UNIVERSITÀ DI TRENTO

Department of Industrial Engineering

XXXIV cycle

Doctoral School in Materials, Mechatronics and System Engineering

Maintenance policies optimization in the Industry 4.0 paradigm

Michele Urbani

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.

December 2021

Maintenance policies optimization in the Industry 4.0 paradigm

Michele Urbani

Email: michele.urbani@unitn.it

Approved by:

Prof. Matteo Brunelli, Advisor
Dept. of Industrial Engineering
University of Trento, Italy

Prof. Mikael Collan
LUT School of Business and Management
Lappeenranta-Lahti University of Technology LUT, Finland

Ph.D. Commission:

Prof. Francesco Pilati
Dept. of Industrial Engineering
University of Trento, Italy

Prof. Debora Di Caprio
Department of Economics and Management
University of Trento, Italy

Prof. Kari Koskinen
Faculty of Engineering and Natural Sciences
Tampere University, Finland

Prof. Yuri Lawryshyn
Faculty of Applied Science and Engineering
University of Toronto, Canada

University of Trento

Department of Industrial Engineering

December 2021

University of Trento - Department of Industrial Engineering

Doctoral Thesis

Michele Urbani- December 2021

Published in Trento (Italy) - by University of Trento

"Success is walking from failure to failure with no loss of enthusiasm."

Winston Churchill

Abstract

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.

Acknowledgements

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 here but shared part of this journey.

Michele Urbani
November 22, 2021
Trento, Italy

Contents

List of Abbreviations

CBM	**C**ondition **B**ased **M**aintenance
CM	**C**orrective **M**aintenance
CPS	**C**yber **P**hysical **S**ystems
CPPS	**C**yber **P**hysical **P**roduction **S**ystems
DT	**D**igital **T**win
GA	**G**enetic **A**lgorithm
IoT	**I**nternet **o**f **T**hings
LCC	**L**ife **C**ycle **C**ost
MOO	**M**ulti **O**bjective **O**ptimization
NED	**N**egative **E**conomic **D**ependencies
PED	**P**ositive **E**conomic **D**ependencies
PM	**P**reventive **M**aintenance
RAMS	**R**eliability **A**vailability **M**aintainability **S**afety
PHM	**P**rognostics **H**ealth **M**anagement

List of publications

This dissertation is based on the following papers and manuscripts.

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 th-rough 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 MikaelCollan 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 Boset-ti 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.

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 wellbeing. 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.

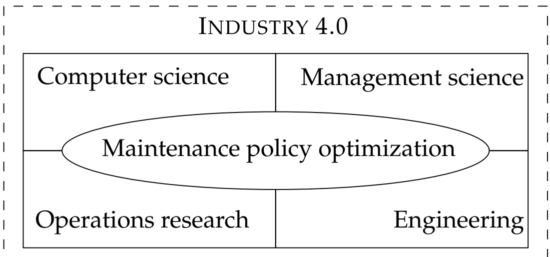

FIGURE 1.1: The map representing the fields covered in this thesis.

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 $\Pr\left(S(t) > L(t)\right) > \rho$, 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.

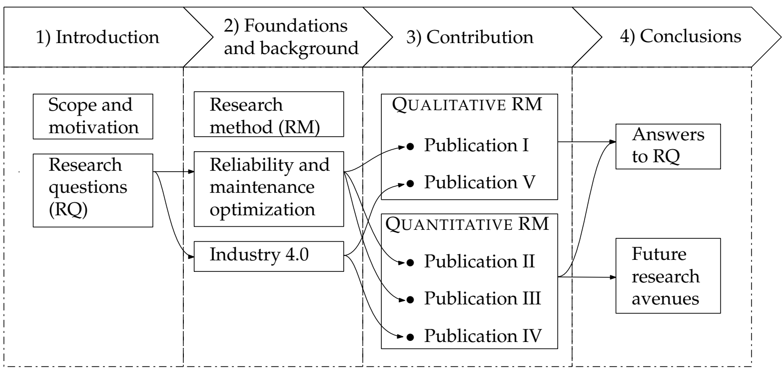

FIGURE 1.2: Contents of this thesis.

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 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".

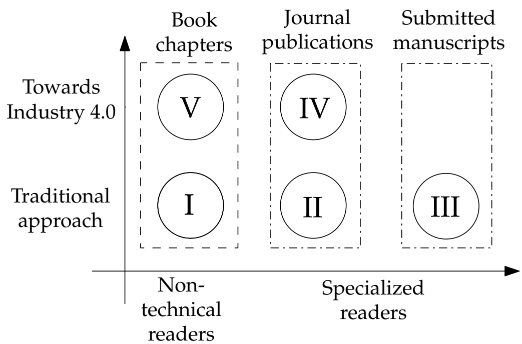

FIGURE 1.3: The research outcomes are represented within a conceptual map, and grouped according to the development approach. I) Publication I, II) Publication II, III) Publication III, IV) Publication IV, V) Publication V.

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.

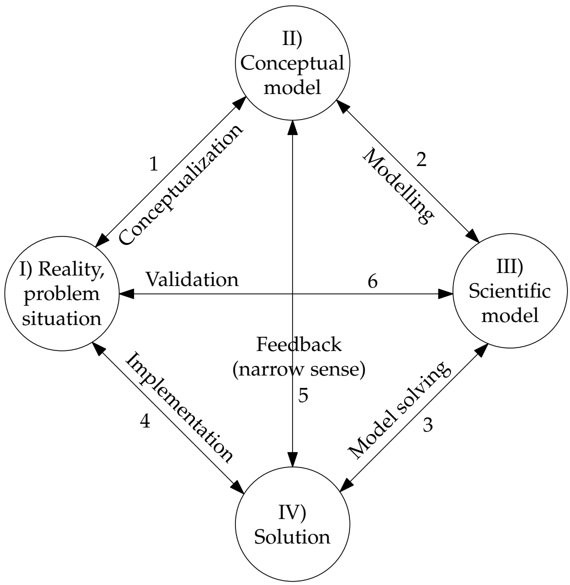

FIGURE 2.1: The systemic approach to problem-solving proposed by Mitroff et al. (1974).

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.

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-buil-ders 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 cooperation (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 subsystem, 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 *maintinir* 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 microproces-sor- and computer-based instrumentation for monitoring of machines allowed the development of the so-called *condition-based* ma-intenance (CBM), which aims at reducing (or even eliminating) unnecessary interventions by doing maintenance on-demand. Since the 2000s, CBM was further developed into prognostics

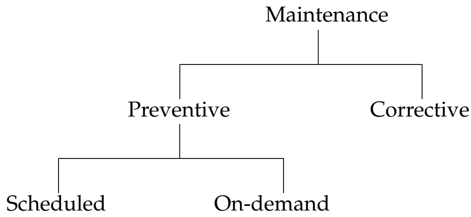

FIGURE 2.2: A schematic representation of maintenance approaches.

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 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, 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,

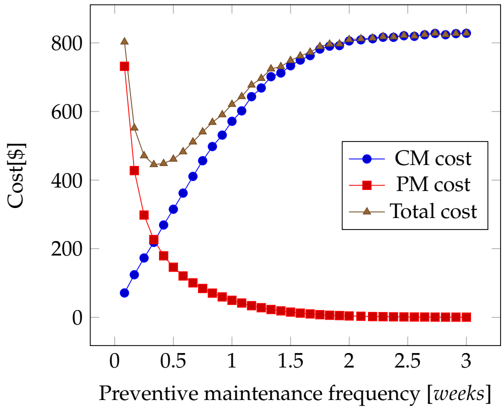

FIGURE 2.3: The cost of CM, PM, and the total maintenance cost as a function of the maintenance frequency. Reprinted from Publication IV.

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 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

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

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

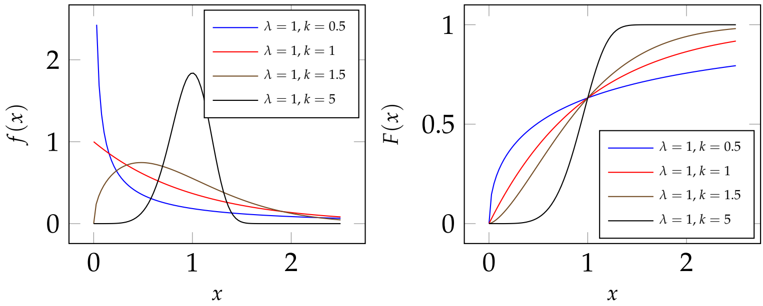

FIGURE 2.4: A few examples of Weibull density and cumulative functions. The case $\lambda = 1$ and $k = 1$ is equivalent to the exponential model.

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-compo-nent 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. Grou-ping 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 achie-ved 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 de-cision-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 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

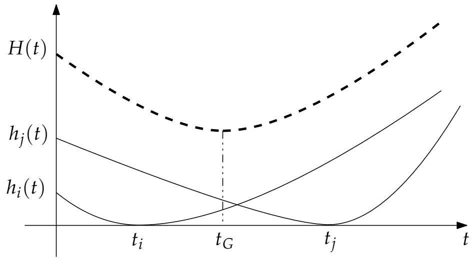

FIGURE 2.5: The incremental expected cost functions for single activities $h_i(t)$ and $h_j(t)$ as a function of the shift time t and the group expected cost of maintenance $H(t)$, which is the sum of $h_i(t)$ and $h_j(t)$ and presents a minimum in t_G.

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 favour 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 simu-lation-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

$$\begin{array}{ll} \text{minimize} & \big(f_1(\mathbf{x}), \ldots, f_p(\mathbf{x})\big) \\ \text{subject to} & \mathbf{x} \in \mathcal{X} \end{array}$$

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 $f_1(\mathbf{x}), \ldots, f_p(\mathbf{x})$ are the $p \geq 2$ *objective functions* of the problem, also denoted by $\boldsymbol{f}(\mathbf{x}) = \big(f_1(\mathbf{x}), \ldots, f_p(\mathbf{x})\big)^T$, and each $f_i : \mathbb{R}^n \to \mathbb{R}$. 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 $i = 1, \ldots, p$ 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 $\mathbf{z} = \boldsymbol{f}(\mathbf{x})$ 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 $i = 1, \ldots, p$ 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-convex, 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 $\mathbf{x}$ called *individuals*, which can be evaluated using a *fitness* function. The latter can be any kind of function that can be fed with $\mathbf{x}$ and returns a scalar value z in the case of a single-objective GA, or a vector $\mathbf{z}$ 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.

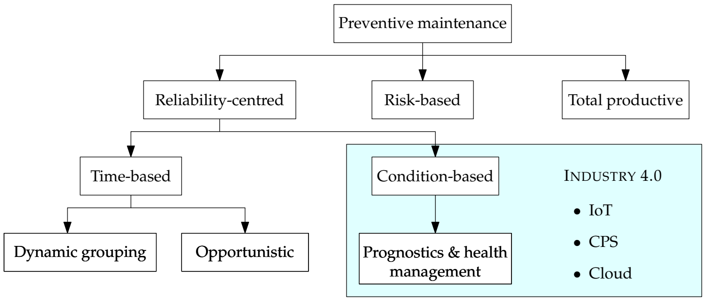

FIGURE 2.6: A schematic representation of preventive maintenance approaches.

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 th-reats. 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 Balc, 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 employement 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 evol-ved 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

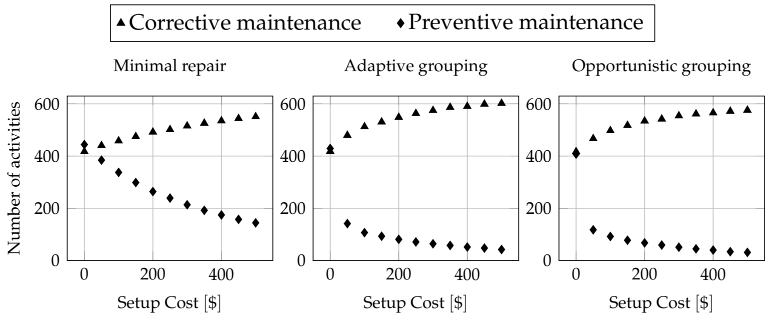

FIGURE 3.1: The number of maintenance interventions caused by PM and CM per each policy. Reprinted from Publication II.

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 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 criteria, 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 autono-mously 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 manu-script.

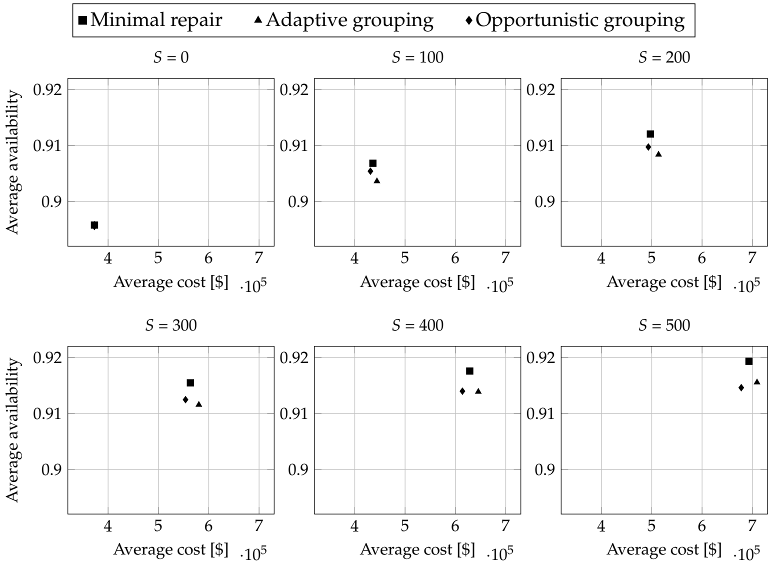

FIGURE 3.2: A bi-criteria comparison of the analysed policies at variable setup costs S. Reprinted from Publication II.

3.3 Publication III (under review)

Maintenance of complex multi-unit systems is often a multi-objec-tive optimization problem. In the manuscript *"An approach for bi-objective maintenance scheduling on anetworked 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.

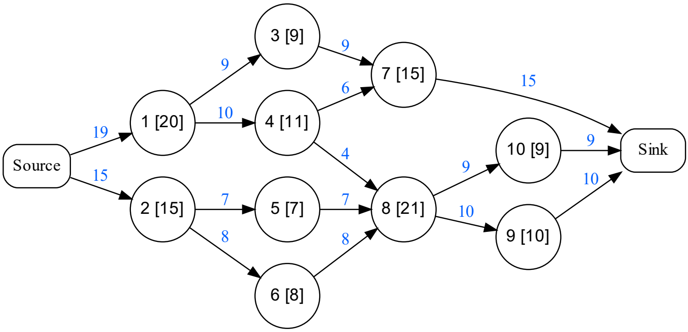

FIGURE 3.3: A graph representing the system analyzed in the paper. The circles represent physical assets and arcs represent the connections between assets. The numbers above the arcs (in blue) describe how the maximum flow of 34 is distributed to the nodes, which are represented by their IDs and capacities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

The multi-unit system of non-identical components was modelled using a directed graph, which described the connections a-mong 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

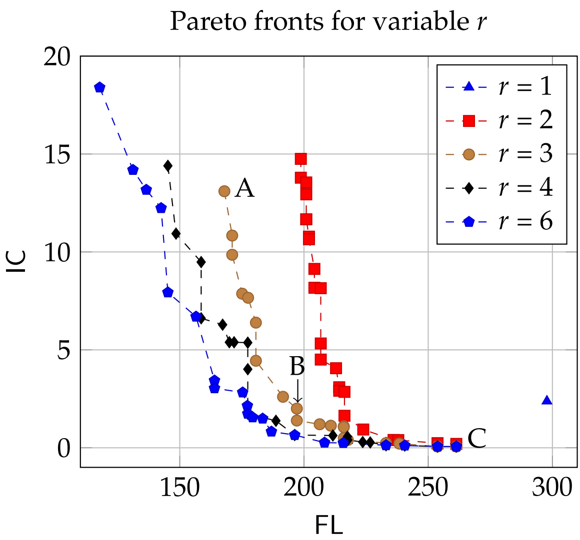

FIGURE 3.4: The frontiers of non-dominated solutions obtained using a variable number of resources. Reprinted from Publication III.

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 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

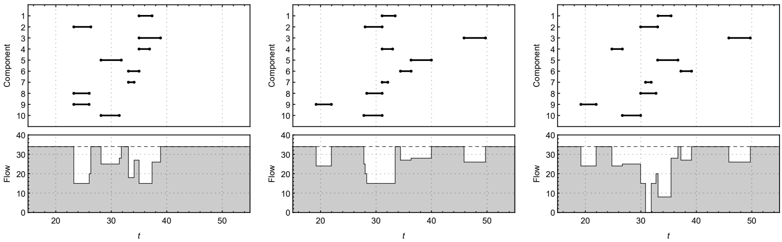

(A) Solution A with IC = 13.09 and FL = 167.95.
(B) Solution B with IC = 1.50 and FL = 199.03.
(C) Solution C with IC = 0.24 and FL = 248.29.

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.

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 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-econo-mic 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

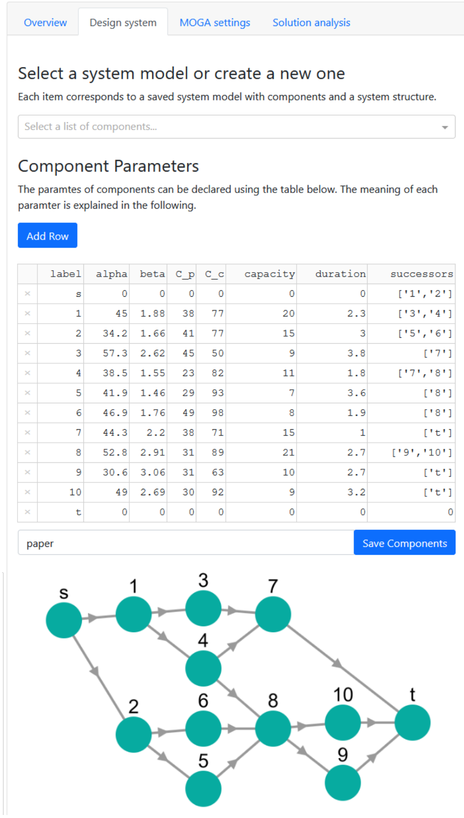

FIGURE 3.6: A screenshot of the "System design" page of the web app realized during the development of Publication III.

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 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-economical 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

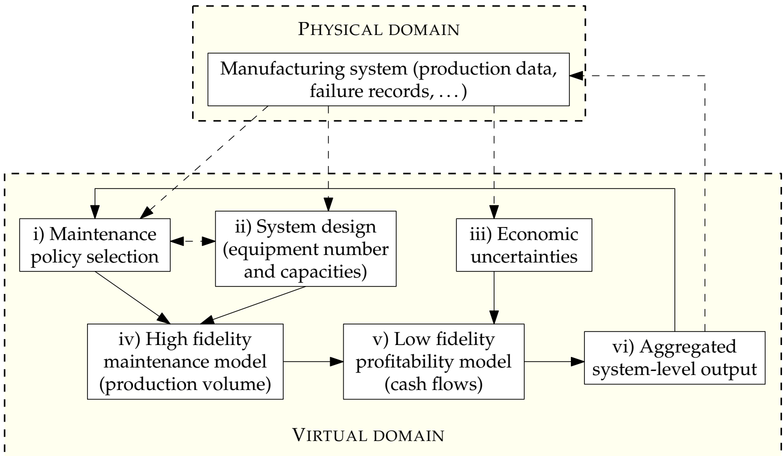

FIGURE 3.7: A schematic representation of the flow of information within the simulated digital twin model. Reprinted from Publication IV.

CPS, DT connects 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 healthcare 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 me-thod 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 abil-

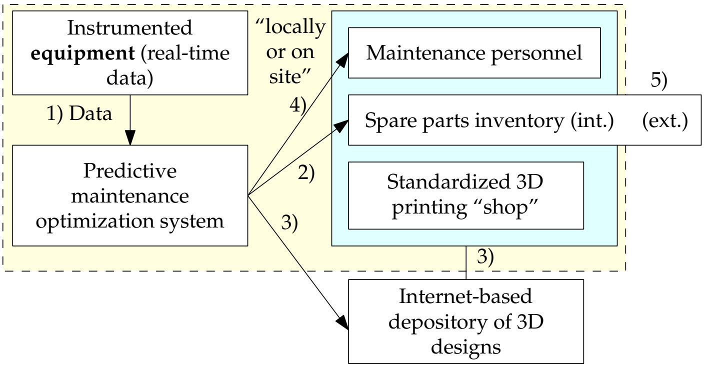

FIGURE 3.8: A blueprint model for PM-oriented maintenance ecosystem. Reprinted from Publication IV.

ity 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 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 perfomance

Continued on next page

Table 3.1 – *Continued from previous page*

Publication	Objective	Research method(s)	Data	Contribution
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.
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 mo-del 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 sche-dules. 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-domi-nated 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 quickly 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. (2021), 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., 2021), 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.

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.

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.

Aven, Terje (2012). *Foundations of Risk Analysis*. John Wiley & Sons.

Baraldi, Piero et al. (2011). “A randomized model ensemble approach for reconstructing signals from faulty sensors”. In: *Expert Systems with Applications* 38.8, pp. 9211–9224.

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.

Barlow, Euan et al. (2021). “A performance-centred approach to optimising maintenance of complex systems”. In: *European Journal of Operational Research* 292.2, pp. 579–595.

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.

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.

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.

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.

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.cdbb.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.

Cimino, Chiara, Elisa Negri, and Luca Fumagalli (2019). "Review of digital twin applications in manufacturing". In: *Computers in Industry* 113, p. 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.

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.

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.

De Jonge, B. and P. A. Scarf (2020). "A review on maintenance optimization". In: *European Journal of Operational Research* 285.3, pp. 805 –824.

Deb, Kalyanmoy et al. (2002). "A fast and elitist multiobjective genetic algorithm: NSGA-II". In: *IEEE Transactions on Evolutionary Computation* 6.2, pp. 182–197.

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.

Derler, Patricia, Edward A. Lee, and Alberto Sangiovanni Vincentelli (2012). "Modeling Cyber–Physical Systems". In: *Proceedings of the IEEE* 100.1, pp. 13–28.

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.

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.

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.

Do Van, Phuc et al. (2013). "Dynamic grouping maintenance with time limited opportunities". In: *Reliability Engineering & System Safety* 120, pp. 51–59.

Ersdal, Gerhard and Terje Aven (2008). "Risk informed decision-making and its ethical basis". In: *Reliability Engineering & System Safety* 93.2, pp. 197–205.

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.

Fonseca, Carlos 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.

Fu, Michael C. (2002). "Optimization for simulation: Theory vs. practice". In: *INFORMS Journal on Computing* 14.3, pp. 192–215.

Gallo, Giorgio (2004). "Operations research and ethics: Responsibility, sharing and cooperation". In: *European Journal of Operational Research* 153.2. Management of the Future MCDA: Dynamic and Ethical Contributions, pp. 468–476.

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.

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.

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.

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 Vocaboulary*. International Standard Organization, Geneva.

Kim, Nam-Ho, Dawn An, and Joo-Ho Choi (2017). *Prognostics and Heal-th 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.

Konak, Abdullah, David W. Coit, and Alice E. Smith (2006). "Multiobjective optimization using genetic algorithms: A tutorial". In: *Reliability Engineering & System Safety* 91.9, pp. 992–1007.

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.

Kuhn, Thomas S. (2012). *The Structure of Scientific Revolutions*. University of Chicago press.

Law, Averill M. (2014). *Simulation Modeling and Analysis. Vol. 5th International Edition*. New York, New York: McGraw-Hill, Inc.

Lee, Edward 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.

Liang, Zhenglin and Ajith Kumar Parlikad (2020). "Predictive group maintenance for multi-system multi-component networks". In: *Reliability Engineering & System Safety* 195, p. 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.

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.

Miettinen, Kaisa (2012). *Nonlinear Multiobjective Optimization*. Vol. 12. Springer Science & Business Media.

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.

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 phenome-non". In: *Interfaces* 4.3, pp. 46–58.

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.

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.

— (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.

Nicolai, Robin P and Rommert Dekker (2008). "Optimal maintenance of multi-component systems: a review". In: *Complex system maintenance handbook*, pp. 263–286.

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.

Papadakis, Ioannis S and Paul R Kleindorfer (2005). "Optimizing infrastructure network maintenance when benefits are interdependent". In: *OR Spectrum* 27.1, pp. 63–84.

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.

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.

Perini, Matteo, Paolo Bosetti, and Nicolae Balc (2020). "Additive manufacturing for repairing: From damage identification and modeling to DLD". In: *Rapid Prototyping Journal* 26.5, pp. 929–940.

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.

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, Fran cesco Di Maio, and Enrico Zio.

Rausand, Marvin and Arnljot Høyland (2003). *System Reliability Theory: Models, Statistical Methods, and Applications*. Vol. 396. John Wiley & Sons.

Rausand, Marvin and Jørn Vatn (2008). "Reliability centred maintenance". In: *Complex System Maintenance Handbook*. Springer, pp. 79–108.

Rocchetta, R. et al. (2019). "A reinforcement learning framework for optimal operation and maintenance of power grids". In: *Applied Energy* 241, pp. 291 –301.

Saaksvuori, Antti and Anselmi Immonen (2008). *Product Lifecycle Management*. Springer Science & Business Media.

Savolainen, Jyrki and Mikael Collan (2020). "How Additive Manufacturing Technology Changes Business Models? – Review of Literature". In: *Additive Manufacturing* 32, p. 101070.

Shafto, Mike et al. (2010). *Modeling, Simulation, Information Technology & Processing Roadmap*. (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.

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.

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.

Stricker, Nicole et al. (2018). "Reinforcement learning for adaptive order dispatching in the semiconductor industry". In: *CIRP Annals* 67.1, pp. 511 –514.

Sun, Bo et al. (2012). "Benefits and Challenges of System Prognostics". In: *IEEE Transactions on Reliability* 61.2, pp. 323–335.

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.

Tuegel, Eric J et al. (2011). "Reengineering aircraft structural life prediction using a digital twin". In: *International Journal of Aerospace Engineering* 2011.

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.

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.

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.

— (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.

Wang, Hongzhou (2002). "A survey of maintenance policies of deteriorating systems". In: *European Journal of Operational Research* 139.3, pp. 469–489.

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.

Wenstøp, Fred (2010). "Operations research and ethics: development trends 1966–2009". In: *International Transactions in Operational Research* 17.4, pp. 413–426.

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.

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.

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.

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.

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.

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.

Zio, Enrico (2007). *An Introduction to the Basics of Reliability and Risk Analysis*. Vol. 13. World scientific.

— (2009). “Reliability engineering: Old problems and new challenges”. In: *Reliability Engineering & System Safety* 94.2, pp. 125 –141.