. 2012.
- [25] Y. Liu and H.-Z. Huang, "Optimal selective maintenance strategy for multi-state systems under imperfect maintenance," *IEEE Trans. Rel.*, vol. 59, no. 2, pp. 356–367, Jun. 2010.
- [26] C. E. Love and R. Guo, "Utilizing weibull failure rates in repair limit analysis for equipment replacement/preventive maintenance decisions," *J. Oper. Res. Soc.*, vol. 47, no. 11, pp. 1366–1376, Nov. 1996.
- [27] R. K. Mobley, *An Introduction to Predictive Maintenance*. Amsterdam, The Netherlands: Elsevier, 2002.
- [28] T. Murino, E. Romano, and P. Zoppoli, "Maintenance policies and buffer sizing: An optimization model," *WSEAS Trans. Bus. Econ.*, vol. 6, no. 1, pp. 21–30, 2009.
- [29] K.-A. Nguyen, P. Do, and A. Grall, "Multi-level predictive maintenance for multi-component systems," *Rel. Eng. Syst. Saf.*, vol. 144, pp. 83–94, Dec. 2015.
- [30] R. P. Nicolai and R. Dekker, "Optimal maintenance of multi-component systems: A review," in *Complex System Maintenance Handbook*. London, U.K.: Springer, pp. 263–286.
- [31] A. Oyarbide-Zubillaga, A. Goti, and A. Sanchez, "Preventive maintenance optimisation of multi-equipment manufacturing systems by combining discrete event simulation and multi-objective evolutionary algorithms," *Prod. Planning Control*, vol. 19, no. 4, pp. 342–355, Jun. 2008.
- [32] F. Pargar, O. Kauppila, and J. Kujala, "Integrated scheduling of preventive maintenance and renewal projects for multi-unit systems with grouping and balancing," *Comput. Ind. Eng.*, vol. 110, pp. 43–58, Aug. 2017.
- [33] N. M. Paz, W. Leigh, and R. V. Rogers, "The development of knowledge for maintenance management using simulation," *IEEE Trans. Syst., Man, Cybern.*, vol. 24, no. 4, pp. 574–593, Apr. 1994.
- [34] H. Pham and H. Wang, "Imperfect maintenance," *Eur. J. Oper. Res.*, vol. 94, no. 3, pp. 425–438, 1996.
- [35] Y. Peng, M. Dong, and M. J. Zuo, "Current status of machine prognostics in condition-based maintenance: A review," *Int. J. Adv. Manuf. Technol.*, vol. 50, nos. 1–4, pp. 297–313, Sep. 2010.
- [36] A. N. Rao and B. Bhadury, "Opportunistic maintenance of multi-equipment system: A case study," *Qual. Rel. Eng. Int.*, vol. 16, no. 6, pp. 487–500, Nov. 2000.
- [37] M. Rausand and A. Høyland, *System Reliability Theory: Models, Statistical Methods, and Applications*, vol. 396. Hoboken, NJ, USA: Wiley, 2003.
- [38] M. Savsar, "Performance analysis of an FMS operating under different failure rates and maintenance policies," *Int. J. Flexible Manuf. Syst.*, vol. 16, no. 3, pp. 229–249, Jul. 2004.
- [39] M. Sheikhalishahi, L. Pintelon, and A. Azadeh, "An integrated approach for maintenance planning by considering human factors: Application to a petrochemical plant," *Process Saf. Environ. Protection*, vol. 109, pp. 400–409, Jul. 2017.
- [40] H. Shi and J. Zeng, "Real-time prediction of remaining useful life and preventive opportunistic maintenance strategy for multi-component systems considering stochastic dependence," *Comput. Ind. Eng.*, vol. 93, pp. 192–204, Mar. 2016.
- [41] F. A. Van der Duyn Schouten and S. G. Vanneste, "Maintenance optimization of a production system with buffer capacity," *Eur. J. Oper. Res.*, vol. 82, no. 2, pp. 323–338, Apr. 1995.
- [42] A. Van Horenbeek and L. Pintelon, "A dynamic predictive maintenance policy for complex multi-component systems," *Rel. Eng. Syst. Saf.*, vol. 120, pp. 39–50, Dec. 2013.
- [43] H. C. Vu, P. Do, A. Barros, and C. Bérenguer, "Maintenance grouping strategy for multi-component systems with dynamic contexts," *Rel. Eng. Syst. Saf.*, vol. 132, pp. 233–249, Dec. 2014.
- [44] H. C. Vu, P. Do, A. Barros, and C. Bérenguer, "Maintenance planning and dynamic grouping for multi-component systems with positive and negative economic dependencies," *IMA J. Manage. Math.*, vol. 26, no. 2, pp. 145–170, Apr. 2015.
- [45] H. Wang, "A survey of maintenance policies of deteriorating systems," *Eur. J. Oper. Res.*, vol. 139, no. 3, pp. 469–489, Jun. 2002.
- [46] H. Wang and H. Pham, *Reliability and Optimal Maintenance*. Cham, Switzerland: Springer, 2006.
- [47] L. Wang, J. Chu, and W. Mao, "An optimum condition-based replacement and spare provisioning policy based on Markov chains," *J. Qual. Maintenance Eng.*, vol. 14, no. 4, pp. 387–401, Sep. 2008.
- [48] R. E. Wildeman, R. Dekker, and A. C. J. M. Smit, "A dynamic policy for grouping maintenance activities," *Eur. J. Oper. Res.*, vol. 99, no. 3, pp. 530–551, Jun. 1997.
- [49] X. Zhao, Z. Lv, Z. He, and W. Wang, "Reliability and opportunistic maintenance for a series system with multi-stage accelerated damage in shock environments," *Comput. Ind. Eng.*, vol. 137, Nov. 2019, Art. no. 106029.
- [50] X. Zhou, L. Xi, and J. Lee, "Opportunistic preventive maintenance scheduling for a multi-unit series system based on dynamic programming," *Int. J. Prod. Econ.*, vol. 118, no. 2, pp. 361–366, Apr. 2009.



MICHELE URBANI received the M.Sc.Eng. degree from the University of Trento, Italy. He is currently a Junior Researcher with the Department of Industrial Engineering, University of Trento. He is also working on the dissertation in maintenance modeling with a focus on how analytics and optimization may affect the field maintenance.



MATTEO BRUNELLI received the Ph.D. degree from Åbo Akademi University, Finland. For a period of five years, he was a Postdoctoral Researcher with the Systems Analysis Laboratory, Aalto University, Finland. He is currently an Associate Professor of mathematical methods for economics with the Department of Industrial Engineering, University of Trento, Italy. His research interests include decision analysis, applied optimization, and mathematical representation of uncertainty.



MIKAEL COLLAN (Associate Member, IEEE) received the D.Sc. degree in information systems from Åbo Akademi University, Finland. He was on the boards of numerous SME companies. He is currently a tenured Professor of strategic finance with the Lappeenranta-Lahti University of Technology, Lappeenranta, Finland. His research interests include business model analysis, investment decision-making, application of analytics in business, industry, healthcare, and application of fuzzy logic in business modeling. He was the Past President of the Finnish Operations Research Society. He is an Ordinary Member of the Finnish Society of Sciences and Letters and the four Finnish academies of sciences.

...

Publication III

Urbani, M., Brunelli, M., and Punkka, A.

An approach for bi-objective maintenance scheduling on a networked system with limited resources

Reprinted with permission from

©n.d., Authors

Publication IV

Savolainen, J., and Urbani, M.

Maintenance optimization for a multi-unit system with digital twin simulation

Reprinted with permission from

Journal of Intelligent Manufacturing

DOI: 10.1007/s10845-021-01740-z

©2021, Springer



Maintenance optimization for a multi-unit system with digital twin simulation

Example from the mining industry

Jyrki Savolainen¹ · Michele Urbani²

Received: 15 June 2020 / Accepted: 13 January 2021
© The Author(s) 2021

Abstract

Optimization of operations and maintenance (O&M) in the industry is a topic that has been largely studied in the literature. Many authors focused on reliability-based approaches to optimize O&M, but little attention has been given to study the influence of macroeconomic variables on the long-term maintenance policy. This work aims to optimize time-based maintenance (TBM) policy in the mining industry. The mine environment is reproduced employing a virtual model that resembles a digital twin (DT) of the system. The effect of maintenance decisions is replicated by a discrete event simulation (DES), whereas a model of the financial operability of the mine is realized through System Dynamics (SD). The simultaneous use of DES and the SD allows us to reproduce the environment with high-fidelity and to minimize the cost of O&M. The selected illustrative case example demonstrates that the proposed approach is feasible. The issues of using high dimensional simulation data from DT-models in managerial decision making is identified and discussed.

Keywords Maintenance optimization · Digital twin · Simulation · Optimization

Introduction

Managing large industrial plants in global competition requires a clear strategic view and a high level of control of operations. Anytime an industry relies on its physical assets, the success of operations is tightly linked to the execution of the right level of maintenance. Maintenance has both the role of keeping an asset in its best condition and to minimize unforeseen system downtime. From a managerial viewpoint, operations and maintenance (O&M) cannot be thought in isolation from the economic context within which every industry operates: a coordinated view of O&M should aim at reaching the right amount of responsiveness and throughput of a sys-

tem that is required by the prevailing market conditions. In this research, the issue is investigated using an example from the metal mining industry, where efficient real-time management of operations is essential to meet the production targets, but where ultimately macro-economic variables, mainly the price of the metal, play the key role in bottom-line profitability in the long-term.

As stated by Bevilacqua and Braglia (2000) and Mobley (2002), maintenance costs can rise to 60 % of total production costs. This cost item can be affected in the short- and medium-term by planning and optimization - unlike many other major costs of industrial operations that are fixed in nature. Despite the importance to plan operations for the impact on long-term profitability, there is only a limited amount of literature on the topic, except for Topal and Ramazan (2010) who introduced a model to estimate maintenance costs in a 10-years mine lifetime. Furthermore, considering multi-machine environments, the sheer size, and the resulting complexity due to a high number of uncertainties is a major hurdle for model development (West and Blackburn 2017). Addressing a company-wide problem-setting, like managing real-time operations and maximizing long-term profitability

✉ Jyrki Savolainen
jyrki.savolainen@lut.fi

Michele Urbani
michele.urbani@unitn.it

¹ Lappeenranta University of Technology, School of Business and Management, Yliopistonkatu 34, 53850 Lappeenranta, Finland

² Department of Industrial Engineering, University of Trento, via Sommarive 9, 38123 Povo, TN, Italy

in a dynamic economic context, requires the help of both advanced analysis methods and control tools. We address the topic using a Digital Twin (DT) modeling concept that is used here in a meaning discussed by, e.g., Rosen et al. (2015), Grieves and Vickers (2017), to refer to “interconnected and multidisciplinary simulation models usable for operations optimization on a system level”. In a recent review of Kendrik et al. (2020) five use-categories of DTs were identified of which the manufacturing stage and usage stage of a system is addressed in this paper.

A DT is a digital model of a physical entity (Negri et al. 2017; Tao et al. 2018; Redelinghuys et al. 2020) providing human-readable, semantic, data-model of reality (Negri et al. 2017; Kunath and Winkler 2018). These models reside in a high-performance, usually cloud-based, computing environment and they can be used for several types of optimization purposes (Negri et al. 2017; Tao et al. 2018; Cimino et al. 2019; Madni et al. 2019). Kendrik et al. (2020) highlight the importance of the digital counterpart of a system to optimize production performance and maximize profitability, which is the goal of the proposed model. The origins of DT can be traced back to the beginning of the 2010s in the aviation and aerospace industry. The early publications (Tuegel et al. 2011; Shafto et al. 2010) revolved around the possibilities of using ultra-high fidelity models to simulate aircrafts’ maintenance under dynamic operating conditions over the equipment lifetime. In this vein, Kritzinger et al. (2018) highlighted the communication aspect between physical and virtual spaces claiming that only models transmitting data in and out from the virtual space can be regarded as DTs. The latter should not be confused with general digital models (no data connection) or digital shadows (only physical-cyber connection). The required fidelity level in DT models remains debated and, in this research, we agree with the claim by Wright and Davidson (2020) that “digital twins can use any sort of model that is a sufficiently accurate representation of the physical object being twinned”.

This paper focuses on the question of building and utilizing multi-domain simulation models that would integrate O&M simulation optimization with the overall profitability simulation of industrial operations in a way that could be referred to as Digital Twin. For the sake of brevity, we limit our scientific inquiry to the context of the mining industry. To answer the research question, a two-phase methodological approach is adopted. First, the general properties of a co-simulation framework are investigated, and references to the relevant literature are provided. Second, an experimental DT model is developed on a virtual case study: a metal mine is considered due to its specific nature of O&M, and because its profitability is directly linked to the price of metal(s), which is a macroeconomic variable. To answer the question, this study presents a conceptual “digital twin” for metal mining that connects a detailed, minute-per-minute maintenance

model of mobile equipment to a monthly-level profitability analysis of metal mining operations. In the model, two separate simulation modules are included: an O&M model, which replicates with high fidelity the effects of O&M decisions, and a managerial cash flow (CF) model, which is used to support decision-making at the production system level. In a co-simulation context, both models are treated as separate simulation units (SU) and when these SUs are considered as a whole, a dynamic system is created (Gomes et al. 2018).

This allows us to replicate a DT model’s operational workflow and software pipelining in a controlled environment, where the O&M model optimizes some of the key system parameters before running the CF simulation for the high-level mid-term economical aspects of the system. The complexity of the system under study is a major reason to adopt a DT-inspired view: where it is not possible to express relationships analytically, a DT can help to integrate data from the field with flexible simulation tools, to achieve an overall improvement of the system’s profitability. Therefore, the goal of this work is to:

- Demonstrate that the DT approach in the context of metal mining operations provides a holistic method to optimize its overall operational profitability under economic uncertainty of metal prices and maintenance costs.
- Point out and discuss the limitations of simulation based digital twins, when it comes to managerial decision making based on multidimensional information.

This paper continues with a brief introduction to the concepts of O&M planning in multi-equipment systems and some general considerations about system dynamics methodology in Sect. 2. In Sect. 3, a literature study on the topic of O&M simulation and DT modeling is provided to set the ground for model building. In Sect. 4, a detailed description of the models—namely the O&M module and the CF model—is provided. This is followed by the empirical application of the model, the validation of the proposed model through two experiments, and a detailed analysis of the results in Sect. 5. The paper closes with conclusions and discussion in Sect. 6, where some strategic considerations are derived from the results of numerical experiments.

Theoretical Background

Machine specific maintenance histories can be tracked with high accuracy using data series of sensory information together with maintenance reports from existing databases. A window of opportunity exists to use this accumulated maintenance information in connection with the DT model depicting the behavior of the overall system. From the reliability-theory point of view, a large-scale industrial system can be mod-

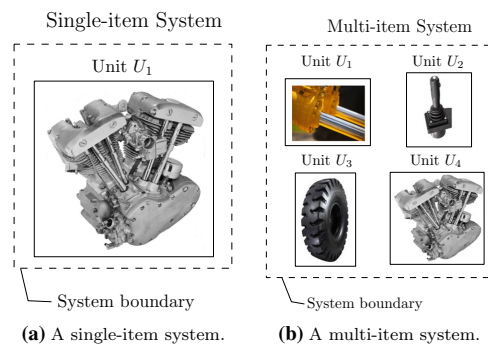


Fig. 1 A schematic representation of single- and multi-item systems

eled using a multi-item system made of several non-identical components, which are characterized by a common set of features, but with proprietary parameters for each feature. Two common examples of such features are the *service time*, and the *time to failure* (TTF)-distributions. In a single-item system, which can be depicted as in Fig. 1a, maintenance can be optimized knowing the TTF distribution and the cost of corrective maintenance.

On the other hand, multi-item systems are sets of components considered as a whole, and they can be represented as in Fig. 1b. One peculiarity of multi-item systems is that very often there is a convenience to carry out maintenance simultaneously on groups of components: since component dependencies of different natures exist – i.e. economic, stochastic, or structural (de Jonge and Scarf 2020) – they can be exploited to minimize maintenance costs and system downtime.

In multi-item systems, maintenance activities and regular operations can be organized according to a maintenance strategy, which determines the rules for scheduling of both. According to Alrabghi and Tiwari (2016), there are two broad classes of strategies: *time-based* maintenance (TBM) and *condition-based* maintenance (CBM) strategies. Both types of strategies include the possibility to perform *corrective maintenance* (CM) and *preventive maintenance* (PM) interventions, where the latter kind of activities are justified by the lower cost of stopping the system and inspect/maintain components before they fail. From the economic point of view, skipping PM can save money in the short-term, but exposes to the risk of more expensive breakdowns in the mid- and long-term.

The major difference between TBM and CBM is the principle that rules decisions: to plan maintenance activities, TBM uses only the work time, whereas CBM exploits also information on the degradation of a component. Depending on the cost and the risk generated by the fault of an item,

both strategies are valuable. Concerning multi-item systems, the state of the art for both types of strategies were reviewed several times in the past (Cho and Parlar 1991; Dekker et al. 1997; Wang 2002; Nicolai and Dekker 2008; de Jonge and Scarf 2020).

A recurrent critique of many multi-item models, which is partly addressed in this paper, is the lack of integration with other fundamental parts of an industrial system—e.g., spare parts and inventory management, human resources management, or planning of operations. Alrabghi and Tiwari (2015, 2016) confirm this by stating that the isolation of maintenance management systems is a limit to their use in practice. The experiment design used in this research resembles the one proposed by Alrabghi and Tiwari (2016) for TBM but contextualized and integrated with higher-level decision-aid tools. The DT framework offers the right testbed for simulation-based production optimization (Uhlemann et al. 2017), and for studying the integration of systems, hence to overcome system isolation.

To deal with the model integration issue, the System Dynamics methodology, originally coined by Forrester (1961), is used in this study. SD is suitable for representing the behavior of complex systems with delays and feedback loops that are constructed using intuitive graphical flowsheet diagrams (Forrester 1994). Within engineering sciences, SD has been traditionally viewed as a high-level managerial method, which is subordinated to fast-to-run, discipline-specific computational models; however, SD has also been applied in several operations research (OR) applications, which were reviewed by Größler et al. (2008). In this paper, the role of the system dynamic model is to serve as a semantic data interface to the overall production system, where all the relevant sub-model(s) can connect.

In this paper we focus our scientific inquiry to the context of metals mining, where the role of equipment reliability is highlighted by the complexity of advanced machinery, and the pressure to meet the production targets (see discussion, e.g., Dhillon (2008)). In real mining systems, data-driven analysis of maintenance policy optimization faces the problem of the reliability behavior of equipment. As a key challenge to maintenance, Hall and Daneshmend (2003) point out that the number of (semi-)mobile equipment hinders the collection of “clean” datasets. Data collection may also be inhibited by the failure of electronic-based hardware (e.g. sensors, wiring, connectors, etc.), which is common in harsh mining environments (Dhillon 2008). The estimation of the near-future degradation state of machines and the forecasting of their most likely end of life require the use of simulation, which is recognized as a main aspect of a DT (Negri et al. 2017; Kritzing et al. 2018; Tao et al. 2018; Cimino et al. 2019). For these reasons, we consider our model eligible to operate as a DT, although the experiments that are presented in the following do not rely on a real-world physical sys-

tem, and a proper product data management system is not implemented.

Literature Study

To clarify the connection of this work with the existing literature, a brief study on the topic of simulation-based DTs and maintenance was conducted. An overview of the models involved in manufacturing system design and operation using DES is provided by Negahban and Smith (2014), who observed that there is an on-going shift to maintenance issues and real-time control. In this vein, we used the following three combinations of keywords to conduct an inquiry on the search engine Scopus: i) “digital twin”, “simulation”, and “maintenance”; ii) “digital twin”, “co-simulation”, “maintenance”; and iii) “co-simulation” and “maintenance”. Based on their relevance to this research, 30 documents were selected and listed in Table 1.

The columns “Digital Twin”, “Maintenance”, and “Co-simulation” are flagged if the keyword represents a relevant topic in the document. The columns “Review”, “Methodology”, and “Application Case” indicate if the document includes a review of the literature, a contribution to methodological aspects, or the presentation of a use case.

Based on the literature study, the number of publications concerning DTs and simulations for O&M optimization increased during the last ten years, as depicted in Figure 2. The majority of the published documents are represented by conference proceedings although the relative share of journal articles has been in a steady increase during the period of 2017–2020. This suggests that the relevance of the topic is being identified in the scientific community.

The content analysis reveals that most of the works (in Table 1) aim at developing technical models of mechanical, electrical, aerospace, and transportation systems, but only a few documents specifically addressed the combination of technical and economic aspects. There seems to be a common understanding that maintenance optimization has a central role in DT models, together with the general aim to improve operations and managerial prediction capabilities. The latter topic heavily relies on the simultaneous use of several simulation tools, but there seems to be little awareness of the co-simulation context that emerges. To verify this observation, our initial research query “i)” was tested by substituting the keyword “simulation” with “co-simulation”: the low number of documents found suggested a lack of general frameworks when co-simulation models are part of a DT.

Outside the context of DTs, the principles and properties of a co-simulation model have been systematically surveyed by Gomes et al. (2018), who highlight the ability to apply separate, “black-box”, simulation units as building blocks of a large (co-)simulation. This aspect is of particular impor-

tance in the real world, where simulation tools for prognostic and health management (Peng et al. 2010; Kim et al. 2016) might come from different developers and they need to be integrated. Several documented industrial applications of co-simulation models within the period 2011–2016 are reported in Gomes et al. (2017).

The documents resulting from query “iii)” are similar to the references mentioned by Gomes et al. (2017) in their literature review. A closer look at these documents reveal that co-simulation models are often “stand-alone” works that do not present a connection with a physical model. In other words, although the potential of co-simulation models in maintenance optimization is clear, there is a lack of research efforts describing how these technical-economic models would be structured and how they would play out. This research work aims at contributing to close this gap by considering simulation optimization of O&M as part of a DT, and by addressing the issue according to the principles of co-simulations.

Data and Methodology

This research addresses the problem of designing a DT, which comes down to the ability to be able to simulate and optimize several models (co-simulation). Such models are not directly integrable due to their fundamental basis (such as software, modeling choice, and the level of detail), and they need to share information in an uncertain/probabilistic environment. Two models are co-simulated in this research: i) an O&M model, and ii) a managerial CF model, which operates high-level decisions based on generated CF and exogenous economic variables. The fleet capacity optimization is used as a means to achieve the economic goals: if there were to be two alternative fleets that meet a production target, the one producing the higher CF would be selected. A schematic diagram of the problem setting is illustrated in Figure 3, where the connections between separate steps are shown.

Inputs of the Digital Twin model consist of the system design and maintenance policy selection, which are marked with (i) and (ii) in Figure 3. These are used to feed the software module (iii) that replicates the operations of a metal mine modeled with a DES. The performance of O&M is evaluated by running a Monte Carlo simulation (MCS) according to the given design and policy selection. The maintenance model can be run with an optimization procedure automatically changing the system design – i.e. the number of resources in the mine, to reach a target value of ore excavated at the minimum cost. The aggregated information produced in (iii) is fed to the managerial feasibility model (v) with economic uncertainties included (iv). The aggregated system output is formed ((vi) in Figure 3), which can be, in case of operating the DT-model continuously, further looped back

Table 1 A summary of the literature study

Reference	Content	Digital Twin	Maintenance	Co-simulation	Review	Methodology	Application Case
Books and Book Chapters							
Boschert and Rosen (2016)	Simulation aspects of DTs are presented.	x		x			
Spiriyagin et al. (2014)	Co-simulation for locomotive design, multibody dynamics, and longitudinal train dynamics is discussed.		x	x		x	
Journal Articles							
Angeljlu et al. (2020)	A DT model for replication of masonry buildings.	x	x			x	x
Lu et al. (2020)	DT for management of buildings.	x	x				x
Du et al. (2020)	A cognitive application of a DT to aid maintenance.	x	x			x	x
Fotland et al. (2020)	DT of offshore vessel cranes.	x	x				x
Li et al. (2020)	Multi-criteria decision-making methods for green performance evaluation methodology of intelligent manufacturing driven by DT.	x	x			x	x
Dowdeswell et al. (2020)	A review of fault detection and diagnostics techniques.		x	x	x		
Szpytko and Duarte (2020)	A complex simulation optimization procedure is developed to improve O&M.	x	x				x
Negri et al. (2020)	A DT-based simheuristic framework for the flowshop scheduling under (failure) uncertainty is proposed.	x	x			x	x
He et al. (2019)	DT for monitoring, diagnosis, and optimized control.	x	x				x
Cimino et al. (2019)	The DT concept is connected to Manufacturing Execution Systems and missing implementation aspects are addressed.	x			x	x	x
Atvaliotis et al. (2019)	A physics-based model is implemented into a DT for RUL prediction.	x	x				x
Ashtari Talkhestani et al. (2019)	The simulation ability of DT is discussed, and the concept of intelligent-DT is presented.	x	x				x

Table 1 (continued)

Reference	Content	Digital Twin	Maintenance	Co-simulation	Review	Methodology	Application Case
Liu et al. (2019)	DT technology is exploited to formulate a real-time maintenance strategy.	x	x		x	x	x
Castañó et al. (2019)	A co-simulation framework is designed and deployed to test sensor reliability in the context of CPSs.		x	x			x
Pradhan et al. (2019)	Co-simulation is used for prediction of wear evolution and vehicle dynamics for maintenance decision-making.			x			x
Mermertas et al. (2019)	Co-simulation for optimization of a large bearing in a steam turbine.		x	x			x
Kaewunruen and Xu (2018)	The DT technology is used to create a replica of historical buildings.	x	x				x
Kontes et al. (2018)	A methodology for improvement of a two-simulation based control is proposed.		x	x			x
Wei et al. (2018)	Co-simulation method for fuel regulator fault diagnosis in a turbopfan engine.		x	x			x
Pradhan et al. (2018)	Rail wheel wear estimation is obtained through co-simulation.		x	x			x
Pradhan et al. (2017)	A co-simulation model for evaluation of rail wheel wear distribution is proposed.		x	x			x
Pombo and Ambrósio (2012)	A high-speed co-simulation procedure for communication between two simulation systems.			x		x	x
Conference Proceedings							
Werner et al. (2019)	A DT model for predictive maintenance applications.	x	x				
Cattaneo and Macchi (2019)	A DT model of a drilling machine obtained by retrofit.	x	x				x
Shubenkova et al. (2018)	DT for failure prediction in industrial vehicles.	x	x				x
Van Der Auwerter et al. (2018)	A review of the use of DTs for construction, validation, and evaluation of the design.	x		x	x		
Boafo et al. (2017)	Co-simulation for real-time safety verification in nuclear power plants.	x	x				
Tuegel et al. (2011)	A new structural modeling concept for the design and maintenance of airframes is proposed.	x	x			x	

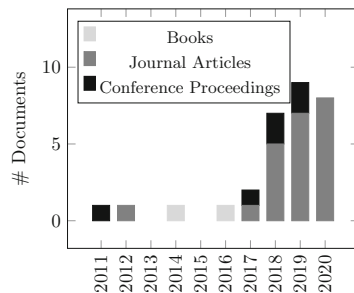


Fig. 2 Number of documents in Table 1 divided per year. “2020” represents the year-to-date in August

to the maintenance module to revise the policy until convergence of results is reached. Once an optimized maintenance policy is found, it can be used to control the physical system.

We acknowledge that the multi-disciplinary model applied in this paper is limited in nature, and we suggest that this model could be referred to as “low fidelity digital twin” (see discussion, e.g., Tuegel et al. (2011)) to distinguish them from the envisioned “full-scale” DT implementations including a wider range of simultaneously operating, high-fidelity, disciplinary sub-models.

Maintenance Module

The justification to use simulation-optimization to model the mine environment in this paper emerges from two reasons: the lack of analytical expressions to model operations, and the need to adapt the configuration of resources to meet the production targets. The module aims at optimizing the maintenance policy, which is a set of heuristic rules to make maintenance decisions. The inherent complexity of the sys-

tem makes it impossible to determine in advance the effects of the proposed maintenance policy, therefore, O&M of a mine’s load and haul process is replicated using DESs in a Monte Carlo simulation experiment. Notwithstanding the possibility to model the environment down to tiny details, the degree of approximation was arbitrarily chosen to provide a realistic amount of complexity in a reasonable amount of time. For a further discussion on the simulation detail-level, the interested reader should refer to, e.g., Zio (2009).

System components are distinguished by type, which defines the available actions when they interact with each other. The elements used to simulate the operations of the mine present a unique behavior, and they are divided in two macro-categories: the first is *server-queue* components, which include shovels, dumpsites or discharge points, and workshops; in this research, server-queue components are represented as in (Law et al. (2000), pp. 12–18). The second category is represented by *agents*, i.e. trucks for transportation of the excavated material around the mine. According to Law et al. (2000) an “agent is an autonomous “entity” that can sense its environment, including other agents, and use this information in making decisions. Agents have attributes and a set of basic if/then rules that determine their behaviors.” The agents can travel between each couple of sites in the mine, and the traveling distance between sites is described by log-normal distributions. This choice allows us to sample the travel time from one site to the other in a realistic way.

The behavior of an agent, i.e. a truck, is characterized by the parameters of the processing time distributions described in the following. A truck is *unreliable* in the sense that it might fail at any moment during operation, and the time to failure (TTF) is a random variable modeled using a two-parameter Weibull distribution

$$W(t; \alpha, \beta) = \frac{\beta}{\alpha} \left(\frac{t}{\alpha} \right)^{\beta-1} \exp \left(- \left(\frac{t}{\alpha} \right)^{\beta} \right) \quad (1)$$

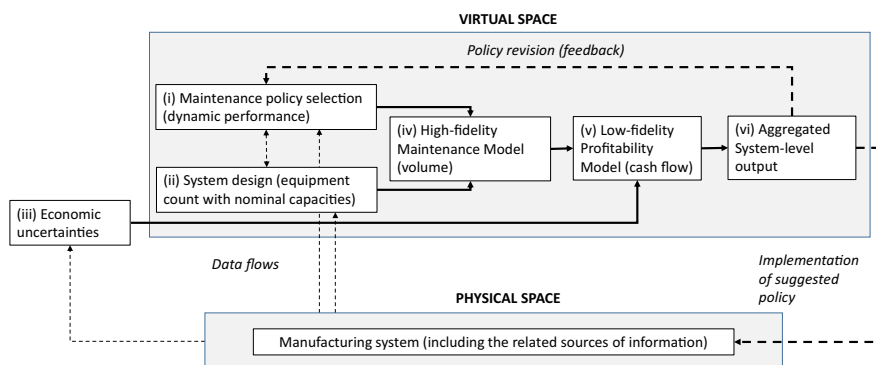


Fig. 3 A schematic illustration of the adopted modeling approach

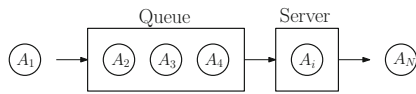


Fig. 4 A schematic representation of a server-queue entity (shovels, dump sites, and maintenance workshops)

where $\alpha > 0$ is the shape parameter, $\beta > 0$ is the scale parameter, and t is the time elapsed since the last maintenance intervention. The time to repair (TTR) is different in case of corrective and preventive maintenance, and it is modeled using a log-normal distribution

$$\text{LogN}(t; \mu, \sigma) = \frac{1}{t \sigma \sqrt{2\pi}} \exp\left(-\frac{(\ln(t - \mu))^2}{2\sigma^2}\right) \quad (2)$$

where the parameter μ is the mean of the distribution, σ is the standard deviation, and t is the duration of the maintenance intervention. Each truck is characterized by its capacity, which varies depending on the truck model, and a cost for both preventive C_P^T and corrective C_C^T maintenance interventions. From the practical point of view, the selected TTR-approach allows us to take advantage of the cumulated maintenance data as the peculiar characteristics of each piece of equipment can be represented.

An agent can be served by server-queue objects, i.e. by shovels, dumpsites, and workshops, which are modeled according to the well-know queueing theory (Law et al. 2000). A server-queue entity presents a waiting room, the so-called *queue*, where the agents, or *customers*, wait their moment to be served by the processor, the so-called *server*. Figure 4 gives a schematized representation of the server-queue object, where the agents are represented by the circles and they join the queue at an unknown arrival rate. Customers are served according to a *first in first out* (FIFO) logic at a serving rate that changes depending on the type of customer.

The three classes of server-queue components present subtle differences. Shovels were modeled as server-queue objects with log-normal serving time distributions, and they presented the peculiar hallmark of *unreliability*: as in the real world, they were subject to the aging process, hence they could unexpectedly fail, or they could be preventively maintained. Therefore, in addition to the serving time, shovels are characterized by a TTF probability density function, which is modeled using Equation 1. When a shovel becomes unavailable due to maintenance, it changes its behavior to that of an agent and it enters the maintenance workshop with maximum priority. The trucks in the queue wait for the shovel to be available again and no other trucks are assigned to that shovel until maintenance ends. Shovels are thus characterized by a cost for corrective C_C^S and preventive C_P^S maintenance in addition to TTF and TTR distributions. As soon as the

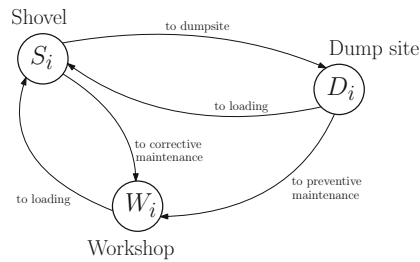


Fig. 5 Scheme of the agent's movements on a map inside the maintenance module between its sub-systems. S_i represents a generic shovel site, D_i represents a dump site, and W_i a workshop

maintenance activity is completed, the shovel is considered available again and trucks can start to join the queue and to be processed.

Workshops are characterized by a FIFO logic with priority for the management of the queue (shovels with maximum priority), and they present a peculiar behavior concerning the processing time of a customer, i.e. a truck or a shovel. The service time is a function of both the type of item served (truck/shovel) and the type of maintenance intervention, i.e. corrective or preventive. Finally, dumpsite components are characterized by a log-normal service time distribution and by the presence of a stockpile; each stockpile has a limited capacity and all the stockpiles feed a single concentrator plant with a specific capacity of material per unit of time. The detailed modeling of the concentrator plant, with equipment such as crushers, conveyor belts, mills, flotation tanks, etc., is left out of the scope of this paper and it is assumed to work without interruptions.

A DES experiment was designed to replicate system operations with a high level of detail. Within the simulation procedure, all the entities interact with each other as described and illustrated in Figure 5. The mine maintenance simulation is initialized by defining the parameters of the probability distributions; the TTF and TTR distribution parameters are listed, together with the costs for maintenance, in Table 3 and Table 4 in Appendix A. Trucks are also characterized by a transportation capacity, which is a random variable sampled from the distributions reported in Table 3 in Appendix A, whereas servers are characterized by a serving time distribution, which parameters are listed for shovels and dumpsites in Table 4 and Table 5 in Appendix A respectively. The parameters listed above remain, together with the duration of the simulation horizon, unchanged for all the runs of the experiment.

Once the simulation is initialized, a truck gets assigned to a target shovel S_i (see Figure 5), thus it travels to the designed site and it joins S_i 's queue. After being processed at a load site, a truck can leave the site due to two reasons: it can

either fail unexpectedly and thus being sent to a workshop W_i in Figure 5 for CM, or it can be sent to a dumpsite D_i in Figure 5 for unloading. After the unloading, a heuristic decides if the agent must be preventively maintained, or if it can continue its regular operation. The decision to submit a truck to PM is based on the age of the truck, namely a TBM policy is adopted. If the threshold value p_i for the i -th agent is lower than the time elapsed since the last maintenance intervention, it undergoes PM, otherwise, it is assigned to a new load site. The maintenance policy can be represented by a list, whose components p_i^j are the PM thresholds for trucks T and shovels S , and it can be represented as follow:

$$P = [p_1^T, p_2^T, \dots, p_{N_T}^T, p_1^S, p_2^S, \dots, p_{N_S}^S] \quad (3)$$

where N_T is the number of trucks, and N_S is the number of shovels in the system.

When a CM or PM intervention is due on a truck, a workshop W_i processes the agent according to the type of maintenance needed and to the TTR distribution of the specific agent. Once the truck has been maintained, its condition is considered “as good as new” from the modeling perspective and it is ready to start a new *mission*. A mission is defined as a chain of actions that includes the travel to a shovel site, the waiting time in queue, the loading and unloading operations.

The shovel’s mission is less detailed than a truck’s mission: each shovel simply operates at its site until a failure occurs, or until it is sent to a workshop W_i for PM. When a truck has been loaded, the age of the shovel is checked against the age threshold p_i^S and, in case the time elapsed since the last CM/PM intervention exceeds p_i^S , it is sent to a workshop W_i with maximum priority, thus preempting each other agent in the queue.

The performance of the system was optimized based on the results of a MCS experiment. Given the stochastic nature of a DES, the problem consists in the minimization of the expected value of the cost of operations $J(\theta)$, and it can be formalized as

$$Z = \min_{\theta} \mathbb{E}[J(\theta)]$$

where θ is a vector containing the system parameters that define the number of trucks N_T and shovels N_S , and the maintenance thresholds P . The problem must be solved under the constraint of reaching a production target X_{min} :

$$\Pr\{X \geq X_{min}\} \geq 0.95.$$

That is, the probability that the output X satisfies the target X_{min} must be greater than 95%; such probability can be calculated using the 95th percentile of the output distribution from the MCS experiment.

To minimize the objective function means to act on two aspects of the model: the number of resources operating in the system and the number of unplanned downtimes. The former is minimized using an enumerative search algorithm, while the second is optimized using a more complex genetic algorithm for search over a stochastic response surface. More details about both procedures are provided in Appendix B.

The code¹ used to implement the algorithms described above is written in Python 3.7 and mostly using *SimPy* simulation library.

Cash Flow Module

The use of system dynamics methodology allows building a compounded, close-to-reality representation of the mining operation that is still easy-to-read and modify compared to writing the model as software code. Detailed SD feasibility models of mining have been introduced in the literature by, e.g., Inthavongsa et al. (2016), Savolainen et al. (2017), who showed the flexibility of the approach and its ability to cope with complexity.

For the sake of brevity, the CF model used in this paper includes only two uncertainties: the metal price and maintenance cost. The simulation horizon is limited to one-year, and a geometric Brownian motion with and without trend is assumed to represent the uncertainty of markets adequately (for discussion see, e.g., Labys et al. (1999), Roberts (2009), Rossen (2015)). An example price simulation used in the experiments is illustrated in Figure 6 with three alternative price trend scenarios for a single random price realization. The uncertainty of maintenance costs are modeled as triangular distributions using expert estimates, which are introduced in more detail in Section 5.2.

A representation of the function block diagram of the applied CF model is provided in Figure 7. The model is divided into two sections: technical and economic models, where the inputs of the mine maintenance module are fed (blocks of the flowsheet marked with blue background). For a full list of parameters see Table 6 in Appendix C.

One of the key output variables of the CF model is the average mill utilization rate. That is, at any point in time, the mill utilization rate is either zero or one depending on the level of the ore stock that is replenished by the truck-shovel system. We exclude the option to increase the size of temporary ore stock giving additional flexibility to maintenance timing, which is often used in small mines. In our case, the stock is limited to $\approx 27,500$ tons of ore, which corresponds roughly to 36 hours of production in the mill.

¹ All the libraries used to realize the simulation optimization experiment are released under a MIT license, and a copy of the code and the relative documentation is freely available at https://github.com/mikiurbi/mine_digital_twin.

Fig. 6 A random metal price array simulated using geometric Brownian motion with 10% volatility and three alternative trend assumptions ($\pm 10\%/yr$ and 0)

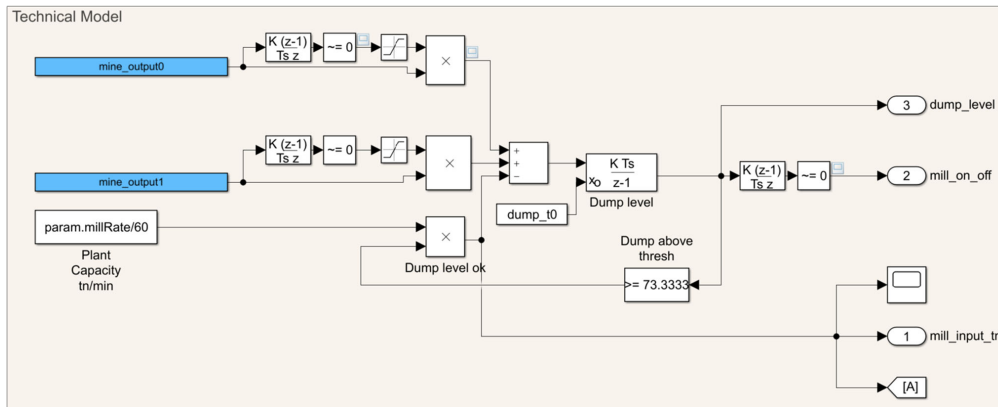
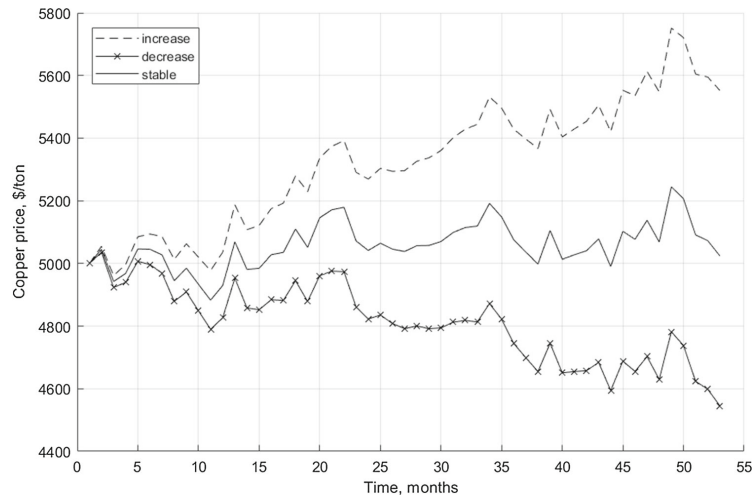


Fig. 7 A flowsheet of the managerial CF model to estimate the economic feasibility of the operations. The input values from the O&M module are marked using a grey background

The key added variables from the CF model include the costs of equipment leasing, fuel costs derived based on the O&M model's indicated operation hours, and other fixed costs such as buildings, and administration. The output price of metal is updated weekly.

We acknowledge that the above-described model construction, including the detailed operations & maintenance model using discrete event simulation and system dynamic cash flow modeling, could be fully implemented in a single software environment. In practice, this is usually not possible, which calls for the DT type of co-simulation approach. The reasons for this can be related to an unwillingness to share confidential financial information (from mine operator

to the model owner), the effort of transferring existing pieces of core software libraries from one environment to another, and importantly, as pointed out by, e.g. West and Blackburn (2017), the uncertainty of financial return of the software product.

Model and Application

In this section, the O&M simulation optimization module's behavior is first validated with two sensitivity analyses and then used in concert with the SD model to run three experiments in a DT system setting. The mine configuration in

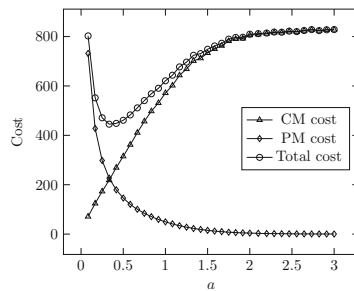


Fig. 8 Sensitivity analysis for the trade-off between CM and PM cost in the maintenance module

these tests varies from one experiment to the other, and the number of components in the system is kept low to avoid over-parametrization of the model (see discussion, e.g., (Zio 2009)).

Maintenance Model Validation

A simple sensitivity analysis was performed to validate that the lower the age threshold for PM, the higher the possibility to avoid unplanned downtimes. On the other hand, the higher the age thresholds, the less effective PM should be in reducing costs. To validate this hypothesis, the maintenance thresholds of all the items were parameterized as follows:

$$p_i^j = a MTBF_i,$$

where i identifies the item, j is the class T for trucks or S for shovels, and $MTBF_i$ is the mean time between failure of the i -th item. The parameter $a \in (0, 3]$ is a scale factor that allows to vary the age threshold p_i^j of all the equipment included in the experiment in question. By parametrizing the age thresholds, it was possible to estimate both the cost of CM and PM for the whole system by changing only the parameter a . In all the other experiments, the cost of CM/PM depends also on C_C , C_P , and on the TTF distributions, but here these parameters are fixed. The sum of the cost of PM and CM at different values of a are plotted in Figure 8.

The fleet used to realize the sensitivity analysis included two trucks and one shovel. When the age thresholds are very low ($\ll MTBF$) the cost of PM is high because PM events are carried out extremely often. However, the cost of PM decreases sharply when a increases and, with maintenance thresholds p_i^j equal to 0.5 times the $MTBF$ values, the cost of CM starts to be higher than the cost of PM, thus making it inconvenient to perform PM more rarely. As it is depicted in Figure 8, the total cost of maintenance presents a minimum cost as a function of PM and CM, which makes clear the

need to optimize the PM age threshold of all the items before running the whole simulation procedure.

A second maintenance model validation experiment was carried out to test the performance of the system with varying maintenance resources; in particular, the difference between two- and three-workshops configurations were analyzed. A total of sixty configurations were tested, namely all the possible combinations of 2 or 3 workshops, 1 to 10 trucks, and 1 to 3 shovels. The statistics used to present the results are the average throughput and the average cost of maintenance obtained from 50 simulations over a 2-year time horizon. For each configuration, the maintenance thresholds p_i^j are optimized and then the DESs are run.

Since the dumpsites present limited capacity, i.e. material excavated cannot exceed the mill production rate, the configurations with two and three maintenance workshops produce different results. As shown in Figure 9, many system configurations deliver the maximum possible amount of material, but at different costs. Interestingly, highly different configurations lead to similar results: for instance, the 2-workshops 3-shovels and the 3-workshops 1-shovel configurations deliver almost the same throughput at the same cost using a similar number of trucks. The two solutions are however very different from a managerial point of view: the investments required to purchase or to rent the equipment, the skilled personnel needed to operate the facilities, and the resilience of the resulting system are meaningful aspects to be considered.

The above-mentioned issues go beyond the reasonable modeling scope of the DES, but these are the issues that can be easily integrated into the managerial profitability model to produce further insights to support operational decision-making. In a dynamic economic environment, provided by the SD, the proposed analysis can be repeated with better implicit knowledge of the production process, such as production targets and planned maintenance, thus producing a probabilistic evaluation of the future scenarios.

Digital Twin Testing and Validation

The first experiment aimed to verify and validate the overall DT approach, whereas the second experiment consisted of a more detailed optimization of the system under the assumption of uncertain maintenance costs. The parameters used in the CF model are illustrative, and they were chosen in a way that approximately 5-6 trucks (max. 10) with 1-2 shovels (max. 3) would satisfy the mill requirements for material tonnage.

The fleet design study was carried out to screen all the possible system configurations that can be produced by ten trucks, three shovels, and the maintenance policies provided in Table 2. In Table 2, the item 'design optimization' indicates if the number of trucks and shovels is set already in

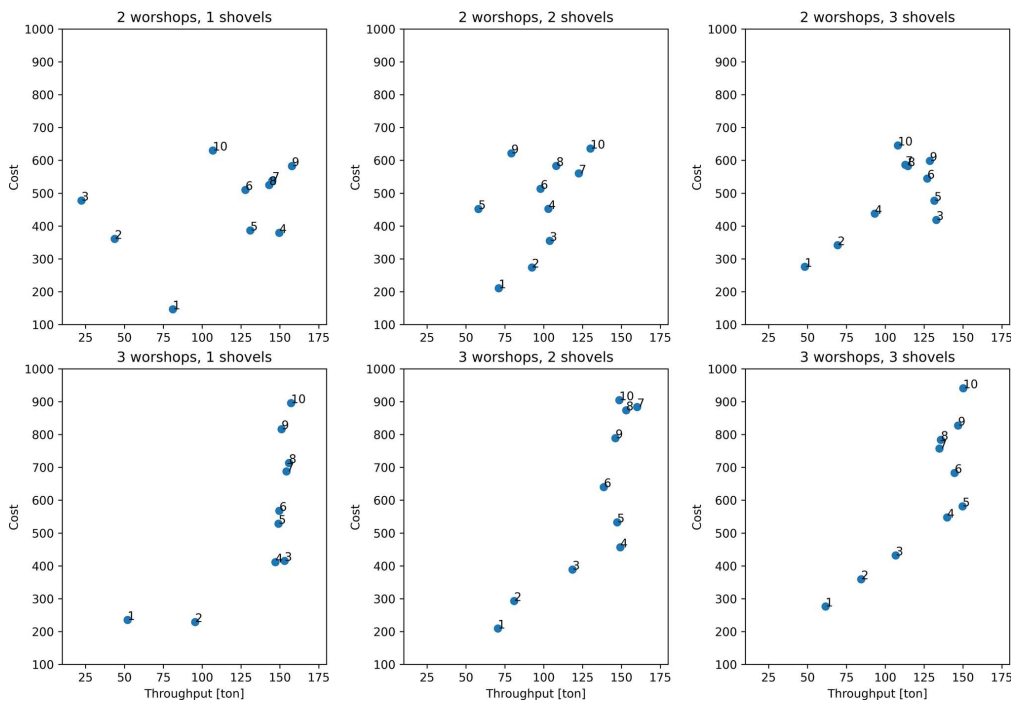


Fig. 9 A multi-criteria comparison of different system configurations. Each point corresponds to a system configuration (number of workshops, shovels, and trucks), and it is characterized by the average cost of maintenance and the average throughput of metal obtained over the simulation horizon

the discrete event simulation. The configurations that have a smaller number of trucks than shovels were discarded manually as irrelevant, when the optimization option was turned off. Three different maintenance policies were used: a “max-corrective” (or “run-to-failure”), a balanced, and a “max-preventive” policy. The first and last policy represent the theoretical endpoints of the available scale of the simulation space: according to the “max-preventive”, a PM action is performed after every mission, whereas according to the “max-corrective” the maintenance threshold is set (de-facto) to as infinite. The balanced policy foresaw one PM event per week of simulation. In a more advanced setting, the balanced policy could be defined by the simulation-optimization algorithm, which searches for the PM thresholds p_i^j that return the minimum expected cost of maintenance for a given configuration. In Experiment 1, maintenance costs were assumed to be fixed (known ex-ante), and the price trends were those displayed in Figure 6. The number of simulations displayed in the last row of Table 2 was determined by the number of combinations, e.g., in experiment one with ‘pre-optimized’ fleet design, the number of combinations to be simulated was

nine as only the number of maintenance policies multiplied by the number of price trends.

Experiment 1 - Fleet design

Results of Experiment 1 without fleet optimization are provided in Figure 10 which shows that the policies “max-corrective” and “balanced”, with 6–10 trucks and 1–2 shovels, would be the most profitable ones. It is noticeable that with the given parameters of fixed costs, and duration of PM- and CM-events, the “max-corrective” policy was favored over the balanced option. The “max-preventive” policy produced negative profits in all cases within the selected set of parameters, which highlighted the need for further maintenance threshold optimization. That is, in this case, to maximize the amount of preventive maintenance leads to the loss of overall cost efficiency due to excess queuing times to the workshop, whereas increasing the number of workshops in the initial design would also have a bloating effect on operation’s costs.

The maximum profit of all tested configurations was reached following the “run-to-failure” policy, which yielded

Table 2 Key parameters used in the experiments

Experiment	Exp. 1	Exp. 2
Experiment Name	Fleet design	Age-threshold optimization
Design optimization	No (Yes)	-
Maintenance policy	{“max-corrective”, “max-preventive”, “balanced”}	“balanced”
Maintenance cost	{Fixed, Uncertain}	Uncertain
PM frequency (estimate), events/wk	1	{0.125, ..., 1} with step size of 0.125
Price trend	{“increase”, “decrease”, “none”}*	{“increase”}*
# Simulations	248 (9)	221

*: incl. 10%/yr volatility, where “increase” = +10%/yr, “decrease” = -10%/yr, “none” = 0%

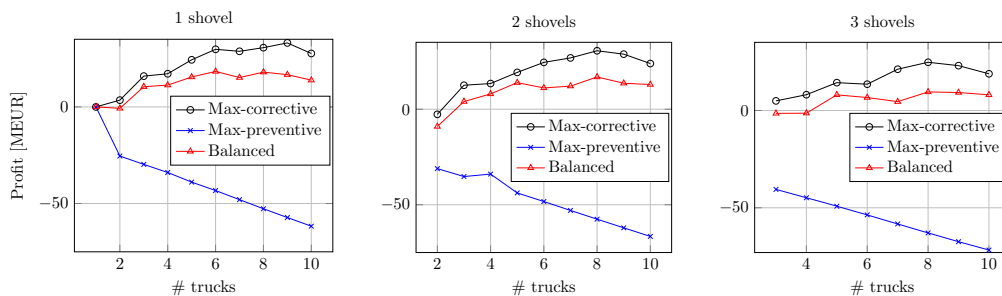


Fig. 10 Experiment 1 results. The profits in millions of units of money over 52 weeks are plotted according to a maintenance policy and separated by number of shovels used in the experiment

a 98% utilization rate of the mill with nine trucks and one shovel, also denoted using the set {9, 1, 98%}. In ranking of results, the price trend (see Table 7 in Appendix D) had a clear effect: increasing price trend (+10%/year) would suggest the {8, 2, 97%} as the second most desirable combination, whereas decreasing and flat trends (-10% and 0) would favor an option for the smaller fleet and lower mill utilization rate {8, 1, 91%}.

To inspect the results of Experiment 1 in more detail, the total number of maintenance events in the case of full PM are plotted in Figure 11. Figure 11a shows that the “max-preventive” policy, with a two-workshops design, is possible only in the case of one truck and one shovel. As the number of trucks increases, the relative share of CM actions goes up since there is not enough capacity in the workshops (Figure 11b) for adequate equipment intake.

Such effect is further highlighted in the case of one shovel and nine or ten trucks: the queuing time spent by trucks (either at loading, unloading, or maintenance) increases, thus making it more probable for them to fail before the next scheduled PM event.

In Experiment 1, the simulation-optimization algorithm in the O&M model was also tested to screen out the infeasible fleet designs already at the beginning of the simulation. The

maintenance optimization, as designed, favored the high mill utilization rate options that were gained with eight to ten trucks and one or three shovels (Table 10).

It is clear that the results of maintenance optimization efforts are uncertain ex-ante even with the assumption of fixed maintenance costs. To take a further step, the role of cost uncertainty was included in the analysis by replacing the fixed maintenance costs with triangular probability distributions in the CF model. These distributions are depicted in Figure 12 using box-plot diagrams (for numerical values see Appendix C), and they represent expert knowledge. In a real case, these distributions could be derived from proprietary maintenance data that are available from the organization’s historical records.

The applied cost distributions of PM and CM differ in shape. To reflect the risk of CM, the distributions of costs have long tails that can produce up to five times the fixed cost, whereas the positive risk of CM is limited to 1% respectively. The PM cost distribution is weighted more in the center of the distribution, thus giving less uncertain results on costs; it is capped to a maximum of 1.1 times the original assumption, and it can go below -10%. From the modeling point of view, this setting creates a strong incentive towards accepting preventive policies over the “max-corrective” given in

Fig. 11 The limited capacity of workshops illustrated using the simulation statistics. Scenarios are taken from results with full preventive maintenance policy settings using an increasing-price array

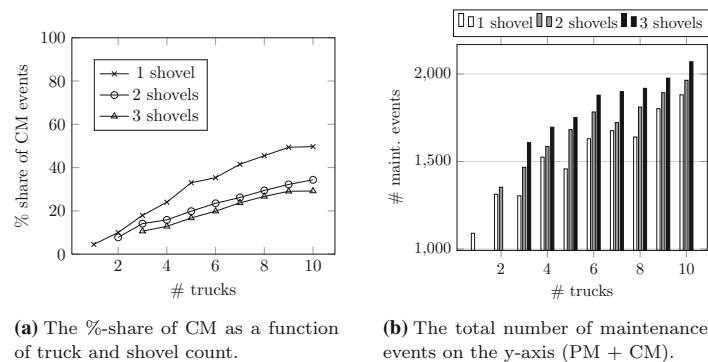
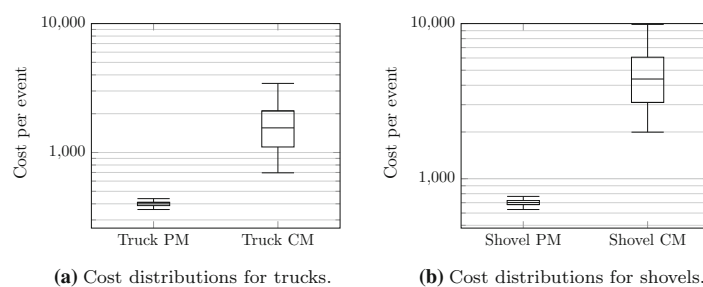


Fig. 12 Box-plot of probability distributions for the prices of maintenance events using 1000 random draws



the first experiment. Another option for the inclusion of cost uncertainty could be to include them in the O&M optimization module that is run before the CF simulation. However, this creates additional complexity to the genetic algorithm used for O&M-model's simulation optimization, and it complicates the user's abilities to interact with the CF simulation experiments within the SD-flowsheet that steers the CF model. A trade-off between model choice is made by using averages of distributions based on 10,000 draws in the optimization model (see Appendix B) and a single random draw in the CF model, which makes it more volatile in terms of results.

Running Experiment 1 with uncertain costs returned the previously suggested outcome {9, 1, 98%} with no PM; this was due to the limited workshop capacity as previously discussed (see Figure 11). Therefore, the question of optimal maintenance policy boils down to finding out whether and what is the optimal time between maintenance events that would keep the amount of CM within reasonable limits.

Experiment 2 – Optimal Timing for Preventive Maintenance

The issue of optimizing preventive maintenance thresholds is addressed in this last experiment. The age threshold, marked as n , for the PM event timing was set as a ratio versus one

round of simulation of the maintenance module. That is, a value of $n = 1$ means that there is approximately one preventive maintenance event per week of simulation in the O&M model, and $n = 0.5$ indicates one PM event every two simulated weeks respectively. In this experiment, n varies from 0.125 to 1.000 with a step size of 0.125.

The visual insights are provided in Figure 13, which shows that the {9, 1, 91.6%} combination with $n = 1$ yields a profit of approximately 30 million. On the other hand, the simulation indicates another option with two trucks less, namely {7, 1, 81.7%} with $n = 0.875$, that has only some 0.5 million profit less than the “best option”. From the decision-making point of view, we can observe that the interpretation and efficient utilization of simulation results becomes difficult because of the high number of co-existing solutions. In this vein, as previously discussed by, e.g., Negahban and Smith (2014), Min et al. (2019), there is an increasing demand to implement meta-model based solutions to simplify simulation data of DT-models into implementable managerial insights. This issue is, however, beyond the scope of this paper.

As depicted in Figure 13, maintenance policies with $n \geq 0.5$ seem to deliver the best performances in terms of profitability and utilization and, furthermore, a positive relationship between the mill utilization and a high num-

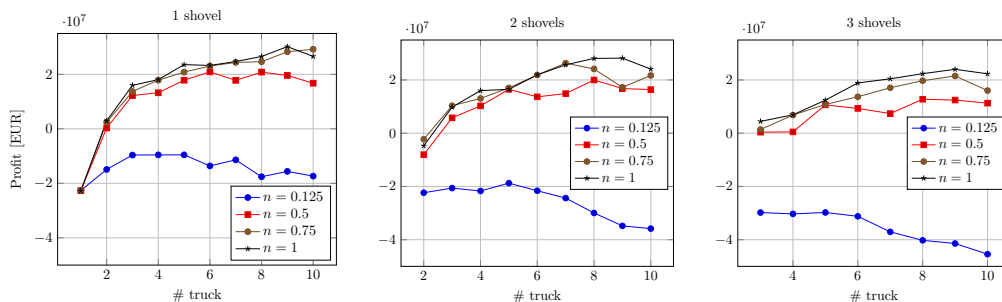


Fig. 13 Profit as a function of fleet configuration and maintenance policy n , where n represents the number of preventive maintenance events per time unit (in this case one week)

ber of trucks exists. On the other hand, based on Figure 13, a reduced number of shovels accounted for an increase of profitability. To verify when adding further equipment is no more beneficial from the profits point of view, Figure 14 is laid out, which depicts the “the profit per truck” as a function of the number of trucks in the system with varying maintenance policy thresholds and shovel counts. Figure 14 suggests that selecting a reduced fleet size could increase the overall profitability of operations (see the middle subplot). The unprofitable configurations, i.e. those showing an unbalanced amount of trucks and shovels, locate themselves into the lower-left quadrants of the plots, where profits are negative. Moreover, Figure 14 highlights that, when investing in three shovels, one should prefer maintenance policies with $n \approx 1.0$, whereas a smaller fleet size could perform at its highest with a $n \in [0.8, 0.9]$.

Result Summary and Analysis

The limited set of results shows that in the mining industry the co-simulation of detailed O&M models for fleet operations in connection with the managerial CF models may provide a way to improve the overall profitability. The potential for economic improvement in this study was brought by the possibility to combine fleet design and maintenance policy, which, based on the detailed simulation, could be able to produce a comparable high revenue per invested unit of money without simply striving to maximize either the rate of excavation or mill utilization that are usually viewed as the most important key performance indicators.

In the model demonstration, a total time frame of one year with weekly simulations in the O&M model was used. For real-world applications with constantly updating operational data, a more frequent optimization might be a more interesting choice, where the managerial considerations could involve the choice of alternative mine plans with less focus on the current market price.

Discussion and Conclusions

This paper presented a DT modeling approach aimed to dynamically optimize the O&M in the mining industry with respect to the uncertain price of the end-product. A two-stage simulation was adopted: firstly, a selected set of mining equipment with their unique failure distributions was chosen, and its capabilities to meet the production demand in the concentrator plant was simulated using discrete event simulation. In the second phase, the aggregated information provided by the DES was fed to the managerial CF model with costs fully accounted to evaluate the overall economic optimum from the operations perspective.

The DES experiment produced useful information on the availability and utilization of mining equipment, and we were able to connect the aggregated output data with the overall profitability of the business via the managerial CF model to form a “digital twin of manufacturing” in mining. Within the limits of the selected simulation values, it seems that in most cases maximizing mill utilization and production throughput would yield the highest expected revenues regardless of the realized price array, but some opportunities for downsizing may exist. The results highlight the role of maintenance as a necessary evil with only little potential to the economic upside. However, these results are highly dependent on the selected values and should be re-evaluated using a more extensive sensitivity analysis of the DT presented or a case example of a real mine.

Some implications to managerial decision making of the mining industry can be suggested. First, simulation-based DT modeling can perform overall operational optimization while considering the stochasticity and high dimensionality of operational data, which, as the second point, allows simultaneous consideration of operational and investment decisions. From the methodological point of view, the whole concept of DT is still taking its shape, and often it is used just to bring extra buzz to promote one’s simulation efforts.

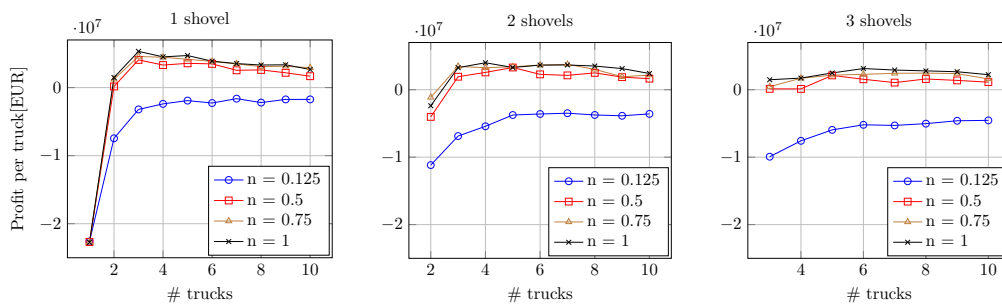


Fig. 14 The profit per truck is plotted as function of the number of trucks for different maintenance policies

We used the term specifically to refer to a holistic system optimization model that consists of several, data-based, self-reliant, and discipline-specific, simulation models. These models run parallelly and they are connected in the virtual space, which could have a (near) real-time connection with the physical process. Whether in the future the approach would be labeled DT or not, the models evidently have novelty value beyond simpler simulations applied today. Insights gained from our modeling efforts corroborate the earlier findings and discussion around the problems in the implementation of DTs. These issues include, but are not limited to, i) the dimensionality and the degrees of freedom of data in large models, ii) a suitable level of data aggregation when transferring it between discipline models, and iii) the validation and verification (V&V) of the results.

A central limitation related to the V&V of this paper is the absence of a true physical mining system; on the other hand, fully virtual modeling enabled testing and running experiments rapidly for scientific research. It should be also acknowledged that the TTF distributions are not able to fully capture the reality of maintenance: additional details could be added in the O&M model by making the failure probabilities dynamic or component-based to reflect the agents' actions better than the static ones used here. As future avenues of research, a case-specifically tailored DES-model could be taken into action, which is connected to the relevant data-gathering systems for up-to-date information as well as the existing CF models. On the other hand, rather than increasing the details of modeling, there is an emerging need to develop meta-model based methods that can interpret the results of DT-simulation in a fast and managerially digestible manner.

Acknowledgements We want to acknowledge the anonymous reviewers for their work, their time, and their comments. These were an important contribution to improve the quality of our work.

This work was supported by the Strategic Research Council (SRC) at the Academy of Finland through the Manufacturing 4.0-project [grant numbers 313349, 313396]. This project has also received funding from the Italian Ministry for Education, University, and Research (MIUR) through the "Departments of Excellence" program.

Funding Open Access funding provided by LUT University (previously Lappeenranta University of Technology (LUT)).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Experiment Parameters

The parameters used to characterize the components used in the simulation experiments are reported in the following. The TTF of trucks and shovels are the same used in (Mena et al. 2013).

Table 3 Parameters for characterization of truck objects. Legend: the letters α and β indicate the shape and scale factor of a two-parameter Weibull distribution

Quantity		TTF		TTR PM		TTR CM		Capacity		C_C^T	C_P^T
Distribution Type		Weibull		Log Normal		Log Normal		Gaussian			
Paramter	Id	α	β	μ	σ	μ	σ	μ	σ		
	1	45	2.1	1.44	0.99	1.87	1.28	220	5	500	60
	2	55	1.4	1.36	1.06	1.76	1.37	220	5	500	60
	3	34	2.5	1.57	0.96	2.04	1.24	220	5	500	60
	4	44	1.3	1.50	1.22	1.95	1.58	220	5	500	60
	5	34	1.9	1.13	0.92	1.46	1.19	220	5	500	60
	6	50	1.5	1.46	1.03	1.89	1.33	220	5	500	60
	7	52	2.3	1.55	1.10	2.01	1.43	220	5	500	60
	8	41	2.1	1.25	0.99	1.62	1.28	220	5	500	60
	9	38	1.3	1.22	1.02	1.58	1.08	220	5	500	60
	10	32	1.6	1.34	1.25	1.74	1.62	220	5	500	60

The letters μ and σ represent the mean and the standard deviation of a Gaussian and a log-normal distribution. C_C^T and C_P^T are respectively the cost of corrective and preventive maintenance of a truck

Table 4 Parameters for the characterization of shovel objects. Legend: the letters α and β indicate the shape and scale factor of a two-parameter Weibull distribution

Quantity		TTF		TTR PM		TTR CM		Capacity		C_C^S	C_P^S
Distribution Type		Weibull		Log Normal		Log Normal		Gaussian			
Paramter	Id.	α	β	μ	σ	μ	σ	μ	σ		
	1	196	2.0	1.44	0.99	1.87	1.2	2.33	0.26	2000	140
	2	187	2.3	1.48	1.19	1.92	1.25	2.21	0.21	2000	140
	3	200	1.9	1.35	1.10	1.79	1.18	2.27	0.19	2200	170

The letters μ and σ represent the mean and the standard deviation of a log-normal distribution. C_C^S and C_P^S are respectively the cost of corrective and preventive maintenance of a truck

B Simulation Optimization Details

To optimize the configuration of the system means to identify the number of trucks and shovels that guarantees to achieve the production target X at the minimum cost. The response surface $\mathbb{E}[J(\theta)]$ is obtained through an MCS experiment and the effect of constrained resources – i.e. the number of maintenance workshops and the maximum capacity of stockpiles – is not easily predictable. We observed that the number of available workshops makes the increase in throughput non-monotonic with respect to an increase in the number of operating trucks and shovels. Therefore, the possibility to guide the search according to some heuristic must be dropped, and given the relatively low number of possible combinations, an enumerative search procedure was implemented. Once the average values of throughput produced by the system were available, the solutions were ranked in ascending order concerning the expected cost of maintenance $\mathbb{E}[J(\theta)]$.

Maintenance thresholds are continuous real variables, which are constrained to be positive. Since there are no

Table 5 Parameters for the characterization of dumpsites and workshops

Quantity		Serving Time		Coordinates	
Distribution Type		Log Normal			
Parameter	Id.	μ	σ	x	y
Dumpsites	1	1.13	0.15	6	5
	2	1.2	0.17	7	3
Workshops	1	–	–	3	0
	2	–	–	6	0

Legend: letters μ and σ represent the mean and the standard deviation of a log-normal distribution

other constraints, an effective tool to find a minimum of the response surface $\mathbb{E}[J(\theta)]$ is the genetic algorithm (GA). Each individual of the population was represented by a vector \mathbf{P} , which was introduced in Eq. (3), and the size of the initial population was set to 50 individuals. The selection of parents occurred according to the “fitness proportionate selection” mechanism, whereas single-point crossover operations and

mutation operations were performed with a preference for the latter (a ratio of 4 mutations to 1 crossover was used). Crossover operations consisted in the selection of a random element of two vectors P , at which the vectors are separated in a head and a tail, and subsequently the tails are swapped. Mutation operations occurred with a probability of 0.2 for each element of P , and when the mutation occurred a quantity sampled from a normal distribution $\sim \mathcal{N}(0, 0.1)$ was added to the element. The stopping criterion used in the GA was the limit on the number of generations, which was set equal to 25. The GA was indeed tested using a different number of generations up to 100 and the algorithm showed no meaningful improvement of the best solution found after 30

generations. The number of DESs required to obtain a reliable estimate of the response surface value returned by an individual was estimated in 50 repetitions. Such requirement makes the execution of the algorithm computationally demanding, therefore a 64-bit Windows Server 2016, Intel Xeon Platinum 8160 CPU 2.10 GHz, with 768 GB of RAM workstation was used to run the algorithm; a single optimization required on average 30 min.

C Cash Flow Model Parameters

Table 6 Parameters used for CF modeling in the simulation experiments. If the cell is empty as the experiment column number increases, then the last given value of the row in question prevails

Parameter	Unit	Exp. 1(a)	Exp. 1(b)	Exp. 2
Simulation timeframe	year	1	-	-
Technical Parameters				
Mill capacity	tn/h	1100	-	-
Cu-content	%/tn	5	-	-
<i>Fuel consumption</i>				
Truck moving	liters/h	100	-	-
Truck idle	liters/h	10	-	-
Shovel operating	liters/h	50	-	-
Shovel idle	liters/h	5	-	-
Economic Parameters				
Cu-price (at $t = 0$)	\$/tn	5000	-	-
Cu-volatility	%/yr	10	-	-
Cu-trend	%/yr	{-10, 0, 10}	-	10
Production unit cost	%/yr	1	-	-
Payable factor	%/tn of Cu	10	-	-
Fuel cost	\$/litre	0.8	-	-
<i>Trucks</i>				
Leasing	\$/pc/month	10000	-	-
Salary cost	\$/h/pc	100	-	-
PM event cost	\$/pc	400	{360, 400, 440} ¹	-
CM event cost	\$/pc	700	{693, 700, 3500} ²	-
PM frequency, n		{0, 1, ∞ }	-	{0.125, ..., 1} with 0.125 step
<i>Shovels</i>				
Leasing	\$/pc/month	10,000	-	-
Salary cost	\$/h/pc	120	-	-
PM event cost	\$/pc	700	{630, 700, 770} ³	-
CM event cost	\$/pc	2000	{1980, 2000, 10,000} ⁴	-
PM frequency, n		{0, 1, ∞ }	-	{0.125, ..., 1} with 0.125 step

^{1,2,3,4}: Triangular distribution; average of 10000 draws: 400¹, 1626², 700³, 4659⁴

D Influence of Price Trends

Table 7 Relative ranking of equipment design options with different price assumptions, when using run-to-failure maintenance policy

Trend	Increase			Decrease			No trend		
	1	2	3	1	2	3	1	2	3
#Shovels									
#Trucks									
1	−22.58			−22.61			−22.60		
2	5.06	−1.05		1.97	−4.30		3.46	−2.73	
3	18.44	15.17	0.00	13.48	9.97	2.40	15.88	12.48	4.86
4	20.02	16.34	11.01	14.34	10.58	5.25	17.09	13.37	8.03
5	27.83	22.74	17.77	20.99	15.83	10.88	24.30	19.17	14.22
6	33.86	28.56	17.20	25.84	20.53	9.87	29.72	24.41	13.42
7	32.96	31.13	25.67	24.67	22.61	16.84	28.68	26.73	21.11
8	35.29	35.34	29.58	26.13	25.92	20.07	30.56	30.48	24.67
9	38.10	33.84	28.06	28.26	23.90	18.12	33.01	28.70	22.92
10	32.67	28.92	23.88	22.75	18.97	13.94	27.55	23.78	18.74

References

- Aivaliotis, P., Georgoulas, K., & Chrysosouris, G. (2019). The use of digital twin for predictive maintenance in manufacturing. *International Journal of Computer Integrated Manufacturing*, 32(11), 1067–1080.
- Alrabghi, A., & Tiwari, A. (2015). State of the art in simulation-based optimisation for maintenance systems. *Computers & Industrial Engineering*, 82, 167–182.
- Alrabghi, A., & Tiwari, A. (2016). A novel approach for modelling complex maintenance systems using discrete event simulation. *Reliability Engineering & System Safety*, 154, 160–170.
- Angeliu, G., Coronelli, D., & Cardani, G. (2020). Development of the simulation model for digital twin applications in historical masonry buildings: The integration between numerical and experimental reality. *Computers and Structures*, 238, 106282.
- Ashtari Talkhestani, B., Jung, T., Lindemann, B., Sahlab, N., Jazdi, N., Schloegl, W., et al. (2019). An architecture of an intelligent digital twin in a cyber-physical production system. *At-Automatisierungstechnik*, 67(9), 762–782.
- Bevilacqua, M., & Braglia, M. (2000). The analytic hierarchy process applied to maintenance strategy selection. *Reliability Engineering & System Safety*, 70(1), 71–83.
- Boafo, E., Nasimi, E., Zhang, L., & Gabbar, H. A. (2017). Co-simulation for real time safety verification of nuclear power plants. *Mechanical Engineering Journal*, 4(2), 15–00532.
- Boschert, S., & Rosen, R. (2016). Digital twin-the simulation aspect. In *Mechatronic futures: Challenges and solutions for mechatronic systems and their designers* (pp. 59–74). Cham: Springer.
- Castano, F., Strzelczak, S., Villalonga, A., Haber, R. E., & Kosakowska, J. (2019). Sensor reliability in cyber-physical systems using internet-of-things data: A review and case study. *Remote Sensing*, 11(19), 2252.
- Cattaneo, L., & Macchi, M. (2019). A digital twin proof of concept to support machine prognostics with low availability of run-to-failure data. *IFAC-PapersOnLine*, 52(10), 37–42.
- Cho, D. I., & Parlar, M. (1991). A survey of maintenance models for multi-unit systems. *European Journal of Operational Research*, 51(1), 1–23.
- Cimino, C., Negri, E., & Fumagalli, L. (2019). Review of digital twin applications in manufacturing. *Computers in Industry*, 113, 103130.
- de Jonge, B., & Scarf, P. A. (2020). A review on maintenance optimization. *European Journal of Operational Research*, 285(3), 805–824.
- Dekker, R., Wildeman, R. E., & Van der Duyn Schouten, F. A. (1997). A review of multi-component maintenance models with economic dependence. *Mathematical Methods of Operations Research*, 45(3), 411–435.
- Dhillon, B. S. (2008). *Mining equipment reliability* (pp. 57–70). London: Springer.
- Dowdeswell, B., Sinha, R., & MacDonell, S. (2020). Finding faults: A scoping study of fault diagnostics for industrial cyber-physical systems. *Journal of Systems and Software*, 168, 110638.
- Du, J., Zhu, Q., Shi, Y., Wang, Q., Lin, Y., & Zhao, D. (2020). Cognition digital twins for personalized information systems of smart cities: Proof of concept. *Journal of Management in Engineering*, 36(2), 04019052.
- Forrester, J. W. (1961). *Industrial dynamics*. Cambridge: MIT Press.
- Forrester, J. W. (1994). System dynamics, systems thinking, and soft OR. *System Dynamics Review*, 10(2–3), 245–256.
- Fotland, G., Haskins, C., & Rølvåg, T. (2020). Trade study to select best alternative for cable and pulley simulation for cranes on offshore vessels. *Systems Engineering*, 23(2), 177–188.
- Gomes, C., Thule, C., Broman, D., Larsen, P. G., & Vangheluwe, H. (2017). Co-simulation: State of the art. [arXiv:1702.00686](https://arxiv.org/abs/1702.00686).
- Gomes, C., Thule, C., Broman, D., Larsen, P. G., & Vangheluwe, H. (2018). Co-simulation: A survey. *ACM Computing Surveys*, 51(3).
- Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In *Transdisciplinary perspectives on complex systems: New findings and approaches* (pp. 85–113). Cham: Springer International Publishing.
- Größler, A., Thun, J. H., & Milling, P. M. (2008). System dynamics as a structural theory in operations management. *Production and Operations Management*, 17(3), 373–384.
- Hall, R. A., & Daneshmend, L. K. (2003). Reliability and maintainability models for mobile underground haulage equipment. *CIM Bulletin*, 96(1072), 159–165.
- He, R., Chen, G., Dong, C., Sun, S., & Shen, X. (2019). Data-driven digital twin technology for optimized control in process systems. *ISA Transactions*, 95, 221–234.

- Inthavongsa, I., Drebenstedt, C., Bongaerts, J., & Sontamino, P. (2016). Real options decision framework: Strategic operating policies for open pit mine planning. *Resources Policy*, 47, 142–153.
- Kaewunruen, S., & Xu, N. (2018). Digital twin for sustainability evaluation of railway station buildings. *Frontiers in Built Environment*, 4, 77.
- Kendrik, Y., Hong, L., Pai, Z., & Chun-Hsien, C. (2020). A state-of-the-art survey of digital twin: Techniques, engineering product lifecycle management and business innovation perspectives. *Journal of Intelligent Manufacturing*, 31, 1313–1337.
- Kim, N. H., An, D., & Choi, J. H. (2016). *Prognostics and health management of engineering systems: An introduction*. Cham: Springer International Publishing.
- Kontes, G. D., Giannakis, G. I., Sánchez, V., Agustin-Camacho, D., Romero-Amorrotu, A., Panagiotidou, N., et al. (2018). Simulation-based evaluation and optimization of control strategies in buildings. *Energies*, 11(12), 3376.
- Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine* 51(11):1016–1022, 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018.
- Kunath, M., & Winkler, H. (2018). Integrating the digital twin of the manufacturing system into a decision support system for improving the order management process. *Procedia CIRP*, 72, 225–231.
- Labys, W., Achouch, A., & Terraza, M. (1999). Metal prices and the business cycle. *Resources Policy*, 25(4), 229–238.
- Law, A. M., Kelton, W. D., & Kelton, W. D. (2000). *Simulation modeling and analysis* (Vol. 3). New York: McGraw-Hill New York.
- Li, L., Mao, C., Sun, H., Yuan, Y., & Lei, B. (2020). Digital twin driven green performance evaluation methodology of intelligent manufacturing: Hybrid model based on fuzzy rough-sets AHP, multistage weight synthesis, and PROMETHEE II. *Complexity*, 2020, 3853925.
- Liu, Z., Chen, W., Zhang, C., Yang, C., & Chu, H. (2019). Data super-network fault prediction model and maintenance strategy for mechanical product based on digital twin. *IEEE Access*, 7, 177284–177296.
- Lu, Q., Parlikad, A. K., Woodall, P., Ranasinghe, G. D., Xie, X., Liang, Z., et al. (2020). Developing a digital twin at building and city levels: Case study of West Cambridge campus. *Journal of Management in Engineering*, 36(3), 05020004.
- Madni, A. M., Madni, C. C., & Lucero, S. D. (2019). Leveraging digital twin technology in model-based systems engineering. *Systems*, 7(1), 7.
- Mena, R., Zio, E., Kristjanpoller, F., & Arata, A. (2013). Availability-based simulation and optimization modeling framework for open-pit mine truck allocation under dynamic constraints. *International Journal of Mining Science and Technology*, 23(1), 113–119.
- Mermertas, U., Hagemann, T., & Brichart, C. (2019). Optimization of a 900 mm tilting-pad journal bearing in large steam turbines by advanced modeling and validation. *Journal of Engineering for Gas Turbines and Power*, 141(2), 021033.
- Min, Q., Lu, Y., Liu, Z., Su, C., & Wang, B. (2019). Machine learning based digital twin framework for production optimization in petrochemical industry. *International Journal of Information Management*, 49, 502–519.
- Mobley, R. K. (2002). *An introduction to predictive maintenance*. Woburn: Elsevier Science.
- Negahban, A., & Smith, J. S. (2014). Simulation for manufacturing system design and operation: Literature review and analysis. *Journal of Manufacturing Systems*, 33(2), 241–261.
- Negri, E., Fumagalli, L., & Macchi, M. (2017). A review of the roles of digital twin in CPS-based production systems. *Procedia Manufacturing* 11:939–948, 27th International Conference on Flexible Automation and Intelligent Manufacturing, FAIM2017, 27–30 June 2017, Modena, Italy.
- Negri, E., Pandhare, V., Cattaneo, L., Singh, J., Macchi, M., & Lee, J. (2020). Field-synchronized digital twin framework for production scheduling with uncertainty. *Journal of Intelligent Manufacturing* pp 1–22.
- Nicolai, R. P., & Dekker, R. (2008). Optimal maintenance of multi-component systems: A review. In *Complex system maintenance handbook* (pp. 263–286). London: Springer.
- Peng, Y., Dong, M., & Zuo, M. J. (2010). Current status of machine prognostics in condition-based maintenance: A review. *The International Journal of Advanced Manufacturing Technology*, 50(1–4), 297–313.
- Pombo, J., & Ambrósio, J. (2012). Multiple pantograph interaction with catenaries in high-speed trains. *Journal of Computational and Nonlinear Dynamics*, 7(4), 041008.
- Pradhan, S., Samantaray, A. K., & Bhattacharyya, R. (2017). Application of semi-hertzian approach to predict the dynamic behavior of railway vehicles through a wear evolution model. *Journal of Friction and Wear*, 38(6), 437–443.
- Pradhan, S., Samantaray, A. K., & Bhattacharyya, R. (2018). Prediction of railway wheel wear and its influence on the vehicle dynamics in a specific operating sector of Indian railways network. *Wear*, 406–407, 92–104.
- Pradhan, S., Samantaray, A. K., & Bhattacharyya, R. (2019). Multi-step wear evolution simulation method for the prediction of rail wheel wear and vehicle dynamic performance. *Simulation*, 95(5), 441–459.
- Redelinghuys, A., Basson, A., & Kruger, K. (2020). A six-layer architecture for the digital twin: A manufacturing case study implementation. *Journal of Intelligent Manufacturing*, 31, 1383–1402.
- Roberts, M. C. (2009). Duration and characteristics of metal price cycles. *Resources Policy*, 34(3), 87–102.
- Rosen, R., von Wichert, G., Lo, G., & Bettenhausen, K. (2015). About the importance of autonomy and digital twins for the future of manufacturing. *IFAC-PapersOnLine*, 48(3), 567–572.
- Rossen, A. (2015). What are metal prices like? Co-movement, price cycles and long-run trends. *Resources Policy*, 45, 255–276.
- Savolainen, J., Collan, M., & Luukka, P. (2017). Analyzing operational real options in metal mining investments with a system dynamic model. *The Engineering Economist*, 62(1), 54–72.
- Shafto, M., Conroy, M., Doyle, R., Glaessgen, E., Kemp, C., LeMoigne, J., et al. (2010). Modeling, simulation, information technology & processing roadmap. *National Aeronautics and Space*. Administration.
- Shubenkova, K., Valiev, A., Shepelev, V., Tsiulin, S., & Reinau, K. H. (2018). Possibility of digital twins technology for improving efficiency of the branded service system. In: *2018 global smart industry conference (GloSIC)*, Chelyabinsk (pp. 1–7).
- Spiryagin, M., Cole, C., Sun, Y. Q., McClanachan, M., Spiryagin, V., & McSweeney, T. (2014). *Design and simulation of rail vehicles*. Boca Raton: CRC Press, Taylor & Francis Group.
- Szpytko, J., & Duarte, Y. S. (2020). A digital twins concept model for integrated maintenance: a case study for crane operation. *Journal of Intelligent Manufacturing* 1–19.
- Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., & Sui, F. (2018). Digital twin-driven product design, manufacturing and service with big data. *The International Journal of Advanced Manufacturing Technology*, 94(9–12), 3563–3576.
- Topal, E., & Ramazan, S. (2010). A new MIP model for mine equipment scheduling by minimizing maintenance cost. *European Journal of Operational Research*, 207(2), 1065–1071.
- Tuegel, E. J., Ingrassia, A. R., Eason, T. G., & Spottswood, S. M. (2011). Reengineering aircraft structural life prediction using a digital twin. *International Journal of Aerospace Engineering*, 2011, 154798.

- Uhlemann, T. H. J., Lehmann, C., & Steinhilper, R. (2017). The digital twin: Realizing the cyber-physical production system for Industry 4.0. *Procedia CIRP*, 61, 335–340.
- Van Der Auweraer, H., Donders, S., Hartmann, D., & Desmet, W. (2018). Simulation and digital twin for mechatronic product design. In *Proceedings of ISMA 2018—international conference on noise and vibration engineering and USD 2018—international conference on uncertainty in structural dynamics* (pp. 3547–3565).
- Wang, H. (2002). A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, 139(3), 469–489.
- Wei, X., Li, B., Yang, X., & Wang, X. (2018). Fault co-simulation of fuel regulator in a certain type of turbofan engine. *Journal of System Simulation*, 30(10), 3923–3932.
- Werner, A., Zimmermann, N., & Lentjes, J. (2019). Approach for a holistic predictive maintenance strategy by incorporating a digital twin. *Procedia Manufacturing*, 39, 1743–1751.
- West, T. D., & Blackburn, M. (2017). Is digital thread/digital twin affordable? A systemic assessment of the cost of DoD's latest Manhattan project. *Procedia Computer Science*, 114, 47–56.
- Wright, L., & Davidson, S. (2020). How to tell the difference between a model and a digital twin. *Advanced Modeling and Simulation in Engineering Sciences*, 7(1), 1–13.
- Zio, E. (2009). Reliability engineering: Old problems and new challenges. *Reliability Engineering & System Safety*, 94(2), 125–141.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Publication V

Urbani, M., and Collan, M.

Additive manufacturing cases and a vision for a predictive analytics and additive manufacturing based maintenance business model

In: *Technical, Economic and Societal Effects of Manufacturing 4.0: Automation, Adaption and Manufacturing in Finland and Beyond.*

Ed. by Mikael Collan and Karl-Erik Michelsen

Cham: Springer International Publishing, pp. 131-148

Reprinted with permission from

©2020, Springer



Additive Manufacturing Cases and a Vision for a Predictive Analytics and Additive Manufacturing Based Maintenance Business Model

Michele Urbani and Mikael Collan

I INTRODUCTION

In the previous chapter we have seen that the literature on additive manufacturing business models can in broad strokes be divided into four different directions. To illustrate the real-world status quo with examples we discuss in this chapter two real-world cases in the context of using additive manufacturing technology in the production of medical prostheses and in the refurbishment of metal dies and discuss the business model aspects of both of these cases. The third part of the chapter is used to discuss a more

M. Urbani

Department of Industrial Engineering, University of Trento, Trento, Italy

e-mail: michele.urbani@unitn.it

M. Collan (✉)

LUT University, Lappeenranta, Finland

e-mail: mikael.collan@lut.fi

© The Author(s) 2020

M. Collan, K.-E. Michelsen (eds.), *Technical, Economic and Societal Effects of Manufacturing 4.0*,

https://doi.org/10.1007/978-3-030-46103-4_7

visionary case, where additive manufacturing, together with predictive maintenance, allows one to rethink how the modern system of maintenance could work.

2 ADDITIVE MANUFACTURING USED IN ENHANCING HEART SURGERY

The use of additive manufacturing in healthcare applications has flourished in the past two decades [1] and the market share of additive manufacturing in this market was \$6.1 billion in 2016 [2]. The market share of additive manufacturing has caught a steady uptrend, a dramatic increase towards \$21 billion in 2020 is expected.

The improvement of old and the creation of new techniques for 3D printing, together with the development of new purpose-specific materials for the healthcare sector have made possible the deployment of a range of patient-specific applications [3]. For instance, the customization of healthcare products and services such as the realization of customized prosthesis, and the development of case-sized in-vitro models would not have, in many cases, been possible without the developments made in additive manufacturing technologies.

In this chapter we discuss the use of additive manufacturing in the treatment of heart disease from the points of view of the medical procedure involved and the technical solution that additive manufacturing can offer. Cardiovascular diseases (CVD), which are the focus of this real-world example, are a top ranked cause of death on a global level. All in all they were responsible for 17.9 million of deaths in 2015, with almost a 50% increase from 1990 [4]. The added value of additive manufacturing in the process is discussed. The chapter is based on the real-world collaboration between the Department of Industrial Engineering of the University of Trento (Italy) and the Azienda Sanitaria per i Servizi Sanitari (public company in charge of the provision of healthcare services) of the autonomous Province of Trento (Italy).

Atrial Fibrillation: The Condition and the Surgical Intervention

Persons who suffer from a condition known as the non-valvular atrial fibrillation (AF) may be subject to the occurrence of blood clots, which

after being formed within the left atrial appendage (LAA) can enter the blood stream and cause a stroke, or other vascular complications. Many patients are regularly treated with oral anticoagulants, which are aimed at preventing the formation of the blood clots. Unfortunately, this therapy is not always possible, since some individuals have low tolerance of anticoagulants or the risk of bleeding problems caused by the anticoagulants is too high.

An alternative to using anticoagulants is to permanently seal off the LAA—the procedure does not solve the problem with AF, but it helps to prevent blood clot formation through the closure of the LAA.

The surgical intervention in question is called *left atrial appendage occlusion*, also known as the left atrial appendage closure. It is a non-invasive non open-body surgical intervention. There are a number of techniques that can be used to occlude the LAA, one of the most recently introduced practices consists of placing an implant via trans-esophageal echo guidance. The purpose of the implant is to ensure the closure of the LAA and to impede the flow of blood. The intervention is carried out under general anesthesia and is similar to the implantation of a stent.

Since the geometry of the human heart is different for each patient, the *size* and *shape* of the implants to be installed are of fundamental importance. In this context, the decision-maker is the surgeon, who bears the responsibility for the outcome of the surgery. The standard process to treat atrial appendage occlusion begins with a computed tomography (CT) scan of a patient's chest. This allows the doctor to create a rough estimate of the shape and size of the implants that will be implanted. CT is an effective tool and provides a set of cross sectional images of the scanned body along the sagittal plane. The set of images can be processed via a 3D-software and a model of the heart can be created—this allows the patient's LAA to be inspected along the desired direction and gives the surgeon a better basis for decision-making. While the CT scan images and the 3D model are helpful, it remains difficult to foresee the practical difficulties that may arise during the operation.

In the current practice, implants are mass-produced according to standardized shapes and sizes, which forces the surgeon to choose from among a set of possible implant alternatives. With the aid of the CT scan, the doctor can narrow down the set of implants that could fit the heart of a given patient, but there is no full-proof way to in advance identify the right implant alternative.

In practice the fit of the implant is directly tested on the patient during the surgery. Once the right dimensions have been found and the final decision on the implant made, the implant is implanted. The regular procedure foresees that patients can be released after a 24 hour recovery which is typically followed by a 45 day anticoagulant treatment. The success of the intervention cannot be determined immediately after the execution, due to the time required by the human body to adapt to the presence of the implant, in fact the verification takes place during the weeks following the surgery through periodic checks

In case the procedure fails, the operation is typically repeated and the implant is replaced with a better fitting one, thus subjecting the patient to a second intervention. The failure of the process may have severe consequences for the patient, such as pericardial effusion, incomplete LAA closure, dislodgement of the device, and other risks related to catheter-based techniques [5].

Enhancing the Procedure with the Help of Additive Manufacturing Technology

Additive manufacturing can be used to reduce the uncertainty connected to making the selection of the implant used and to assist in planning the surgery ex-ante. Contrary to what the reader might expect, the target of using additive manufacturing in this context is not the creation of a tailor-made implant, but finding a best fitting mass produced implant. While it is logical to expect that once a 3D model of the patient's heart is available an implant of the right size and shape could be additively manufactured, however the low cost of the mass-produced implant and their availability on the market does not make it profitable to print them. Instead a real-size copy of the patient's heart is printed, based on the 3D model obtained from the CT scan. This allows the surgeon to test fit of the (different) mass-produced implants. The process is relatively low cost and provides the advantage of being able to perform pre-surgery testing of fit and by reducing uncertainty increases the chances of a successful operation. The life-size printed model of the heart also allows the surgeon to study the heart and the execution without any pressure, which may prevent practical difficulties during surgery to become overwhelming.

The 3D model produced from the CT images is practically print-ready, no processing by a human operator is needed. A software tool provided by the 3D printer supplier analyses the model and schedules the work for the

printer and adds the print-support structures needed, before the model can be sent to the machine for printing. The printing technology used is stereolithography and it is one of the oldest 3D printing technologies. Patented in the United States in 1986, stereolithography uses a generic photopolymer—specifically a thermosetting resin monomer to build the printed object layer-by-layer on a build platform. Each layer is solidified by a UV laser beam that moves all over the cross-section and is used to solidify the resin. Once a layer has been completed the build platform is lowered and a new layer of material is injected on the cross section—before the new layer is solid, the excess liquid resin is wiped out in order to obtain the right layer-thickness; this also ensures that the surface is even enough for the application of the forthcoming layers. Finally, the printed heart must be finalized by removing the printing supports mechanically (by hand).

The material used (resin) does not have to be bio-compatible as it is not used in contact with the human organs and it can be chosen based on needed mechanical properties of which elasticity is the most important in being able to replicate human tissue-like behaviour during the testing and surgery planning phase. The resin typically used is the softest provided by the supplier and has a Shore-hardness of 50A. Compared to many other photopolymers for 3D printing, the substance is very elastic and it can reach 160% of elongation before breaking. The printing process including the post-treatment of the polymer takes a few hours depending on the complexity of the printed object. The short printing-time together with a good surface finish of the end product, print resolution is up to 100 microns, make stereolithography a suitable technology for this application.

The Business Model Perspective

The application of additive manufacturing described above presents features that are also found in previous literature [6] as hallmarks of successful implementation of additive manufacturing in practice—low production volume, customization of the product, on-demand production, the availability of the (3D) model, and the (modest) cost of the printing equipment.

The cost of a suitable stereolithographic printer ranges between €3000 and €5000, which in the context of the healthcare sector is a rather affordable investment for most medium and large hospitals, if knowledgeable personnel is already in place. The cost of the resin monomer is approximately 200 €/l, which translates to a material cost of some tens of euros

per printed heart. In the context of this case, the region of Trentino—Alto Adige in the North of Italy, the yearly number of operations of the described type is less than one hundred; the region has approximately one million inhabitants.

The low volume of printed hearts produced may nevertheless make it unreasonable for a single hospital to acquire a printer, in which case renting the printing capacity could be a more cost-effective solution. Many research-centres and universities are likely to own a suitable printer and a partnership between hospitals and local research-institutes are a logical way to create win-win partnerships around this theme. More importantly, research-institutes typically have the manpower and expertise to manage the printing process. Starting from a scratch, the time required to acquire the knowledge to manage a polymeric printing process is reasonable—with a few weeks of training, a person is able to manage the whole process.

To summarize, in the present context it does not make sense to talk about a proper business model for this additive manufacturing application, however there is a clear benefit to using additive manufacturing to enhance the treatment of atrial fibrillation. On a scale of multiple hospitals and some thousands of printed hearts on an annual basis there might be a profitable niche for a specialized AM manufacturer. There is always business relevance in being able to provide superior techniques for medical purposes that lower the risks of surgical interventions—this refers to what has previously been characterized as incremental change in adopting additive manufacturing [7].

3 REFURBISHING METAL DIES WITH 3D-PRINTING

In the context of manufacturing, it is very common to have processes that require a physical contact between a manufacturing machine and the processed product. When the contact is made in order to modify the shape of the processed item it is inevitable that the part of the machine that makes the contact, the so called *die*, will be subject to wear. The die can be made of different materials, but here we concentrate on metal dies. Many different surface treatments and improvements in the grade of the base material for dies have been designed in order to limit the effects of wear on metallic dies. While the metals used are hard, on the long run it is inevitable that a die loses nominal geometry, or the surface of the die becomes defective.

The more severe defects are typically surface cracks, sub-surface cracks, and the loss of material from the surface of the die. Generic loss of the

nominal geometry is a type example of less severe defects that occurs with time in the most highly stressed areas of the die. Change of die geometry is a common cause of low quality in the end product. The general picture is that wear causes the end product of the process involving the die to decrease in quality and at a point when the quality-decrease at a limit to acceptable level the die used must be replaced or refurbished.

Due to the above issues, some metallic parts of manufacturing systems must in practice constantly be monitored through the inspection of the quality of the finished items, or through the inspection of the status of the parts themselves. Control performed by way of inspecting finished items is typically based on comparing the produced parts to the design specifications and when the design tolerances are no longer satisfied action must be taken. There are opportunities for predictive maintenance in these cases—minimizing the number of unsatisficing end products is a cost issue.

Maintenance of a system component most often requires stopping the machine, which causes a loss in productivity—this is why the need to repair machines quickly and efficiently is as old as the manufacturing industry itself. Here we concentrate on the maintenance of metal dies used in manufacturing and especially on the refurbishing of metal dies by way of additive manufacturing technologies.

Refurbishing Metal Dies

The current practice of refurbishing metallic dies is based on manual labour. After the defective die has been found via a visual inspection of the production line, it is removed and prepared for maintenance. If the defect is a surface-crack the damaged region of the die will typically consist of an irregular surface on which work cannot be done—machining is first done to remove irregular surfaces, this is done by an operator by hand with a milling machine. In this task the die must be carefully placed in the milling machine and the position of the die and the machine must be calibrated. The end-result is a cavity with a smooth surface.

The cavity is then filled with a suitable filler-metal, typically a manual electro-welding process is used. After filling the cavity the die undergoes re-machining so that the original required (nominal) geometry is re-obtained. This means that a milling machine is again used, after loading the die and calibration of the position on the machine. An error in the positioning of the die in the milling machine will compromise the success of the whole operation. While the manual refurbishing of the die is

relatively inexpensive the risks related to the positioning and calibration of the die in the milling machine remains a problem. The main phases of die-refurbishment are visible in Fig. 1.

The Hybrid Manufacturing Approach to Refurbishing Metal Dies

Thanks to recent developments in the field of manufacturing equipment development, new *hybrid* equipment has become available. A hybrid manufacturing workstation embeds two, or more, manufacturing technologies within it. Typically this means that the elements of *both additive and subtractive manufacturing are present in the same system*. The clear advantage of a hybrid workstation is that as it is able to perform a large number of operations the set-up costs are typically lower. Specifically, only one instance of pre-processing (including calibration) is needed if the hybrid workstation is able to perform an operation, for which multiple machines are otherwise needed—this may dramatically reduce the time consumption as well as the risks related to pre-processing. Hybrid work stations are operated by software designed specifically for these machines. The downside of modern hybrid machines is their relatively high cost.

The hybrid workstation used in refurbishing dies is a DMG Mori Lasertec 65—the workstation integrates laser-deposition melting technology with a 5-axis milling station. The station is able to automatically change between the laser- and the milling-heads. Limitations that the workstation has have to do with the volume and the weight of the worked-on parts (\varnothing 500 mm \times 400 mm; 600 kg)—this kind of limitations are “real” in terms of the workstation not being able to handle larger and heavier objects; as technology is developed further these limitations are slowly relaxed, but the limitations mentioned are on a “good modern



Fig. 1 The three main phases of refurbishing metal dies. (Original photos Matteo Perini)

level". In heavy industry the die component size can still be too large to fit into hybrid workstations for quite a while. The workstation is able to handle various metals and alloys that include stainless steel, nickel-based alloys (Inconel 625, 718), tungsten carbide matrix materials, bronze and brass alloys, chrome-cobalt-molybdenum alloys, stellite, and tool-steel. The CAM/CAD software used is the Siemens NX.

The process of refurbishing dies with the hybrid workstation begins with the setting up and calibration of the damaged die in the workstation and is followed by a 3D-scan and the subsequent construction of a virtual model of the damaged die. A separate software is used for 3D-scanning. The accuracy of the virtual model depends on the resolution of the scanner. The accuracy is a relevant issue, as the more accurate the model is the more accurately it can be decided, which parts of the damaged areas need to be removed—typically the more is removed the more needs to be added later on. If the metal alloy used is very expensive the ability to use less material may have a positive effect on the total cost.

After the decision has been taken, the virtual model is compared to a model of the original (nominal) geometry of the die. With the original model and the virtual model of the damaged die it is possible to obtain the difference between the two and "instruct" the workstation to reconstruct the nominal geometry. The accurate reproduction of the original topology by additively filling the cavities to be repaired is the result of a focused research project at the University of Trento that developed a new method [8] and supporting software that translates the topological difference to a set of machine-understandable instructions that the CAM software is able to read.

The laser deposition melting solution used allows a homogenous distribution of metallic powders, which occurs under the protection of a shield gas that protects the process from oxidation. The system construct is such that a separate work chamber with a controlled environment is not needed—this makes the process faster than is typically the case. It must be observed that the die typically consists of two metallic parts—the part that can be called a "saddle" that is connected to the machine and the "contact part" that is made of a harder metal and that is attached to the saddle and that is the part of the die that is in contact with the produced parts. The actual additive manufacturing procedure is divided into three layers, where the first layer is the (material of the saddle part of the) refurbished die, the second layer is called a dilution zone and it is a mixed material made partly of the original die saddle metal and partly of the filler metal (contact part),

and the third layer is fully made of the filler metal. This three-part procedure is able to produce a very durable *non-porous* and *crack-free* metallic solid—the refurbished die can be said to be “as good as new”, which is the best possible end-result.

The Business Model Perspective

Similar to the heart implant example above, also in the case of die refurbishing the uniqueness of the procedure and the product are key elements—that is, the unique faults in the dies offer a possibility for additive manufacturing to be competitive. Furthermore, as the dies are typically constructed of two metals the reconstruction process of a die is not simple and the ability to refurbish dies to “as good as new” state requires handling high product complexity in an efficient way, something that is possible with the hybrid workstations presented above. This also means that if dies are refurbished en masse that there is always an element of customization to the work—identifying the refurbishing procedure for the various kinds of faults allows something that resembles mass customization. If a relatively expensive hybrid workstation is acquired with a profit in mind it is clear that the workstation should have an as high as possible rate of utilization. This means that there should be a number of different dies (and other parts) for the refurbishing of which the processes should be well-known and ready.

In such a case, even a single hybrid workstation could act as a part of a number of maintenance supply chains and in essence function as a machine-as-a-service (MaaS). The workstation could be purchased through a leasing contract by the customers, who pay an annual fee for the use of the machine, or the machine is acquired by one “player” who then sells the capacity of the machine to others—there are many possible types of ways to organize the availability and the sale of the capacity of the workstation. In the case of the refurbishing metal dies the workstation can, e.g., be a part of a die maintenance chain that consists of predictive maintenance system in place at one or several manufacturing facilities that use(s) metal dies and that is able to refurbish-on-demand. Persona et al. [5] write about maintenance outsourcing and the resulting effects on supply chain organization.

For a manufacturing company the number of dies that need refurbishing on recurring basis must be large enough to warrant the relatively high costs of acquiring and operating a working station, which indicates that

such a move would make sense only for large-enough operations. If a workstation is present on-site any logistics costs are reduced—this may have a marginal positive effect on the cost side, however the potential to shorten downtimes with on-site refurbishing may have a more remarkable positive effect. In a broader perspective the adoption of additive manufacturing technologies must be regarded as a strategic choice for a firm. Purchasing AM capable machinery, such as the hybrid workstation, an organization makes a long-term commitment into a new technology, which not only includes the cost of equipment, but demands the acquisition of the related human talent. We refer the interested reader to Weller et al. [9] for additive manufacturing cases in maintenance applications.

One option in this space is to outsource the maintenance of the metal dies and buy “dies as a service”. There are specialized firms that exclusively sell industrial maintenance capacity and in a sense machine availability—typically in these cases the production facilities belong to customer (manufacturing company) and the service provider is in charge of their good functioning. This option will be discussed more in the remaining part of this chapter, where disruptive maintenance-related business models that rely on digitalization and excellence in additive manufacturing are presented.

4 PREDICTIVE MAINTENANCE AND ADDITIVE MANUFACTURING: JOINT BUSINESS MODEL

Broadly speaking, predictive maintenance is the practice of scheduling and performing maintenance in a way that predicts failures and is hence able to contribute to minimizing production downtimes, maximizing component lifetimes, and to minimizing maintenance costs, we refer the interested reader to see [10]. The indirect benefits that predictive maintenance brings include the potential to use maintenance resources more efficiently, the ability to carry a lower inventory of spare parts, and the important ability to make “tougher” production-related promises to customers. These benefits accrue to both the owner of the maintained system and to the organization responsible for the maintenance that can also be the same organization. The ability to make maintenance more efficient is a source of lasting competitive advantage.

Predictive maintenance is winning ground in manufacturing (and elsewhere) due to the instrumentation of manufacturing equipment that

allows automating the collection of condition data. Based on the data collected predictive models can be tuned in a way that enables the accurate prediction of the timing of equipment failures and the construction of smart maintenance schedules.

Different architectures for predictive systems exist, perhaps the most prevalent at the time of writing are “monitoring-based” systems that track deviations in the system captured by sensors and alert as they appear. Different types of deviations may have different types of “fingerprints” and known the tell-tale signs of a deviation allows the correct classification of the deviation and the correct prediction of an incoming fault. These systems are evolving in the sense that their ability to identify failures becomes better with time as more and more data is accumulated and the patterns that distinguish the different failures become better known. In essence these systems utilize “machine learning”.

Smart means in this context also the ability select a good maintenance policy that keeps the level of unexpected component failures (and stop-pages) at an acceptable level. Smart increasingly means also being able to answer to more difficult questions such as: “once equipment is shut down for maintenance, what else than only the minimum necessary maintenance should one do?”—questions of this type and finding answers to them is difficult and requires system-size modeling for maintenance optimization.

Maintenance optimization work typically includes the modeling of the maintained system and the individual maintainable components (including modeling the wear and tear) and the optimization of the system maintenance based on the model. Bundling maintenance actions in an optimal way is a complex optimization problem and requires considerable computing power and good modeling. So far the typical target of maintenance optimization has been a single system, however, it is clear that the optimization of multiple systems simultaneously offers added benefits. If issues such as workforce scheduling are also taken into consideration in the optimization the complexity of the optimizable problems increases, but on the other hand so do the potential rewards.

One can without a doubt make the claim that the sophistication needed in maintenance optimization is at par with the sophistication needed in the rest of the Manufacturing 4.0 paradigm—someone might even go as far as to say that smart maintenance can be seen as a part of the paradigm when the maintenance context is manufacturing.

*Predictive Maintenance Based Business Model
for Additive Manufacturing*

What makes predictive maintenance different from “typical maintenance” is that due to the instrumentation in place even ad hoc (un-expected) failures can be predicted—in other words there is typically sufficient time to react between acknowledging that a component is about to malfunction and the actual time the component breaks. This period of grace that results from the predictability of faults can be utilized to render the manufacturing operation more efficient by determining the optimal maintenance actions that are performed, when the component that is known to malfunction is changed and by making ready the preparations for the said actions to be performed—including procuring the needed components that need to be exchanged. In this context the procuring the components is the key issue, because the new replacement components can be taken from a (local) storage if they are available, brought to the failing machine from a storage or production location further away, or produced on-site (or near) by additive manufacturing.

Implications of enforcing and making stronger the connection between predictive maintenance and additive manufacturing are quite remarkable—in cases where a failing component, or a spare-part, can be manufactured in time and on-demand for the maintenance action to take place, there may be and there most likely are savings to be made. In the case, where the alternative is transporting a spare-part from far away, which is by far not unheard of. If the on-site (additive) manufacturing of spare parts becomes the trend, the logistics of spare-parts becomes less of an issue and in fact the “logistics middleman” can be even completely cut out. Spare parts logistics are replaced by the logistics of the much less expensive and non time-critical logistics of the materials needed to produce the spare part(s) on location and the digital logistics of the information needed to print the spare part.

One can observe that also the need for the storage of spare parts is diminished as only spares that cannot be printed on demand must be stored—as time passes it can be expected that the selection of materials available for additive manufacturing grows wider and the quantity of non-printable parts grows smaller. Generalizing and perhaps being slightly polemic one may surmise that if there is a revolution by additive manufacturing, then it surely must also be a revolution in logistics. As logistics costs are not insignificant there is a clear potential for savings immediately,

when the cost of production of spare parts by way of additive manufacturing become competitive. We refer the reader interested in the supply chain effects of additive manufacturing to see [11].

The above described predictive maintenance—additive manufacturing symbiosis requires quite seamless informational collaboration between the activities of maintenance (and operation of the maintained equipment), which typically require full knowledge of the design of the said machine, and of the parts production for the machine. In other words, the collaboration of a number of stakeholders in the process is necessary in a way that is very fast, and in the best case automatic.

Automation means that there is a need for a standard system level “rules of play” that govern the informational and trade exchanges taking place within the system, including a joint understanding and pre-acceptance of the involved costs. With the costs we refer, among other things, to the cost of the rights of use of the “recipe” or the digital plans required to print the spare parts, whose IPR typically resides with the original equipment manufacturer.

The fact that a number of things need to be pre-planned and pre-accepted creates a great a “natural” hurdle, when (multiple) separate organizations need to reach consensus—it is therefore likely that the first working systemic solutions that incorporate these technologies in the way envisioned above are formed by actors that already control the different steps of the maintenance and spare-parts production whole and are therefore able to benefit from any and all efficiency increases and cost savings related to process changes.

Blueprint for a Vision

Instrumented equipment is able to digitally transmit real-time information about the condition of perishable parts to what is called “predictive maintenance optimization system” in Fig. 2. The idea is that a sophisticated maintenance analytics system is able to utilize data coming from the sensors located in the production equipment (#1 in Fig. 2), to create results by utilizing modern condition-based maintenance and predictive maintenance models (for ad-hoc failures), and to use the results as input in a sophisticated maintenance action optimization. Modern optimization systems are able to intelligently group maintenance actions to realize potential cost savings from performing multiple maintenance actions simultaneously. Putting smart maintenance planning automatically into

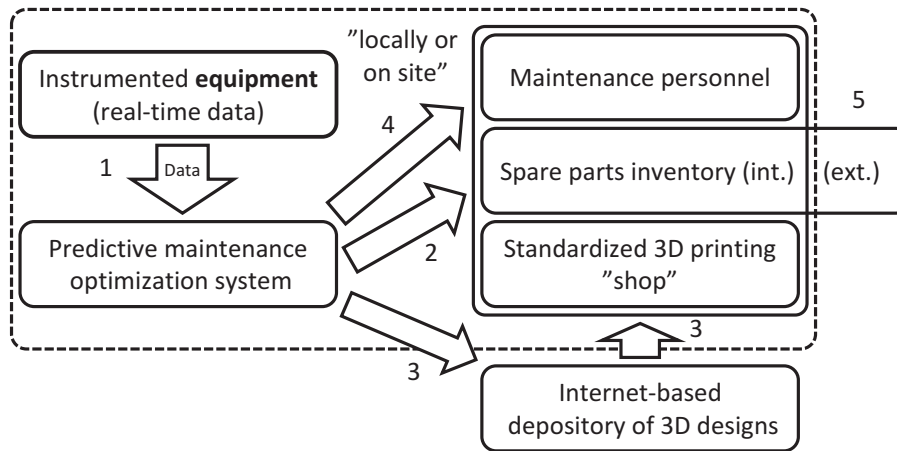


Fig. 2 Additive manufacturing “3D printing shop” as a part of a predictive maintenance oriented maintenance ecosystem

action means that the optimization system has the ability to check the local spare parts inventory for existing parts that are needed (#2 in Fig. 2), to order needed designs from the on-line OEM depository of 3D designs for the parts that need to be printed, and reserving time slots for printing from the 3D printing facility (#3 in Fig. 2).

The maintenance optimization system can also reserve (and in some cases even optimize) the maintenance personnel resources needed and schedule the actual maintenance (#4 in Fig. 2). In contrast to the traditional model, where the locally non-available parts would be searched for and ordered by the automatic system and then shipped to the location from an external warehouse possibly on another continent (#5 in Fig. 2) the additive manufacturing based model can allow for all physical actions to be performed on location. It is clear that a hybrid of the “old and the new” is a state that may be in place for a long time and where the smart optimization system is ultimately able to decide whether to order spare parts from an external warehouse or from a local 3D-printing facility based on minimizing a cost function that may include, e.g., time penalties.

In an ideal world the optimization system is able to create a circumstance, where the costs are minimized, optimal amount of maintenance is carried out, parts are ready just-in-time, and personnel resources are optimized. The driving forces behind reaching this kind of a state are the development of digital instrumentation in equipment (IoT), development of smart predictive maintenance systems that are coupled with advanced

maintenance optimization systems that are digitally connected to resource management systems. The vision presented includes an Internet-based depository of 3D designs as a component—such depositories already exist for hobbyist designs, as of yet serious B2B depositories have not emerged.

Many different kinds of business model possibilities exist within the vision, the envisioned whole can be realized within the “realm” of a single actor, or by way of collaboration of specialized single actors.

5 CONCLUSIONS

This chapter has concentrated on presenting two real-world cases of how additive manufacturing can be used to enhance existing processes that otherwise demand precision manual labor and/or cannot be performed as well. Both of the real-world cases show that there is potential in additive manufacturing in places, where sophisticated tailoring of what is done is required, and where precision is a key factor. In both the cases the business model aspects had not been fully explored due to the exploratory and piloting nature of the activities performed, but it remains quite clear that with a high-enough demand for the presented activities there is a profitable business case to be made. If a specialized know-how is created around an additive manufacturing resource, the resource can be leveraged to service multiple different clients. It must be observed that in the same way as with any production technology, if the utilization rate of the equipment used is low the chances of reaching profitability remain a challenge—the laws of production economics do not change.

There seems to be a place for visioning additive manufacturing based business models that combine additive manufacturing with other technologies, such as predictive maintenance, as presented in this chapter. The ability of additive manufacturing to deliver on-demand is an important factor from the point of view of efficiency gains it is able to bring to the business of which additive manufacturing is a part of. When coupled with “control” technology that is able to make just-in-time orders and to optimize processes the ability to produce just-in-time can be exploited effectively. The prospect of locally manufacturing with additive manufacturing technologies through a global web of digital information is an interesting one and puts pressure on mass-production and long-haul logistics based business models.

Industry-grade 3D-printers can be thought of as platform investments that service more than one client and that draw from a world-wide resource

of 3D-printing designs. At this time serious commercial business to business depositories of 3D-printing designs do not exist and the business model is still in its infancy. Many issues remain to be solved in the (digital) collaboration between the original equipment manufacturers to whom the 3D-printing designs belong to, the secure distribution and pricing of designs, and the (trusted) network of 3D-printing resources that can service clients globally.

REFERENCES

1. S. H. Huang, et al., "Additive manufacturing and its societal impact: a literature review," *International journal of advanced manufacturing technology*, vol. 67, pp. 1191–1203, 2013.
2. I. Campbell, et al., "Wohlers Report 2017 – Additive Manufacturing and 3D Printing State of the Industry Annual Worldwide Progress Report," Wohlers Associates Inc. , Fort Collins, Colorado 2017.
3. A. A. Zadpoor and J. Malda, "Additive Manufacturing of Biomaterials, Tissues, and Organs," *Annals of biomedical engineering*, vol. 45, pp. 1–11, 2017.
4. H. e. a. Wang, "Global, regional, and national life expectancy, all-cause mortality, and cause-specific mortality for 249 causes of death, 1980–2015: a systematic analysis for the Global Burden of Disease Study 2015," *Lancet*, vol. 388, pp. 1459–1544, 2016.
5. P. B. Sick, et al., "Initial Worldwide Experience With the WATCHMAN Left Atrial Appendage System for Stroke Prevention in Atrial Fibrillation," *Journal of American college of cardiology*, vol. 49, pp. 1490–1495, 2007.
6. A. Persona, et al., "Remote control and maintenance outsourcing networks and its applications in supply chain management," *Journal of operations management*, vol. 25, pp. 1275–1291, 2007.
7. J. Savolainen and M. Collan, "Additive manufacturing technology and business model change – a review of literature," *Additive manufacturing*, Vol. 32, Art. 101070, 2020.
8. M. Perini, et al., "Additive manufacturing for repairing: From damage identification and modeling to DLD," *Rapid Prototyping Journal*, 2020 (ahead of print).
9. C. Weller, et al., "Economic implications of 3D printing: Market structure models in light of additive manufacturing revisited," *International Journal of Production Economics*, vol. 164, pp. 43–56, 2015.
10. R. K. Mobley, *An introduction to predictive maintenance*, 2nd ed. Oxford, UK: Butterworth-Heinemann, 2002.
11. D. Thomas, "Costs, benefits, and adoption of additive manufacturing: a supply chain perspective," *International journal of advanced manufacturing technology*, vol. 85, pp. 1857–1876, 2016.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



ACTA UNIVERSITATIS LAPPEENRANTAENSIS

- 956.** NGUYEN, HOANG SI HUY. Model based design of reactor-separator processes for the production of oligosaccharides with a controlled degree of polymerization. 2021. Diss.
- 957.** IMMONEN, HEIKKI. Application of object-process methodology in the study of entrepreneurship programs in higher education. 2021. Diss.
- 958.** KÄRKKÄINEN, HANNU. Analysis of theory and methodology used in determination of electric motor drive system losses and efficiency. 2021. Diss.
- 959.** KIM, HEESOO. Effects of unbalanced magnetic pull on rotordynamics of electric machines. 2021. Diss.
- 960.** MALYSHEVA, JULIA. Faster than real-time simulation of fluid power-driven mechatronic machines. 2021. Diss.
- 961.** SIEVINEN, HANNA. Role of the board of directors in the strategic renewal of later-generation family firms. 2021. Diss.
- 962.** MENDOZA MARTINEZ, CLARA. Assessment of agro-forest and industrial residues potential as an alternative energy source. 2021. Diss.
- 963.** OYEWO, AYOBAMI SOLOMON. Transition towards decarbonised power systems for sub-Saharan Africa by 2050. 2021. Diss.
- 964.** LAHIKAINEN, KATJA. The emergence of a university-based entrepreneurship ecosystem. 2021. Diss.
- 965.** ZHANG, TAO. Intelligent algorithms of a redundant robot system in a future fusion reactor. 2021. Diss.
- 966.** YANCHUKOVICH, ALEXEI. Screening the critical locations of a fatigue-loaded welded structure using the energy-based approach. 2021. Diss.
- 967.** PETROW, HENRI. Simulation and characterization of a front-end ASIC for gaseous muon detectors. 2021. Diss.
- 968.** DONOGHUE, ILKKA. The role of Smart Connected Product-Service Systems in creating sustainable business ecosystems. 2021. Diss.
- 969.** PIKKARAINEN, ARI. Development of learning methodology of additive manufacturing for mechanical engineering students in higher education. 2021. Diss.
- 970.** HOFFER GARCÉS, ALVARO ERNESTO. Submersible permanent-magnet synchronous machine with a stainless core and unequal teeth widths. 2021. Diss.
- 971.** PENTTILÄ, SAKARI. Utilizing an artificial neural network to feedback-control gas metal arc welding process parameters. 2021. Diss.
- 972.** KESSE, MARTIN APPIAH. Artificial intelligence : a modern approach to increasing productivity and improving weld quality in TIG welding. 2021. Diss.
- 973.** MUSONA, JACKSON. Sustainable entrepreneurial processes in bottom-of-the-pyramid settings. 2021. Diss.
- 974.** NYAMEKYE, PATRICIA. Life cycle cost-driven design for additive manufacturing: the frontier to sustainable manufacturing in laser-based powder bed fusion. 2021. Diss.

975. SALWIN, MARIUSZ. Design of Product-Service Systems in printing industry. 2021. Diss.
976. YU, XINXIN. Contact modelling in multibody applications. 2021. Diss.
977. EL WALI, MOHAMMAD. Sustainability of phosphorus supply chain – circular economy approach. 2021. Diss.
978. PEÑALBA-AGUIRREZABALAGA, CARMELA. Marketing-specific intellectual capital: Conceptualisation, measurement and performance. 2021. Diss.
979. TOTH, ILONA. Thriving in modern knowledge work: Personal resources and challenging job demands as drivers for engagement at work. 2021. Diss.
980. UZHEGOVA, MARIA. Responsible business practices in internationalized SMEs. 2021. Diss.
981. JAISWAL, SURAJ. Coupling multibody dynamics and hydraulic actuators for indirect Kalman filtering and real-time simulation. 2021. Diss.
982. CLAUDELIN, ANNA. Climate change mitigation potential of Finnish households through consumption changes. 2021. Diss.
983. BOZORGMEHRI, BABAK. Finite element formulations for nonlinear beam problems based on the absolute nodal coordinate formulation. 2021. Diss.
984. BOGDANOV, DMITRII. Transition towards optimal renewable energy systems for sustainable development. 2021. Diss.
985. SALTAN, ANDREY. Revealing the state of software-as-a-service pricing. 2021. Diss.
986. FÖHR, JARNO. Raw material supply and its influence on profitability and life-cycle assessment of torrefied pellet production in Finland – Experiences from pilot-scale production. 2021. Diss.
987. MORTAZAVI, SINA. Mechanisms for fostering inclusive innovation at the base of the pyramid for community empowerment - Empirical evidence from the public and private sector. 2021. Diss.
988. CAMPOSANO, JOSÉ CARLOS. Integrating information systems across organizations in the construction industry. 2021. Diss.
989. LAUKALA, TEIJA. Controlling particle morphology in the in-situ formation of precipitated calcium carbonate-fiber composites. 2021. Diss.
990. SILLMAN, JANI. Decoupling protein production from agricultural land use. 2021. Diss.
991. KHADIM, QASIM. Multibody system dynamics driven product processes. 2021. Diss.
992. ABDULKAREEM, MARIAM. Environmental sustainability of geopolymer composites. 2021. Diss.
993. FAROQUE, ANISUR. Prior experience, entrepreneurial outcomes and decision making in internationalization. 2021. Diss.



ISBN 978-952-335-742-6
ISBN 978-952-335-743-3 (PDF)
ISSN-L 1456-4491
ISSN 1456-4491
Lappeenranta 2021