Designing Manufacturing Systems for Distributed Control

by

Dipl.-Logist. Henning Blunck, MSIE (USA)

A Thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in International Logistics

Approved Dissertation Committee:
1. Supervisor: Prof. Dr.-oec. Julia C. Bendul (Jacobs University Bremen)
2. Supervisor: Prof. Dr. Marc-Thorsten Hütt (Jacobs University Bremen)
3. Supervisor: Prof. Dr. Dieter Armbruster (Arizona State University)

Date of defense: June 6th, 2017

Department of Mathematics & Logistics

Statutory Declaration

Family Name, Given/First Name	Blunck, Henning
Matriculation number	20330788
Type of thesis	PhD-Thesis

English: Declaration of Authorship

I hereby declare that the thesis submitted was created and written solely by myself without any external support. Any sources, direct or indirect, are marked as such. I am aware of the fact that the contents of the thesis in digital form may be revised with regard to usage of unauthorized aid as well as whether the whole or parts of it may be identified as plagiarism. I do agree my work to be entered into a database for it to be compared with existing sources, where it will remain in order to enable further comparisons with future theses. This does not grant any rights of reproduction and usage, however.

This document was neither presented to any other examination board nor has it been published.

German: Erklärung der Autorenschaft (Urheberschaft)

Ich erkläre hiermit, dass die vorliegende Arbeit ohne fremde Hilfe ausschließlich von mir erstellt und geschrieben worden ist. Jedwede verwendeten Quellen, direkter oder indirekter Art, sind als solche kenntlich gemacht worden. Mir ist die Tatsache bewusst, dass der Inhalt der Thesis in digitaler Form geprüft werden kann im Hinblick darauf, ob es sich ganz oder in Teilen um ein Plagiat handelt. Ich bin damit einverstanden, dass meine Arbeit in einer Datenbank eingegeben werden kann, um mit bereits bestehenden Quellen verglichen zu werden und dort auch verbleibt, um mit zukünftigen Arbeiten verglichen werden zu können. Dies berechtigt jedoch nicht zur Verwendung oder Vervielfältigung.

Diese Arbeit wurde noch keiner anderen Prüfungsbehörde vorgelegt noch wurde sie bisher veröffentlicht.

Location, Date

Henning Blunck

Acknowledgements

"The more human beings proceed by plan the more effectively they may be hit by accident."

Friedrich Dürrenmatt,
The Physicists

I shall consider myself a fan of making plans.

The work presented in this thesis by contrast, leads us to the edge of planning. It considers systems too complex to be effectively anticipated and planned by a single mind or even computer. It investigates situations, where production system designers feel compelled to trust the emergence of behavior, rather than prescribing it. It is a journey, well outside the comfort zone for aficionados of "the one best way".

It has been a similar journey for me.

Coming to the end of the journey, I can only express my gratitude and appreciation for those who have accompanied me, pointed to promising alleys, warned me about potential roadblocks ahead, or have cheered me on along the way. First, I would like to thank Prof. Bendul for her support in the research documented here and many other aspects of my time as a research associate in the Workgroup for Production & Logistics Networks at Jacobs University Bremen. I remain grateful to my second and third supervisors, Prof. Hütt and Prof. Armbruster, for making the research presented here as transdisciplinary and multi-faceted as it hopefully is. I hope to have been minted by their curiosity and cheerful approach to the scientific endeavor. I thank my colleagues in the workgroup, past and present, for their help, inspiration, (support of) geekiness, and outright denial of frustration and despondence even when the chips were down.

Finally, this work, like many others, would not have been possible without the enthusiastic work that is done all over the world by developers of free software and those who share their knowledge through blogs and online communities. It would not be without them that scientists today can set out to perform, analyze, and communicate experiments as those reported here. Thank you!

Abstract

The distribution of control capabilities and functions among autonomous system components has attracted extensive research in the fields of logistics and production planning & control (PPC). It relies on deep-rooted concepts of emergence and self-organization that are being increasingly applied as a system of thought to describe the genesis of order in networks of interacting systems (be them biological, social, or engineered). Their emergent nature, however, renders much of the traditional, reductionist knowledge about the design of manufacturing systems and their control void, opening a gap in the understanding that is already threatening the industrial adoption of distributed PPC approaches. The current thesis addresses this particular research gap. It is driven in especially by the frequently expressed hypothesis that a combination of classical, centralized production control and new, distributed forms can yield optimal performance. This hypothesis is explored through a combination of interdisciplinary literature review and minimal model investigations. In a first step, the literature is analyzed and categorized to form a classification model of decisions pertaining to the design of manufacturing systems and their control. This describes a parameter space within which a combination of centralized and distributed control can be attained. The thesis then discusses two minimal models that investigate in more detail particular design decisions. First, Chapter 4 applies Cellular Automata on networks of different structure to investigate the role of control network hierarchy on the performance of agents in simple, distributed problem solving settings, finding not only a performance peak at "medium" levels of hierarchy, but also developing a mechanistic understanding for it. The second quantitative model in Chapter 5 borrows from findings in algorithmic game theory to explore how the emergent behavior of selfish agents can be reconciled with the established ideal in manufacturing system design to set target utilization levels for machines. The findings of this thesis support a design approach for distributed Production Planning & Control (PPC) systems based on evidence and analysis, instead of experience and experimentation. It enhances our understanding of the success factors of distributed control in production environments and beyond. It can advance the development of "emergence engineering" by providing a deeper understanding of the target-driven design of Complex Adaptive System (CAS).

Contents

List of Abbreviations

ABM Agent-Based Model.

AI Artifical Intelligence.

CA Cellular Automaton.

CAS Complex Adaptive System.

CFA Continuous Capacity and Flow Assignment.

CGP Constraint Generating Procedure.

CIM Computer-Integrated Manufacturing.

CLT Complexity Leadership Theory.

CNDP Continuous Network Design Problem.

ConWIP Constant Work-in-Process.

CPS Cyber-Physical System.

CSP Constraint Satisfaction Problem.

DAI Distributed Artificial Intelligence.

DCSP Distributed Constraint Satisfaction Problem.

DDM Distributed Decision Making.

DES Discrete Event Simulation.

DSA Distributed Search Algorithm.

FIFO First In – First Out.

FMS Flexible Manufacturing Systems.

GC Graph Coloring.

GCD Graph Coloring Dynamics.

GST General Systems Theory.

HMS Holonic Manufacturing Systems.

HPP Hierarchical Production Planning.

IoT Internet of Things.

MAS Multi Agent System.

MILP Mixed-Integer Linear Program.

MRP Material Requirements Planning.

MRP II Manufacturing Ressource Planning.

MSA Method of Successive Averages.

NDP Network Design Problem.

NE Nash Equilibrium.

OR Operations Research.

PFSL Preprocess First, Schedule Later.

PPC Production Planning & Control.

PROSA Product Resource Order Staff Reference Architecture.

RAS Resource Allocation System.

SA Simulated Annealing.

SMAS Self-Interested Multi Agent Systems.

WIP Work-in-Process.

List of Figures

List of Tables

Chapter One

Introduction

> *"What the advocates of unilateral centralization and unilateral devolution have both failed to realize is that the two processes must go hand in hand. History has, by its successive swings between the two extremes, clearly demonstrated that alone, neither is viable, whether it be in the social or in the technical field"*
>
> Hatvany (1985, p. 103)

On deciding, whether to entrust one central omniscient planning entity for the planning and control of manufacturing operations or to leave such tasks to multiple distributed entities, the history of Production Planning & Control (PPC) has seen the scale tip to both sides over the last 80 years.[1] This chapter will highlight how researchers and practitioners alike have turned (again) in the last 10 to 15 years to distributed PPC to cope with increasing complexity and variability in production environments. It will show a gap in knowledge concerning the design of such systems (Section 1.3), especially with respect to the widely held, but insufficiently investigated claim that a combination of classical hierarchical PPC systems and distributed approaches can deliver superior performance (Section 1.2).

Given this research gap and hypothesis, the chapter outlines the theoretical background and methodological approach adopted by this thesis and introduces the research questions (Section 1.4). The structure of the remaining chapters is outlined in Section 1.7.

1.1 Of Pulls and Pushes: A Changing Production Landscape for Production Planning & Control

1.1.1 Application Pull: Increasing Complexity and Flexibility Require New Approaches to Production Planning & Control

Koren (2010, Ch. 1) argues that global manufacturing changes its paradigm gradually from concepts of mass production (selling only a few products in large quantities) and mass customization (allowing the customer to choose from a variety of pre-determined options

[1] c.f. Sections 2.2.1 and 2.2.2.

to individualize the product) to *global manufacturing*, where globalized demand and supply leads to oversupply and increasing demand for regionalized (adapted to the needs and living conditions of a region) and even individualized (tailored toward individual customers) products.

It is then not surprising that many authors point to similar trends shaping future manufacturing (and hence PPC) systems: Windt and Hülsmann (2007) mention heterogeneous markets, demand volatility, international competition, short product life-cycles with frequent with product modifications on short notice, and a high number of variants in small lot sizes as key changes in the market and product environment. The National Research Council in the USA has come to a similar assessment of the trends affecting manufacturing companies: It mentions a shift toward creativity and the ability to produce individualized products to meet sophisticated customer demand in a more global and competitive market environment as the trends affecting manufacturing over the next 20 years (National Research Council 1998). To be successful in the future, companies will have to cope simultaneously with increasing complexity, scalability stress, uncertainties, and the necessity to adapt rapidly to changes in customer demand[2], all the while being challenged to meet tougher customer demands in terms of delivery-reliability and flexibility as well as keeping check on logistic cost (Windt and Hülsmann 2007).

For a long time, companies have been trying to cope with increasing external complexity by internalizing complexity: Wiendahl and Scholtissek (1994) provide empirical evidence for the increased complexity in the product, operations, and organization of production until the mid-1990s. This has traditionally been accomplished through the application of Material Requirements Planning (MRP) and Manufacturing Ressource Planning (MRP II) algorithms, where the different functions of PPC (c.f. Section 2.1.3) are solved sequentially through distinct, specialized, and centralized decision computer software, constrained by the decisions of their predecessors. As Section 2.2 will show in greater detail, such systems are well equipped to manage production in calm, stable production environments, but they are challenged where frequent adjustments and quick responses are necessary. However, owing to above-described developments in the market environment, the need for PPC systems to react swiftly to unexpected changes has become more significant over the past decades (T'kindt and Billaut 2006, Ch. 1.6.1; Gudehus and Kotzab 2009, Ch. 20.4.2). Going forward, it is widely assumed that to cope with the ever-increasing volatility and complexity, fundamental changes to the organization and control of production will become necessary. Monostori et al. (2006, p. 698) note: "Various solution proposals unanimously imply that the future of manufacturing lies in the loose and temporal federations of cooperative autonomous production entities". This view is shared by multiple authors including Bennett and Dekkers (2005), National Research Council (1998), Shen et al. (2006b), Windt and Hülsmann (2007), and Koren (2010, Ch. 9). Thus, managing integration and partnership among such autonomous entities is a key challenge for future logistics and manufacturing systems (Enarsson 2006, Ch. 1.7) requiring new coordination mechanisms (c.f. Sections 2.1 and 2.1.2) implemented in novel PPC systems. The decentralization of existing manufacturing systems and their

[2]Nof et al. 2006; similar conclusions in Maturana and Norrie 1996; McFarlane and Bussmann 2000, 2003; National Research Council 1998; Rahimifard 2004; Shen et al. 2006b; Brückner 2000, Ch. 2.2.2.

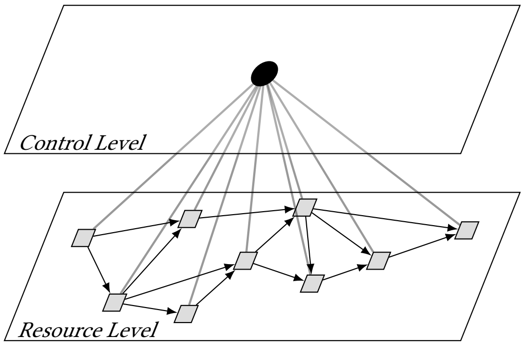

(a) Completely centralized control with only one controller.

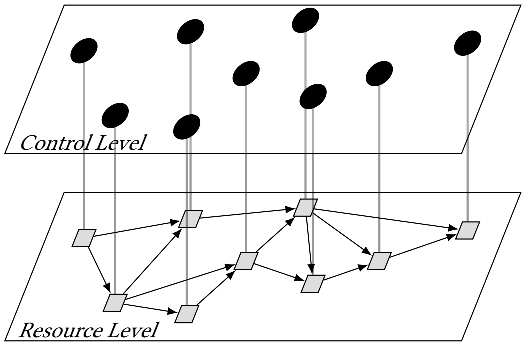

(b) Decentralized control. No communication between agents. All decisions are made on purely local information.

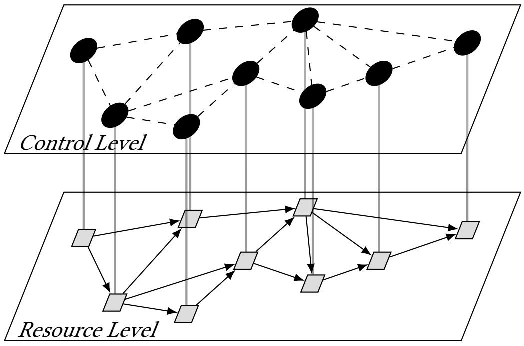

(c) Distributed control. Control is structurally decomposed into separate controllers for each resource. Controllers communicate among each other.

Figure 1.1: Three different configurations of control over a networked system of resources. Arcs represent material flow. Dashed lines represent information exchange. Gray lines represent control loops. Adapted from Becker et al. (2011).

control leads to a decomposition of the classically assumed hierarchy in PPC and a shift towards *self-organization* (Lasi et al. 2014, c.f. also Section 1.5.1).

As manufacturing system designers and planners across industries demand new solutions for their PPC challenges, and the devolution of planning and control authority, is commonly thought of as a feasible response, there is a clear practical motivation to better understand the working principles and success factors of such novel PPC systems. In accordance with definitions in control theory (c.f. Bakule 2008; Ge et al. 2017), this thesis will define as *hierarchical*, a control architecture, where only one decision-making entity is active at any point of time. Multiple parallel decision-making entities, which do not seek coordination (i.e. do not at least exchange information, c.f. Table 2.1), will be said to form a *decentralized* control architecture, while information exchange among them constitutes a *distributed* control architecture. Figure 1.1 gives a visual depiction of centralized, heterarchical, and distributed control architectures. Between the extremes of hierarchical and distributed (or "heterarchical") control architectures, this thesis will identify a design space for *semi-heterarchical* or *hybrid* control architectures (c.f. Fig. 2.2).

1.1.2 Technology Push: Cyber-Physical Systems Herald the Fourth Industrial Revolution

Technological advances and decreasing prices of computer and communication technology, as well as the advent of additional technologies at the interface of the digital and physical world (such as virtual reality, embedded sensors, additive production technologies, radio-frequency identification (RFID), etc.) have (re-)fueled the idea to equip physical products with computational and communication capacities, collectively forming a Cyber-Physical Systems (CPSs) (Rajkumar et al. 2010). Previously, the idea of *embedded systems* distributed computation capacity to physical items; however, only with the additional communication capacity can these former “Black Boxes” now interact with each other and operate in systems (Lee 2008). Collectively, such “intelligent” objects with the ability to communicate and coordinate (Meyer et al. 2009) form an *Internet of Things (IoT)* (Atzori et al. 2010; Gubbi et al. 2013).

The combination of the above-mentioned *application pull* in the domain of PPC (Section 1.1.1) with the *technology push* through CPS has lead to a renewed effort to promote digitization in industrial environments and the call to establish “networked production” systems (term from The Economist Intelligence Unit 2014; c.f. also Lasi et al. 2014). It is more commonly marketed as the advent of the “fourth industrial revolution”.[3]

To reap the expected benefits from the digitization of production, a multitude of national and international initiatives and research schemes have been set up. The European Commission (2016) has found initiatives in virtually every member state (c.f. Fig. 1.2), although a closer look by Gronau and Theuer (2015) has revealed stark contrasts with respect to funding level, -sources, and goals. Outside the continent, initiatives like the “Advanced Manufacturing Partnership 2.0” (USA), “Made in China 2025” (PR China), and “Manufacturing Innovation 3.0” (South Korea) have comparable goals (Blanchet and Rinn 2016). Notable international efforts include the ‘Physical Internet’ (PI, or π for short) initiative, which is primarily focused on transportation solutions powered by CPS (Ballot et al. 2012; Montreuil 2011; Montreuil et al. 2013; Sarraj et al. 2014).

Germany was among the first countries to launch research activities in this field with efforts to investigate networked production starting around 2010 (Gronau and Theuer 2015). In 2013 the German Academy for Science and Engineering (acatech) coined the term “Industrie 4.0” (Industry 4.0) for research activities aiming at the application of IoT concepts in production environments (Kagermann et al. 2013; Lasi et al. 2014). Blanchet and Rinn (2016) estimate that adopting the vision of Industry 4.0 can increase the return on capital employed by 25% and create over 10 million new jobs in Europe over the next 20 years.[4] Overall, McKinsey & Company (2015) estimates that Industry 4.0 can positively influence a wide range of performance measures, including a 3 – 5% increase in productivity, 20 – –50% reduction of inventories, and an increase in labor productivity

[3] The first three were triggered by water-/steam-powered mechanical manufacturing in the late 18th century, the division of labor in the early 20th century (c.f. Taylor 1911), and computer-based automation in the 1970s (Kagermann et al. 2013)

[4] The net figure (after accounting for job losses due to automation) is less impressive: about 1.4 million (Blanchet and Rinn 2016).

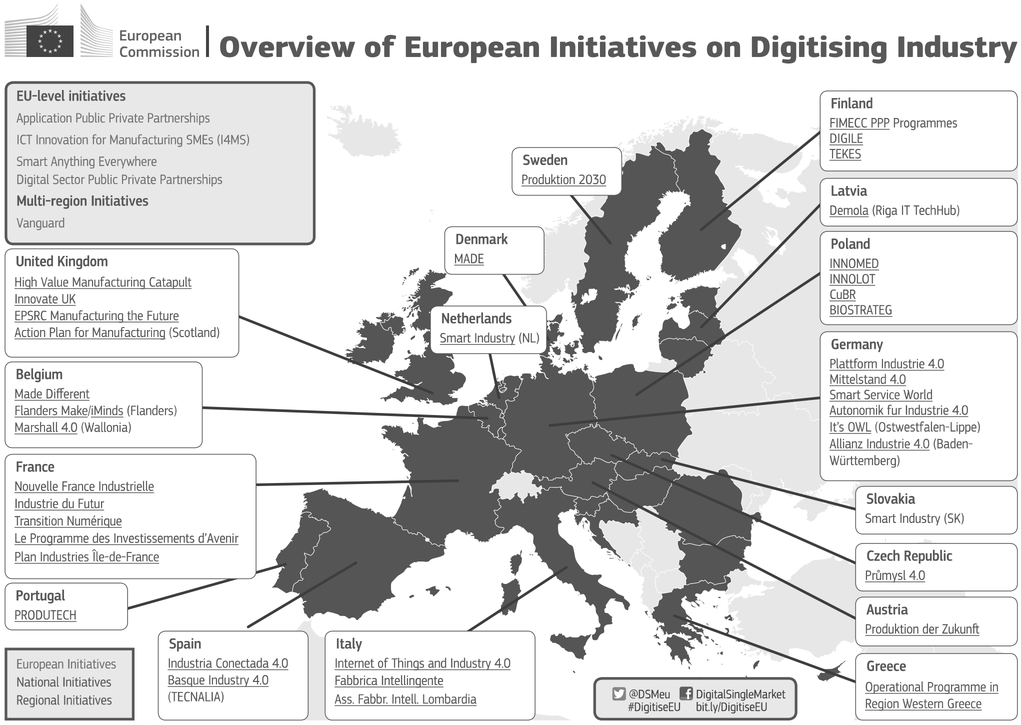

Figure 1.2: Regional, national, and European initiatives on digitizing industry (figure from European Commission 2016). Reproduced with permission (Ref. Ares(2017)160805).

by 45 – –55% (Fig. 1.3). Bauer et al. (2014) even expect productivity gains in Germany of up to 30% for some industries, including general manufacturing from 2013 to 2025.

While Industry 4.0 is used as an umbrella term to describe multiple aspects of digitization in production environments (c.f. Brettel et al. 2014, for a review), self-organization is a key concept within Industry 4.0 (Delfmann et al. 2017; Lasi et al. 2014) which (so it is hoped) “will lead to the emergence of dynamic, real-time optimised, self-organising value chains” (Kagermann et al. 2013, p. 20). Changing the organization of production toward self-organizing “Smart Factories” (Lasi et al. 2014; Zuehlke 2010) also seems to be the mostly anticipated application scenario for Industry 4.0 (in Germany), where Boston Consulting Group (2016) found that 64% of the queried companies (above $50 million in annual turnover) had already made or were investigating steps in that direction.

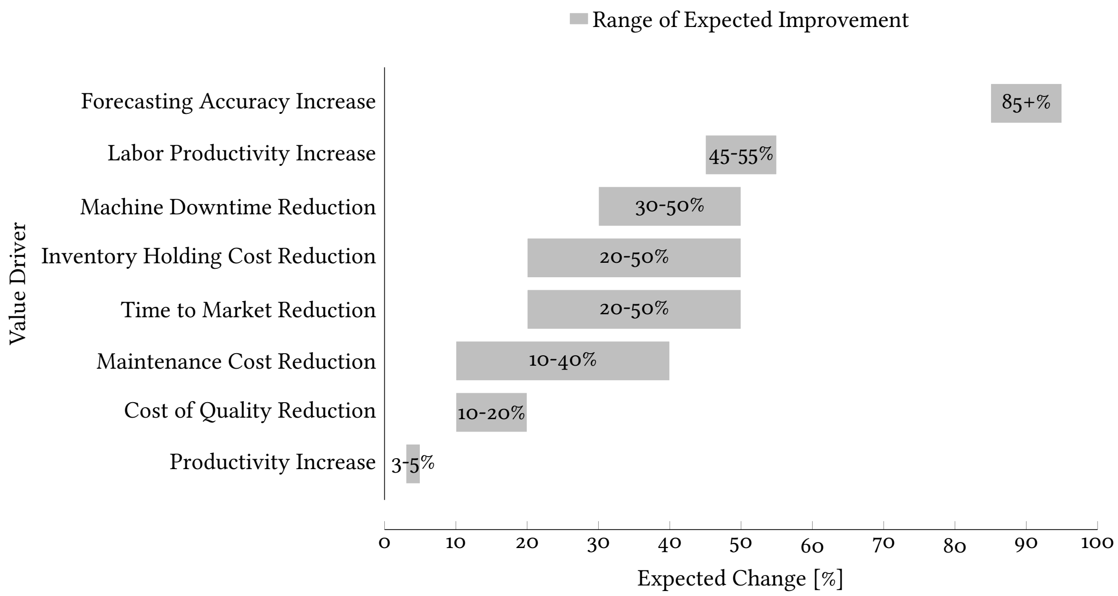

Figure 1.3: Expected Impact of Industry 4.0 on a manufacturing performance measures (McKinsey & Company 2015)

1.2 Hypothesis: A Combination of Centralized and Distributed Production Control Yields Best Performance

With many promises and significant funding allocated to the development and implementation of distributed control approaches, it is important to note that distributed (in particular: fully heterarchical) control architectures are unlikely to reliably yield excellent performance either. As will be further elaborated in Section 2.2.2, distributed control systems may exhibit erratic, unpredictable system behavior and lead to unsatisfactory levels of target achievement (c.f. e.g. Mařík and McFarlane 2005; Shen et al. 2006a; Trentesaux 2009). In fact, these negative traits (among other factors, like high investment cost, lack of standards, high communication effort between agents, ...) have so far impeded the adoption of distributed control architectures in practice.[5]

The observed problems with both hierarchical and (strictly) heterachical/distributed control approaches have contributed to the rising interest in the investigation of *hybrid* or *semi-heterarchical* control approaches. Cardin et al. (2015) and Pach et al. (2014) use the label "hybrid" to designate control approaches that "are intended to capitalize the advantages of reactive and predictive/proactive approaches, while limiting their drawbacks" (Cardin et al. 2015, p. 3).[6] As Section 2.3 will further elaborate, the hypothesis that a combination of hierarchical and distributed control traits may prove to be advantageous

[5] c.f. e.g. Askin and Goldberg 2002, Ch. 12.1.2; Dilts et al. 1991; Mařík and Lažanský 2007; Mařík and McFarlane 2005; McFarlane and Bussmann 2003; Shen et al. 2006a; Trentesaux 2009; Vánca 2014.

[6] The term *semi-heterarchical* is also used, e.g. by Berger et al. (2010), Zambrano Rey (2014), and Zambrano Rey et al. (2013).

regarding the overall system performance has widely been articulated (c.f. e.g. Cochran and Kaylani 2008; Giret and Trentesaux 2015; Philipp et al. 2006; Zambrano Rey et al. 2014). A particularly visual expression of this idea was presented in Philipp et al. (2007) and Scholz-Reiter et al. (2009a), who have visualized a *hypothesized* curvilinear relationship between the "degree of autonomous control" and logistic target achievement as shown in Fig. 1.4. The authors assume that the logistics performance (z-axis) reaches a maximum for medium levels of autonomous control (x-axis), with the overall performance (and the shape of the relationship between autonomous control and performance) being further modulated by the degree of complexity in the system (y-axis, going through the paper plane).

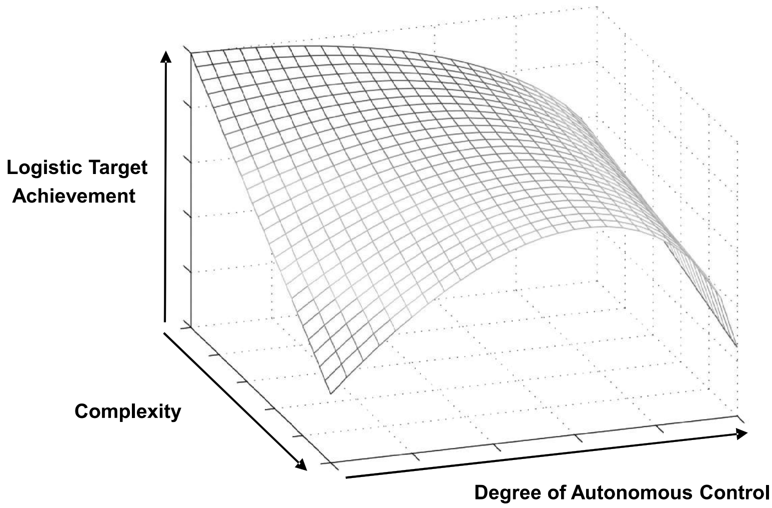

Figure 1.4: Hypothesized relationship between degree of autonomous control, complexity, and performance (Philipp et al. 2007; Scholz-Reiter et al. 2009a). Reproduced with permission (RightsLink® License Number: 4021260904170).

1.3 Research Gap: Designing Manufacturing Systems for Distributed Control

The assumed advantages of hybrid control architectures can only be realized if designers are able to purposefully engineer PPC systems along the entire spectrum of hybrid control architectures. Here current research is found to be insufficient. Distributed PPC systems — this section will show — require new and fundamentally different design guidelines as compared to their hierarchical ancestors. Research activity in the domain of distributed PPC and related disciplines has hitherto been unable to provide such new design guidelines beyond institutionalized experience and "gut feeling".

Hierarchical PPC systems are designed according to a top-down *reductionist* approach where problems are approached through problem decomposition and system implementation by a single design team. System behavior in this context can be "hard-coded" into the

	Reductionist View	Constructionist View
Motivation	Devise a system of interacting agents that work together to solve a common problem	Devise agents which interact to further their own needs
Measure of Goodness	Overall system performance	Performance of individual agents
Agents' Decision-Making Function	Benevolence — accept all requests made	Individual utility maximization
System Coherence	Carefully engineered by single design team	Emergences out of interplay between agents
Main Drawback	Fails to fully exploit concept of autonomous agents — too much system level design	System behavior defined through human-intensive refinement of individual agents

Table 1.1: Comparison of Reductionist and Constructionist approach with Multi-Agent System Design (Kalenka and Jennings 1999).

system. Designers of distributed control architectures, on the other hand, build production control systems by designing the basic behavior of agents *first* and by relying on the concepts of *self-organization* and *emergence* to yield a coherent global system behavior from these building-blocks (bottom-up, *constructionist* approach) (c.f. Section 1.5.1). "In other words, the decision-making knowledge stored locally in the agents causes the global behavior of the system to operate in a way that cannot be precisely predicted" (Mařík and Lažanský 2007, pp. 1371 f.). Table 1.1 compares the reductionist and constructionist approaches to system design.

While an exact top-down approach (which local rules collectively yield a desired system-level behavior x) would certainly be most helpful to system designers, it has repeatedly been shown to be extremely difficult and computationally exhaustive even for very simple model classes, such as CAs, to derive such cause-and-effect relationships between local rules and global behavior (c.f. Deutsch and Dormann 2005, Ch. 4.4). Designing distributed control systems then requires new and unique design considerations from the system designer.[7] The main challenge is to combine local decision-making behavior in such a way that a globally coherent behavior emerges that is aligned with the objectives for that production system (c.f. Cavalieri et al. 2000; Crutchfield and Mitchell 1995; Delfmann et al. 2017). For Kalenka and Jennings (1999, p. 136), "what is required is the ability to exploit the conceptual power of autonomous agents (as in the constructionist view), but to ensure the overall system performs in a coherent manner (as in the reductionist view)". However, they consider these two demands contradictory (ibid.).

[7] so stated by Cavalieri et al. 2000; Delfmann et al. 2017; Dilts et al. 1991; Plattform Industrie 4.0 2015; Rogers and Brennan 1997; Scherer 1998; Trentesaux 2009; Van Dyke Parunak 1999; Verstraete et al. 2008a.

Defining any type of control system requires making a series of design decisions that collectively define a *control architecture* (c.f. Section 2.2). For agent-based PPC systems, Brückner (2000, Ch. 2.2.3), Lesser and Corkill (1981), and Shen et al. (2006a) all identify a similar set of major design decisions. These include determining (1) which objects are represented by agents, (2) agent encapsulation (form of problem decomposition), (3) agent modeling (knowledge base, computation-, and communication capabilities), (4) system structure (communication and interaction structure among agents), and (5) coordination and negotiation protocols (how do agents find a common solution).

Practical advice to address these design decisions, however, is practically absent and the design of distributed PPC systems is generally assumed these days to entail significantly more experience and "black art" than what we have come to expect from reductionist approaches (Kalenka and Jennings 1999). Monostori et al. (2006, p. 714) state "When appealing for emergent functionalities in complex systems, we seem to go back to the old traditions of engineering. Long ago, the words like 'machine', 'mechanical', or 'engineer' did not refer to rationality but rather to trickery, artifice and machinery [...]. When dealing with complex systems in this way, we try to elicit effects which are beyond the limits of our actual knowledge."

The absence of a scientifically founded engineering approach for distributed systems was already observed in the domain of Multi Agent System (MAS), the scientific discipline concerned with the investigation of systems of multiple independent agents (Section 2.1.2). Jennings and Wooldridge (1995, p. 366) note: "Surprisingly little work has been undertaken on methodological aspects of agent-based systems, and yet if this technology is to be a commercial success, then designers must have a structured way of developing well-engineered agents and agent systems. This work needs to identify how robust and flexible individual agents can best be designed and how these well-designed components can then be combined to give an overall system that is similarly robust and flexible. The problem of providing a system-level description is made more difficult by the fact that many of its properties can only be observed at runtime." And while Müller (1997, p. 230) points to the existence of several general design approaches discussed in the literature, he finds that "currently there are no rationals and no criteria which may be used to support the decision for a special approach or technique". In the absence of rationale with predictive power about the behavior of agent systems, the development process for agent-based PPC systems is today largely driven by experience (Farid and Ribeiro 2015; Van Dyke Parunak 2000) and/or simulation experiments that seek to predict the collective behavior of agents, giving this "trial-and-error"-based methodology to the development of PPC architectures a far more pronounced role in the design of agent-based PPC systems as compared to their hierarchical counterparts (Mařík and Lažanský 2007).

Against this backdrop, the need to develop new design methodologies, aimed at the design of agent (or CPS) based PPC systems, has repeatedly been expressed.[8] Botti and Giret (2008, p. 20) observe that "To date, almost all of the applications in the Holonic Manufacturing Systems (HMS) field have been built using no design or development method" (assessment shared by Farid and Ribeiro 2015). Where methodologies (at least in

[8] c.f. e.g. Delfmann et al. 2017; Farid and Ribeiro 2015; Giret and Trentesaux 2015; Plattform Industrie 4.0 2015; Botti and Giret 2008, Ch. 2.3.5; McFarlane and Bussmann 2003.

foundations) exist, they are built on experience and intuition (Farid and Ribeiro 2015; Van Dyke Parunak 2000). Instead, case-based comparisons are performed. Experimentation testbeds have repeatedly been proposed (c.f. Cavalieri et al. 2003; Rogers and Brennan 1997; Trentesaux et al. 2013), but no standard has emerged so far. Trentesaux (2009) claims that for more comprehensive design methodologies to emerge, the development of "emergence engineering" is a necessary precursor, allowing to control the general direction of emergent properties. Namatame and Sasaki (1998, p. 189) specify the questions to be answered: "We need to understand how to set up the architecture of an agent as a component of a complex system suitable for evolution, how self-interested behavior evolves to cooperative behavior, and how the goal structure of each agent can be self-modified in order to achieve the common goal."

The observed absence of design guidelines is fundamentally connected with the observed unpredictability of distributed PPC systems. In particular, an understanding of the functional principles behind distributed control and its success factors is missing. Accordingly, Monostori et al. (2006, p. 714) believe that "Getting back to what is now considered principled engineering, we need further research in characterizing interactions that may or may not produce emergent phenomena, explore its root causes (such as the dimensionality and connectivity of agents, the flow of information among them, and the propagation of constraints) and develop predictive theories." Tsoukas (1996) agrees that for researchers to pose "if, then"-statements, studies on aggregate systems must be conducted that would render such statements plausible and reliable for application.

Philosophical discussions (as by Abbott 2006), give "us the hope, at least, that engineering emergence is not an impossible dream" (Stepney et al. 2006, p. 5). Moreover, some more practical attempts have been made to provide a more solid footing for the development of distributed control architectures in the context of PPC. Zapf and Weise (2007) present an approach for the offline engineering of agent societies that essentially constitutes an iterative refinement process, steered by a genetic algorithm, to guide the system behavior to the desired behavior. Scherer (1998) derives design principles for (human-operated) production systems from cybernetics theory. Brückner (2000) derived design principles from synthetic ecosystems as candidate principles for the design of production control systems to achieve desirable global properties, including robustness, agility, and flexibility, yet without providing further evidence for the validity of these principles. Focusing on agent-based production control, the design methodology *ALEM* (autonomous logistics engineering methodology) is presented in Scholz-Reiter et al. (2009b), based on the UML notation used in classic software/system design approaches. The determination of the system structure and agent abilities are seen as central steps in the 8 step procedural model (ALEM-P), but no further recommendations are given as to how these decisions should be made. Fischer et al. (2003) present an abstract specification of a holonic multi-agent system, which, they say, can be applied also to HMS. The paper, however, gives little instruction as to how to solve the practical design problems described above. Similarly. and in the same year, Leitão et al. (2003) use a petri-net based approach to formally specify agents in ADACOR, an HMS flavored PPC reference architecture, again without giving concrete design guidelines. Both approaches aim at formal validity and not at system-performance. Most recently Farid and Ribeiro (2015) proposed using principles and methods from axiomatic design to control reconfigurable production systems through a multi-agent controller. Following the ideas of axiomatic design, they first develop (qualitative) design principles which are observed while design alternatives are described in terms of a knowledge base and constraint matrices. The approach does not make any statements about the likely performance of different design alternatives, nor does it address the above-mentioned apparent necessity for combining heterarchical and hierarchical design traits.

1.4 Research Questions

Sections 1.2 and 1.3 have shown problems concerning the practical implementation of distributed PPC systems that are closely tied to the lack of knowledge about the design of such systems. The academic community then should gain a better understanding of how to design distributed PPC systems in a constructionist bottom-up fashion to preserve the advantages associated with their distributed, reactive nature while ensuring sufficient levels of target-achievement and overall well-mannered system behavior. Monostori et al. (2006, p. 713, highlighting in original was removed) state that "the key issue is how we can engineer multi-agent systems that exhibit purposive, goal-directed oriented behavior at the system level by relying on their emergent nature. Generally, how can we design for emergence?" The research-guiding question picks up this line of thought:

> Q_0 (Research-guiding question)*: "How can both the plant and the controller of manufacturing systems be designed to achieve high logistics performance under distributed control?"*

Given the inherently intractable properties of design problems (Grünig and Kühn 2013), and the vast amount of literature streams and research fields associated with the design and analysis of MAS (c.f. Section 3.2.1), this question requires a multi-faceted, stepwise approach, which is reflected in the chapters of this thesis as well as in the derived sub-questions in the following.

In particular, the course of research presented in this thesis is guided by the above-discussed hypothesis that, given the advantages and disadvantages of both hierarchical and distributed PPC, a combination of hierarchical and distributed (production) control should achieve (under broad conditions) an optimal system performance (Section 2.3). This thesis will investigate this hypothesis both qualitatively and quantitatively.

Hierarchical and distributed control will be shown in Section 2.2 to differ on a variety of dimensions, including the form of problem decomposition, applied coordination principles, etc.. A "combination" of these two poles can be achieved along a variety of dimensions as well, which need to be known and understood, when outlining the design space of hybrid production control architectures. Research question Q_1 aims at this broad, qualitative understanding:

> Q_1*: "Which design decisions concerning both controller and plant impact the duality between hierarchical and distributed control in PPC?"*

The discussion so far, as well as the following classification in Chapter 3 (in particular Section 3.3.2), highlights that notions of hierarchy play a central role not only in the differentiation between architectural styles, but also in the discussion on how to "tame" distributed control architectures. Thus, it is appropriate that the concept of hierarchy should be given special attention in this thesis. Research question Q_2 places it in the context of the hypothesized advantages of hybrid control architectures.

> Q_2: *"Which degree of hierarchy in the controller results in the highest logistics performance of a manufacturing system and why?"*

Finally, the research-guiding question intentionally goes beyond designing the distributed control of a (given) production system. Thus, it is not only acknowledged that the physical structure of a production system is likely to have an effect on the performance of a distributed control system (c.f. Section 3.3.1), but the question is raised, if a the physical structure of a production system can be designed in such a way that distributed control exhibits desirable properties. Research question Q_3 extends the research questions to be discussed in this thesis in this direction.

> Q_3: *"How can a production plant be designed to entice selfish agents to exhibit predictable and desirable emergent system properties?"*

With respect to research question Q_0, the first derived research question aims at a qualitative description and formalization of the design space. An answer to research question Q_1 will help the designer of manufacturing systems under distributed control to understand the problem and decision alternatives available to them, as they try to maximize performance. Research questions Q_2 and Q_3 pick up a particular design dimension and subject them to rigorous quantitative analysis.

1.5 Approach

The investigations in this thesis are based on several interrelated ideas and theories put forward to discuss and quantitatively analyze manufacturing systems. They are presented here.

The concepts of *emergence and self-organization* (Section 1.5.1) provide the conceptual basis to understand the rise of macro-level patterns in distributed systems. It has revolutionized our understanding of systems, towards CAS (Section 1.5.2). CAS theory does not only describe such systems, but it has galvanized a set of models and disciplines that have set out to investigate them, some of them will be drawn upon in this thesis (Section 3.2.2). Sections 1.5.3 and 1.5.4 introduce two research fields that are actively concerned with the prescriptive design of distributed systems — a challenge facing also designers in the production domain.

1.5.1 Emergence and Self-Organization

Reductionism, the idea that observed system-behavior can and should be explainable and eventually controllable through decomposition into subsystems and subproblems (a priori analysis) and reduction of the phenomenon (a priori reduction) (Fuenmayor 1991) has long dominated the analysis of natural and man-made systems and shaped disciplines like Operations Research (OR) and classical understandings of PPC (Ackoff 1973, 1974; Bunimovich 2001; De Wolf and Holvoet 2005; Innes and Booher 1999). While its merits remain unchallenged and so far unparalleled (Barabási 2012; Paslack 1991, Ch. 2.1), the principle of reductionism is considered to have reached its limits as researchers have turned their attention to larger and more complex systems, noting that in fact not all observed system-level behavior can be explained through its parts (Bar-Yam 2002; Barabási 2012).

To understand how a system can be "more than the sum of its parts"[9], researchers have discovered two principles: *self-organization* and *emergence*. Despite subtle differences, both self-organization and emergence are dynamic processes — i.e. they arise over time (De Wolf and Holvoet 2005; Goldstein 1999). Mostly (especially in the context of systems engineering), both phenomena are observed and exploited jointly (De Wolf and Holvoet 2005); they are applied today in many scientific and engineering disciplines (Paslack 1991, Ch. 1; Vec et al. 2006; Windt and Hülsmann 2007).

Self-Organization describes the spontaneous evolution of order (collective, spatio-temporal patterns) in dynamical (far from equilibrium) systems.[10] Bonabeau et al. (1999, p. 9) define self-organization as "a set of dynamical mechanisms whereby structures appear at the global level of a system from interaction among its lower-level components". It is central to the notion of self-organization that order arises from simple rules of interaction and in the absence of an external control in the process (i.e. the order emerges autonomously), which means that "there is no dichotomy between the organizer and the organized" (Deutsch and Dormann 2005, p. 30) in self-organized systems.

Emergence, on the other hand, "refers to the arising of novel and coherent structures, patterns" (Goldstein 1999, p. 49). The focus here is (a) on the novelty of the emergent phenomenon (it was not explicitly encoded in the behavior of the subsystems) and (b) on the difference in the levels between the contributing system parts (micro) and the resulting, observed property (macro) (Abbott 2006; De Wolf and Holvoet 2005; Holland 2002; Stepney et al. 2006). Goldstein (1999) requires systems to fulfill four requirements (nonlinearity, self-organization, beyond equilibrium state, attractors) to be able to exhibit emergence. Holland (1998, Ch. 7) define a general class of systems called Constraint Generating Procedures (CGPs) that (he claims) exhibit emergence. A system(-model) is a member of the class if it *generates* (i.e. exhibits during run-time) *procedures* (dynamic behaviors) subject to *constraints*. In particular, MAS represent a suitable modeling approach to capture such emergent dynamic behavior (Bonabeau 2002).

[9] An idea that dates back to ancient Greek philosophy, presumably Aristotle (Klir 1991, Ch. 3; von Bertalanffy 1972).

[10] Bonabeau et al. 1999, Ch. 1.2.2; De Wolf and Holvoet 2005; Goldstein 1999; Marzo Serugendo et al. 2004; Paslack 1991, Ch. 1; Prigogine and Stengers 1984; Toffler 1984.

Conceptualizing distributed PPC as a process of self-organization and emergence helps this research to broaden the view beyond the originally considered domain. It allows to identify additional research streams, where similar problems have been subject to academic debate which can be drawn upon (c.f. e.g. Section 3.2.2).

1.5.2 (Complex Adaptive) Systems Theory

"Intellectual movements meant to replace reductionism with an appreciation for modeling interactions instead of simplifying them away" (Anderson 1999, p. 219) have a surprisingly long history:

As Anderson (ibid.) reiterates, the investigation of complex systems has reverberated through the 20th century, starting with debates on holism and gestalt theory (c.f. De Wolf and Holvoet 2005; Goldstein 1999; Klir 1991, Ch. 3; Paslack 1991, Ch. 2.7) after World War I and finding renewed attention after World War II, with the advent of closed-loop controllers that led to the rise of *cybernetics* (Ashby 1961; c.f. also Klir 1991, Ch. 3) and the evolution of the General Systems Theory (GST) (Ackoff 1974; c.f. von Bertalanffy 1972, for a review by the originator of the theory). The third wave of scientific engagement with complex systems was caused by the advent of *chaos theory* in the 1960s, sparked by Lorenz's discovery that even simple, well-understood, and deterministic relationships could produce behavior that may appear random at large scales, but are yet bound by the laws of physics and mathematics (Levy 2000).

CASs theory is different from the above-mentioned study of dynamical systems in that it does not consider the evolution of a system, expressed as mathematical expressions, over time, but seeks to understand complex systems through the interaction of its components (Anderson 1999; McMillan 2008, Ch. 3). Unlike in chaos theory, where the state of chaos was the subject of interest in itself (Levy 2000), CAS theory investigates the rise of *order* through self-organization and emergence (Anderson 1999).[11] Compared to the equilibrium-based assumptions about systems in the reductionist approach, CAS theory emphasizes the systems at the "edge of chaos", showing bottom-up emergent system properties and adaptation to its environment (Anderson 1999; Schneider and Somers 2006). A central feature of CAS is their adaptiveness due to both the interplay between the system and its environment (Choi et al. 2001) as well as the ability of agents to change their behavior over time (adapting and learning) (Holland 2002). For Anderson (1999, p. 219) "The hallmark of this perspective is the notion that at any level of analysis, order is an emergent property of individual interactions at a lower level of aggregation".

CAS thinking has been applied to manufacturing systems (Nilsson and Darley 2006), manufacturing networks/supply chains (Brintrup et al. 2015; Choi et al. 2001; Surana et al. 2005)[12], as well as the study of social systems (Bonabeau 2002; Castellani and Hafferty 2009; Plowman et al. 2007).

[11] Complexity Science and CAS can, therefore, be seen as the intersection of General Systems Theory (GST) (*systems thinking*) and cybernetics (Castellanos 2012, Ch. 2).

[12] Solow and Szmerekovsky (2006) mention supply chains as a prominent example where complex system understanding can be transferred to a business context.

A CAS perception of production will be applied especially in Chapter 4, where it provides the underpinning for the applied modeling approach and the performed analysis (e.g. the quest to identify *order* in Section 4.5.1). The bottom-up nature of CAS also helps to motivate the research idea followed in Chapter 5, where the thesis studies the physical design of networks that arise in anticipation of agents' decision-making.

1.5.3 Complex Leadership Theory (CLT)

Complexity and CAS-theory have had a significant and intensely debated impact on organization theory (Anderson 1999; McKelvey 1999). Here (not unlike research in the domain of PPC, c.f. Section 2.2.2) researchers started to contemplate in the 1990s whether decentralized, non-hierarchical networks could be better organizational structures for highly adaptive and innovative organizations (Levy 2000; Stacey 1995) and hence the scientific field has (partially) moved away from reductionist thinking (McMillan 2008, Ch. 2; Marion and Uhl-Bien 2001), replacing the ideas of the GST with CAS for the conceptualization of organizations (Schneider and Somers 2006). Today, CAS are a frequently used conceptualization also in organizational research (c.f. e.g. Stacey 1993).

Such new perception requires new theories about the role and functioning of leadership in complex systems (Uhl-Bien et al. 2007). Uhl-Bien et al. (ibid., p. 298) note that "Leadership models of the last century have been products of top-down bureaucratic paradigms. These models are eminently effective for an economy premised on physical production but are not well-suited for a more knowledge-oriented economy". This perception has lead to the rise of Complexity Leadership Theory (CLT) as a new model of leadership in complex systems, suitable "for the knowledge era" (Uhl-Bien et al. 2007, p. 299; Marion and Uhl-Bien 2001, c.f. also). Leadership under this new framework is not a priori assigned to individual agents/roles, but an emergent property (Uhl-Bien et al. 2007). Morgan (2006, Ch. 8) argues that one of the consequences of complexity perceptions of organizations is to "Rethink what we mean by organization, especially the nature of hierarchy and control" (ibid., p. 255).He argues that in settings where organization cannot be externally imposed, hierarchy can be "generated by the need to cluster and direct activities to address the contingencies at hand" (ibid., p. 256). This leads to a changed role of leaders(hip): instead of controlling the evolution of the system, leaders (now) *enable* the future development of the system (Marion and Uhl-Bien 2001; Morgan 2006).

In that, Complexity Leadership Theory (CLT) picks up a long -held debate within organization theory: since the "Hawthorne Experiments" in the 1930s (c.f. Section 2.2.2), there has been ongoing interest in an assumed "control vs. autonomy duality" (Thomas et al. 2005b), where an "optimal mix" between control and autonomy (or exploitation and exploration) is classically assumed (c.f. March 1991), but has more recently been challenged by a view that advocates a cyclical change between control and autonomy (Thomas et al. 2005b).

Another feature of CAS conceptions of leadership and organization, relevant for this thesis is that it makes organization theory open to computerized experiments. Lewin et al. (1998) argue that one key advantage of the CAS approach is to open social sciences to analytical approaches, e.g. from physical sciences. They point, in particular, to the possibility of developing an understanding of the phenomena of emergence in social systems: “The real potential of this modeling approach, however, is in identifying which parameters are important in establishing the emergent culture of the workplace and creating simulations that help firms to discover the reliable interventions [...] that management should make” (Lewin et al. 1998, p. 37). Hence, analytical models are actively promoted as sources to build and verify theories about leadership (Carley 1995; Carley and Gasser 1999; Hazy 2007, 2008; Hazy et al. 2007). In particular, agent-based models are recommended to capture the dynamics of complex organizations (Lazer and Friedman 2007; Lichtenstein et al. 2006; Siggelkow and Rivkin 2005) and to investigate the relevant impact factors to influence emergent behavior (Lewin et al. 1998; Lichtenstein 2007). Examples of the use of minimal models to explore the impact of organizational structure on performance are given e.g. in Carley (1997), Lazer and Friedman (2007), Siggelkow and Levinthal (2003), and Siggelkow and Rivkin (2005) (c.f. Section 4.1.2).

This thesis will apply CLT in Chapter 3 as one field of research — outside PPC — which considers the active design of systems (organizations) in which individuals cooperate toward a goal, and in Chapter 4, where it provides testable hypotheses for the role of leaders in the context of CASs.

1.5.4 Game Theory

Where decisions are made in the presence of other actors, agents have to consider the others’ responses in their decision making (Klein and Scholl 2011, Ch. 1.4.3). Game theory is an essential tool for modeling such interactions between (selfish) decision makers and the resulting collective behavior (c.f. Ossowski 1999, Ch. 2.3.2; Monostori et al. 2015; Argoneto et al. 2008, Ch. 2). It is an important instrument to understand emergent behavior as it analytically captures the reactive relationship between agents’ decisions and hence “encourages a more careful look at emergence in rule-governed systems” (Holland 1998, p. 42).

Game theory cannot only provide descriptive insights into situations where individual players (agents) interact, but also yield prescriptive hints for system design (Marden 2016; Marden and Shamma 2015). The branch of game theory concerned with the purpose-driven design of algorithms to shape game outcomes is called *algorithmic game theory* (initiated by Nisan and Ronen 1999; more recent overviews given in Marden and Shamma 2015; Nisan et al. 2007). It has been applied to resource allocation problems, especially in the domain of traffic planning (this thesis will build on this work in Chapter 5).

This thesis will draw upon game theory especially in Chapter 5, as a mean to predict analytically the emergent behavior of (infinitely) many selfish agents. Like CLT, game theory will also be drawn upon in Chapter 3 repeatedly, e.g. as motivation for a measure of myopic behavior (Section 3.1.3).

1.6 Methodology

To address the research gap, based on aforementioned theoretical foundations, this thesis applies two different methodological approaches, with a literature foundation laying the groundwork for subsequent quantitative analysis through minimal model experiments.

1.6.1 From Existing Literature to a Design Space Conceptualization

This thesis draws upon multiple streams of literature to conceptualize, model, and analyze manufacturing systems. The disciplines reviewed include — besides the literature published in the domain of PPC — game theory, statistical physics, scheduling theory, and organization theory. This broad base allows to conceptualize a design space for hybrid PPC systems in Chapter 3, situated between the "poles" of purely hierarchical and purely distributed control, and to identify design decisions that can affect the positioning of a given system between the two.

The result of the literature reviews in Chapters 2 and 3 is the classification model shown in Section 3.6. It aggregates the understanding developed in Chapter 2 and the findings of the reviewed disciplines of Chapter 3 into an actionable frame of reference for designers of manufacturing systems under distributed PPC. The classification model helps a system designer to find the right balance between hierarchical and distributed control for the case of application.

1.6.2 Exploring Complex Systems Through Minimal Model Simulations

Analytical experiments in this thesis are performed using (more or less) minimal models that can be applied and analyzed as abstractions of manufacturing systems.

Over the years, a rich set of model classes for the investigation of manufacturing systems has been developed.[13] Within this large model landscape, a strong argument is to be made for highly abstract, minimal models, when seeking to discover the mechanics behind observed phenomena, such as the success or failure of distributed production control approaches: Holland (2002, p. 29) says about the analysis of CAS: "Accordingly, anything considered irrelevant to the question should be considered a detail that can be eliminated from the model. A model much like a political cartoon, makes its points by exaggerating certain features while eliminating incidentals." Generally, a more simple model with less independent parameters will also make it easier to "discover and understand the subtle effects of its hypothesized mechanisms" (Axelrod 2007, p. 24) through simulation.

Minimal models can be useful even when they do not directly resemble the system under investigation or leave out important aspects of the real-world system. They may be used

[13] c.f. Cassandras and Lafortune 2008, Ch. 1.3.3 and 10.3; Reveliotis 2005, Ch. 1; Kuehnle 2007; Leung and Suri 1990; Papadopolous et al. 1993, Ch. 1.5; Zimmermann 2008, for reviews and classifications.

instead akin to *gedanken experiments* in physics: an intellectual exercise aimed not at the collection of data comparable with a real world system, but as an investigation of the dependence between observed behavior and starting conditions (design decisions) (Holland 1998, Ch. 12). They have a strong track record, in particular, in the investigation of complex systems where contributions like the *Sandpile Model* (Bak et al. 1987) that introduced the concept of *self-organized criticality*, or the *Kuramoto Model* of interacting oscillators (Kuramoto 1984) as a minimal model of synchronization processes, have found ample applications far beyond the original model setting.

The explanatory power of minimal models has also been discussed from a science-philosophical perspective by Batterman (2002), Batterman and Rice (2014), and Burton and Obel (1995). For Batterman and Rice (2014, p. 375), minimal models are "a class of explanatory models that are explanatory for reasons that have largely been ignored in the literature. These reasons involve telling a story that is focused on demonstrating why details do not matter". Batterman (2002) points, in particular, to the ability of minimal models to understand emergent properties by arguing that, by describing asymptotic behavior, minimal models in domains plagued with (seemingly) unpredictable behavior can extract *stable phenomenologies*, describing the forces at play in a system of many particles. In conclusion, for Batterman (ibid., p. 22) "a good model is one which doesn't let a lot of [...] details get in the way. In many cases, the fine details will not be needed to characterize the phenomenon of interest, and may, in fact, actually detract from an understanding of that phenomenon".

In the two analytical chapters of this thesis (Chapters 4 and 5), two minimal models will be applied to investigate the emergent collective behavior agents in manufacturing environments. The model in Chapter 4 provides a highly simplified representation of a scheduling problem where operations have to be assigned to machines. However, interdependencies between operations and differences within the operations and machines are neglected. Chapter 5 considers the flow of products through a shop with process alternatives. The model considers the queuing caused by agents arriving at machines and the response of intelligent products to perceived differences in process path "attractiveness", but it only models a fluid approximation of the long-term average routing behavior, ignoring e.g. differences in processing times and the difference between adjusting the capacity of one machine and acquiring multiple, parallel machines.

1.7 Thesis Outline

In the next chapter (Chapter 2), the concepts of "hierarchical" and "distributed" control are introduced in more detail. A review of their historical development, strengths, and weaknesses will allow this thesis to develop the research hypothesis more broadly. Chapter 3 will introduce the concept of *Myopia* to describe the design space between them. Using evidence from a broad range of disciplines, the chapter concludes with a classification model that associates decisions along the decision steps in manufacturing system design and operation with possibilities to increase or decrease the degree and/or impact of myopic behavior, thereby describing a design space for hybrid PPC systems between the (stereotypical) poles of hierarchical and fully decentralized control.

The remaining chapters then explore specific design decisions using minimal quantitative models. Table 1.2 places these quantitative chapters in the context of the developed classification model of Chapter 3. In particular, Chapter 4 provides unprecedented insights into the mechanisms of coordination among an agent population and the role of hierarchy in that process, thus providing evidence and a conceptual understanding of the role of hierarchy. Chapter 5 investigates how environmental changes — in particular changes to machine capacity — can shape agent behavior and yield predictable and desirable system-level properties.

The thesis is concluded in Chapter 6 with a discussion of the results for the practice of manufacturing system design and in the light of the invoked theoretical foundations. Also, an outlook on recommended future research is given. Figure 1.5 provides an overview of the chapters and their interconnection.

Chapter	Modeling Approach	Applied Countermeasures	Addressed Dimension of Myopia	Research Question
4	CA/statistical physics	Introduction of hierarchy	loss of performance	2
5	(algorithmic) game theory	control of flexibility	unpredictability, high social cost	3

Table 1.2: Overview of the design problems investigated analytically in the context of the classification derived in Chapter 3.

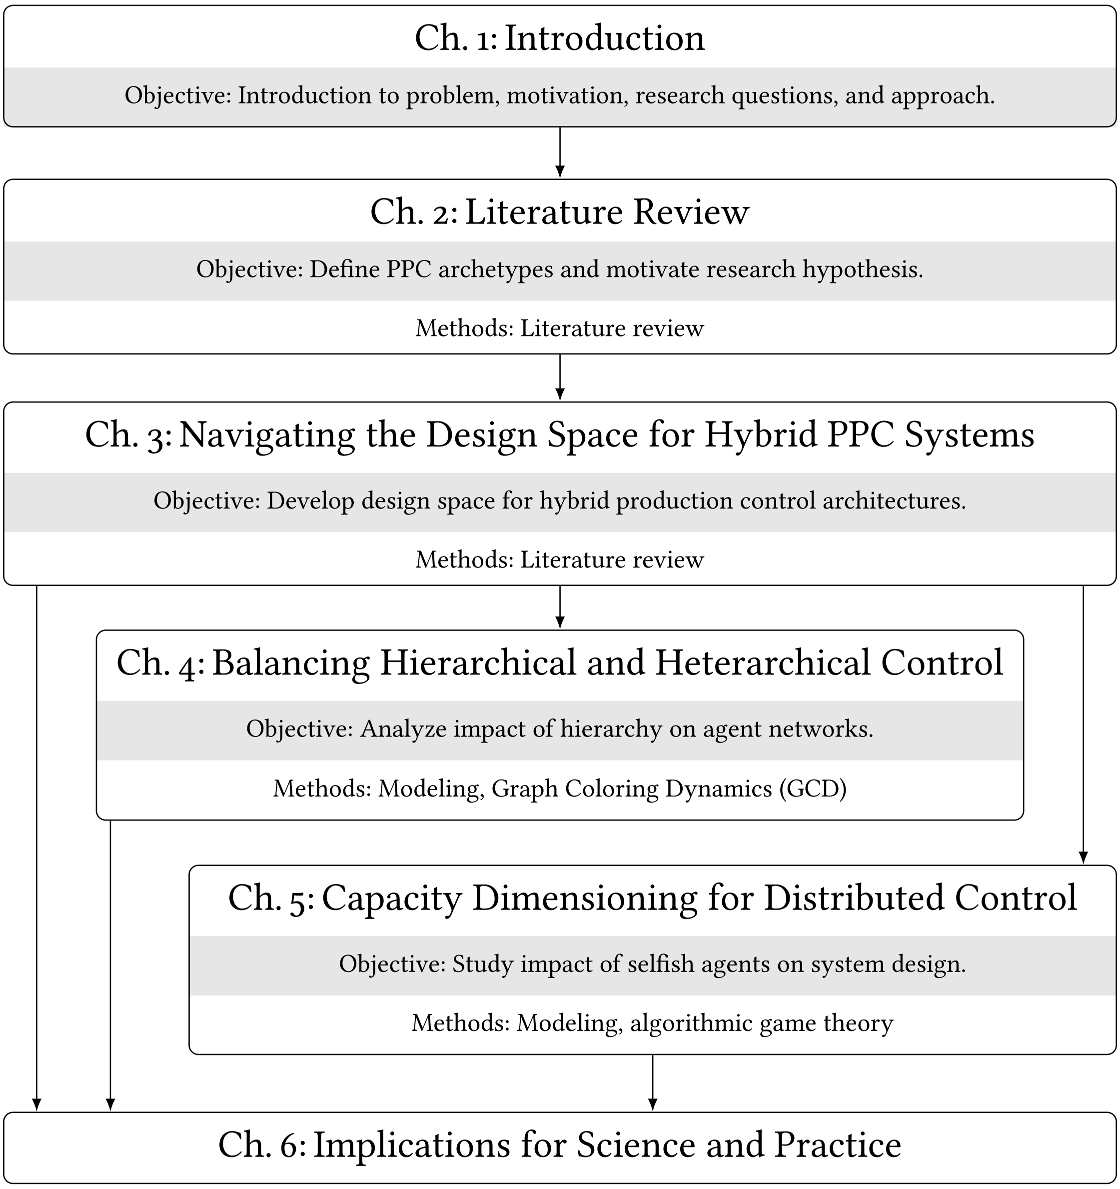

Figure 1.5: Graphical outline of the thesis, describing the interrelationships between the chapters and the addressed research questions (Rq.s) in each chapter.

Chapter Two

Literature Review: The Control of Production Systems

> *"Of primary importance is the fundamental question concerning the choice of control architecture: i.e. is it possible to determine whether a specific control architecture is appropriate for solving a given manufacturing system control problem?"*
>
> Rogers and Brennan (1997, p. 881)

To understand better the differences between hierarchical and distributed control, this chapter first introduces the system control problem in general and in the domain of production systems in particular (Section 2.1). Subsequently, the chapter will introduce in more detail hierarchical and distributed PPC as two architectural styles, each having its own strengths and weaknesses (Section 2.2). The resulting hypothesis that a combination of traits from both paradigms may show superior performance, already introduced in Section 1.2, is then further elaborated and substantiated in Section 2.3.

The chapter lays the foundation of exploring the design space *between* both architectural styles in more detail in the next chapter.

2.1 System Control

2.1.1 Fundamentals

Following Klir (1991, Ch. 2), this thesis defines as a system S as the tuple

$$S = (T, R)$$

of the "things" (or components) that make up the system T and the relations R among the elements of T.

Any system S endowed with solving a task is called a "goal-seeking system" and can naturally be represented as a "control loop", in which the system is decomposed into a controller and a system plant (Mönch 2005, Ch. 2.1; Vrabič and Butala 2012). The controlled system and the controller communicate through exchange of information, with sensor information flowing from the controlled system to the controller and actuator signals flowing in reverse. Baker (1998) specifies this concept for the realm of production systems in which the system to be controlled is the manufacturing plant (comprising of physical entities such as resources and products) and the manufacturing system controller is the PPC system (c.f. Fig. 2.1). For the controller's ability to render decisions based on information obtained from the plant (and its comparison with a goal state), we will also refer to entities that take control functions as decision making entities. The combination of the system plant and the controller constitutes the engineered and controlled system, which is set up and maintained to fulfill the formal cause of production systems: to transform inflowing materials into a transformed material outflux to meet customer demand (Hopp and Spearman 2008, Ch. 6.2.2).

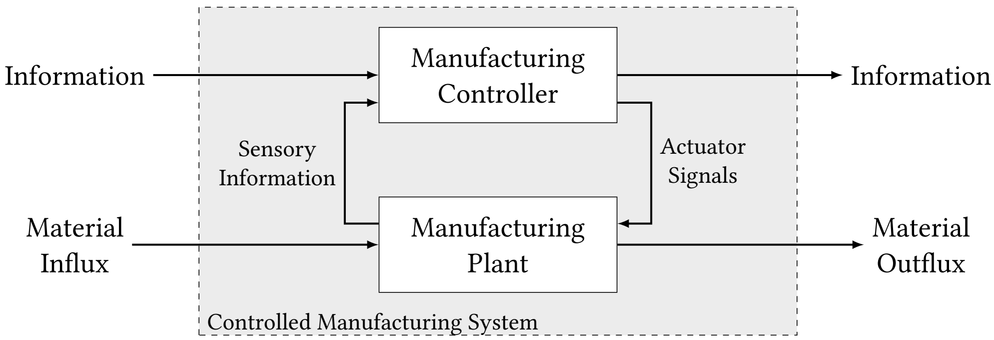

Figure 2.1: Factory Control Model. Adapted from Baker (1998), Sallez et al. (2010), and Trentesaux (2009).

Most real-world scenarios will give decision-making entities not a single, clear-cut, extracted decision problem, which can adequately be addressed by *problem solving* alone. Rather, the decision-making entities will find themselves confronted with "unstructured states of confusion" (or "messes") (Ackoff 1974, p. 5) of multiple interrelated decision-problems. The anticipatory, ex-ante process of considering the best set of decisions to reach a set of goals in the face of such "messes" is generally called *planning*. The process of detecting and responding to deviations during the plan *implementation* is generally referred to as *control* (c.f. Ackoff 1970, 1974, for both statements). In responding to new and unexpected information, control can be understood as "the process of scheduling the activations of information sources, both external (e.g. acquiring new input) and internal (e.g. invoking rules or updating beliefs)" (Pearl 1988, p. 318).

Since the control of larger and more complex systems easily grows beyond manageability, engineers usually apply *problem decomposition* to reduce and manage complexity (Schneeweiss 2003a, Ch. 1; Bakule 2008). Decomposition describes the splitting up of a given decision problem into smaller sub-problems for sequential or parallel execution (distribution in time and/or space). From the point of view of decision science, the solution approach is now an instance of a Distributed Decision Making (DDM) problem (Schneeweiss 2003b). Problem decomposition can be either *functional* (along the lines of the control functions to be implemented) (c.f. Passino 2005; Shen 2002) or *physical* (along the lines of the physical subsystems) (Shen 2002; Shen et al. 2006b; Tharumarajah 2001). The form of problem decomposition will be identified as a key discriminator between hierarchical and distributed perceptions of PPC and a design decision through which the information horizon of decision-making entities — and hence their ability to render globally "optimal" decisions — is impacted (c.f. Section 3.1.2).

2.1.2 Populations of (Selfish) Agents as Controllers

Decision-making entities may, and in most real-world manufacturing systems generally will, include both humans and computerized systems. Over the course of this thesis, a particular type of decision-making entity will repeatedly attract attention: *agents*. Being commonly defined as a computational entity that can perceive its environment, take decisions, and interact with other agents and its environment in order to achieve goals (i.e. is goal-seeking)[14], the agent concept provides the set of in- and outputs as well as the decision-making capacity required for controllers in the context of Fig. 2.1.

Where controllers are perceived as agents, the interaction between multiple controllers establishes an instance of a Multi Agent System (MAS) (Carabelea et al. 2004; Steegmans et al. 2004; Tokoro 1996). While MAS is a commonly used term in computer science, other sciences sometimes refer to the same concept as an Agent-Based Model (ABM) (Bonabeau 2002; Niazi and Hussain 2011; Holland 1998, Ch. 6). If the agent population of a MAS comprises of agents of identical capabilities, the MAS is called *homogeneous*, otherwise *heterogeneous* (Alshabi et al. 2007). According to Nilsson and Darley (2006), Agent-Based Models (ABMs) (or MASs) are particularly well suited as a modeling approach when (1) dynamic systems are distributed in time and space, (2) made up of many interacting autonomous parts, (3) exhibit several objectives and (conflicting) constraints, and (4) emergent phenomena can be expected. These are conditions that are all met for manufacturing systems (ibid.). In particularly they are well suited to model CAS (Castellani and Hafferty 2009, Ch. 5.2.6.3) and systems comprising of multiple, intelligent devices (as in the IoT) (Zhang et al. 2002).

From a software-engineering perspective, agents can be implemented in different ways. In the most simple way, agents simply encapsulate a piece of functionality or software (c.f. Shen et al. 2006b, for this and other implementation alternatives). In the literature on MAS, however, an additional characteristic is deemed essential (Jennings and Wooldridge 1998; Kalenka and Jennings 1999; Wooldridge and Jennings 1995), which will be assumed in this thesis when referring to agent-based approaches to PPC: *autonomy*, the ability and right of agents to render decisions independent of external entities (Falcone and Castelfranchi 2001; Jennings and Wooldridge 1998; Monostori et al. 2006; Windt and Hülsmann 2007). Agents will also be considered *goal-seeking*, i.e. "endowed with goals and that their behavior is guided by an internal (mental) representation of the effects" (Carabelea et al. 2004, p. 103). In conclusion, "an autonomous agent will be defined here as an agent that acts to achieve its own goals." (Conte and Castelfranchi

[14] Jennings and Bussmann 2003; Jennings and Wooldridge 1998; Marzo Serugendo et al. 2004; Monostori et al. 2006; Shen et al. 2006b; Stone and Veloso 2000; Conte and Castelfranchi 1995, Introduction.

1995, p. 46). Such Self-Interested Multi Agent Systems (SMAS), comprising of *individually motivated* agents, have become the focus of various subfields of Artifical Intelligence (AI) research since the mid 1980s (Conte and Castelfranchi 1995, Ch. 3; Green et al. 1997) (c.f. Section 2.2.2).

Several authors have pointed out the (perceived) underlying dichotomy between autonomy and the need to avoid chaos, to meet global constraints, to resolve agent interdependencies, to deal with distributed information and resources, and to eventually attain efficiency at a system level (c.f. Alshabi et al. 2007; Bar-Yam 2002; Doran et al. 1997; Falcone and Castelfranchi 2001; Váncza 2014). It remains to be argued then, to why autonomous (i.e. fundamentally selfish) agents should be assumed in a situation where the overall system behavior is the yard stick for the performance of a control architecture.

With designers of PPC systems being interested in high target achievement (c.f. Section 2.1.3), *not* in the implementation of particular design ideas, agent autonomy possesses no value in itself. There are, nevertheless, good reasons to investigate and understand the behavior of autonomous selfish agents. In (strictly) distributed control systems, cooperation is hard to attain as the existence of a common, shared goal to work towards, is required for cooperative behavior (Doran et al. 1997; Shenker 1995). Maintaining such common goal for the MAS population, however, would require a central hierarchically superior entity, which would make MAS-based control systems unfeasible in situations where the provision of such a central coordinator is unrealistic or impossible (Conte and Castelfranchi 1995, Ch. 3). Conte and Castelfranchi (ibid., p. 164) also remind us that in the beginning, "distributed problem solving has been characterized by pre-compiled cooperation, and MAS relied essentially upon the benevolence assumption" (c.f. also Brainov 1996; Ossowski and García-Serrano 1999; Tokoro 1996). Thus, the goal and the agents' behavior necessary to achieve it was coded into agents' behavior. Such an approach seems feasible where the system size is manageable, emergent behavior is predictable, and all agents are created by the same designer who can guarantee the assumption. In the dynamic value-creation networks foreseen in Section 1.1 however, it can no longer be assumed that all agents were designed for cooperation (Green et al. 1997) and selfish behavior may be required to survive in such open system (Tokoro 1996).

In defending a competition-based MAS implementations, one may also point to conceptual considerations: cooperative solutions are often vulnerable to selfish users (Shenker 1995; Başar and Olsder 1998, Ch. 4.7).[15] This is ever more problematic where coordination is to be attained across company borders (Sandholm 2000). As Ahrens (1996) points out, a competitive interaction may also seem more "natural" for people when vying for scarce resources. Selfish agents, therefore, implement the notion that the interplay of supply and demand observed at the macro-economic level (markets, for example, are highly self-organized, c.f. Ahrens 1996; Papadimitriou and Valiant 2010) can similarly be found within a company where customer orders compete for sparse production capacity (Ahrens 1996, 1998).[16]

[15] An example for a cheating-prone socially optimal solution for the routing games discussed in e.g. in Section 3.3.1 and Chapter 5.

[16] This idea of a "market within an enterprise" dates back to Coase and his milestone work "The Nature of the Firm" (Coase 1937).

Nwana et al. (1997)	Schneeweiss (2003a, Ch. 1)	Malone and Crowston (1994)[17]	Description
	data exchange		Agents exchange information about internal variables.
negotiation			Agents engage in a negotiation process, seeking to find mutually acceptable solution.
contracting	(reactive) negotiation	mutual adjustment	Market-like self-organization, with each decision-making entity adjusting to manage interdependencies.
multi-agent planning		direct supervision	A coherent plan is assembled from all decision-making entities and checked for coherence, conflicts.
organizational structuring	planning	standardization	A-priori structure among agents is given. Agents may instruct and anticipate the behavior of others.

Table 2.1: Forms of agent coordination, sorted by "increasing sophistication of communication and integration" (in the sense of Schneeweiss 2003a, p. 5)

[17] Aggregation of multiple works from organization science literature.

With cooperation not explicitly encoded in the agent behavior, agents need to engage in a process "in order to ensure their community acts in a coherent manner" (Nwana et al. 1997, p. 79). This process is called *coordination* (c.f. also Jennings 1993; Schneeweiss 2003b). The design of coordination processes is a pervasive problem in diverse areas such as social science, organization theory, anthropology, political, and computer science (Malone and Crowston 1994; Nwana et al. 1997). There are several classifications of coordination techniques in the literature: Nwana et al. (1997) name organizational structuring, contracting, multi-agent planning, and negotiation. Schneeweiss (2003a, Ch. 1) points to four consecutive forms of coordination that can be ordered by "increasing sophistication of communication and integration" (ibid., p. 5). The techniques are compared in Table 2.1, maintaining the ordering of Schneeweiss (ibid.). It is generally accepted that achieving coordination in "open" settings in the absence of pre-compiled cooperation is harder to achieve (Ossowski and Omicini 2002). Understanding the different means by which coordination may be attained helps to conceptualize and discriminate the existing paradigms of PPC (Section 2.2). As hierarchical and distributed PPC will be found to sit at opposite ends of the spectrum, Table 2.1 can also indicate which coordination techniques could be applied in hybrid PPC systems.

2.1.3 The Production Planning & Control Problem

The repetitive tasks involved in the management of the value-creation processes of a company (Becker 2012, Ch. 2.1.3; Gershwin et al. 1986; Hopp and Spearman 2008, Ch. 13; Schuh 2006) are collectively known as Production Planning & Control (PPC), although the exact definitions and delineations of these terms vary between authors and research streams (Gudehus and Kotzab 2009, Ch. 2).

PPC tasks span over multiple time horizons — from machine level setup and staffing decisions (with the shortest time horizon) to factory and production network planning and design at the highest time scale (Bahl et al. 1987; Hopp and Spearman 2008, Ch. 13.2.1). Exemplary tasks at the different time scales — referred to as the strategic, tactical, and operational level, as common in planning and control discussions (Miller 2002, Ch. 1.1) — are given in Table 2.2. Notably, the different time scales form a control hierarchy in which the decisions of previous steps influence subsequent decisions at any given point.

Long-term strategic PPC decisions often concern or support *design* decisions where decisions about the number, location, and equipment of production facilities (Bahl et al. 1987) are made to meet anticipated changes in the product mix and volume. Decisions on *production capacity* made at this location significantly impact on the company's bottom line and all subsequent decisions (Hopp and Spearman 2008, Ch. 18.1).

For Gudehus and Kotzab (2009, Ch. 2), the key characteristic of *production planning* is that it generally deals with inaccurate information about demands anticipated for the future. The majority of planning processes within PPC have a medium range (Hopp and Spearman 2008, Ch. 3.2.3) and are generally concerned with translating actual or forecasted customer demand in to production orders and determining release dates for them (c.f. e.g. ibid., Ch. 13.2). *Production control*, on the other hand, (in line with

Time Horizon	Planning Horizon	Decisions
Strategic	years to decades	capacity decisions facility locations
Tactical	weeks to years	work scheduling staffing assignments preventive maintenance purchasing decisions
Operative	hours to weeks	material flow control worker assignment machine setup decisions quality compliance decisions

Table 2.2: Production system management decisions by planning horizon. Table adapted from Bahl et al. (1987) and Hopp and Spearman (2008, Ch. 3.2.1), c.f. also Miller (2002, Ch. 1.1)

Section 2.1.1), considers the short-term steering of released orders (Gudehus and Kotzab 2009, Ch. 2; Hopp and Spearman 2008, Ch. 3.2.4).

The *scheduling* task is situated between (classical) production planning and production control (Gudehus and Kotzab 2009, Ch. 2; Becker 2012, Ch. 2.1.4). It is performed ex ante, yet only considers *actual* production orders (Gudehus and Kotzab 2009, Ch. 2; Hopp and Spearman 2008, Ch. 15). During scheduling, tasks (process steps necessary to complete jobs) are allocated to resources (Pinedo 2008, Ch. 1). Since resources can usually only perform one task at a time, manufacturing systems are commonly modeled as Resource Allocation Systems (RASs) (Reveliotis 2005) where "a finite set of reusable resources [...] are exclusively allocated to a number of concurrently executing processes for the sequential execution of their various processing stages" (ibid., pp. 12 f.). Hence, scheduling commonly entails solving (in parallel or sequence) a *sequencing* and an *allocation* problem. Where such flexibility is given, the allocation problem determines the allocation of tasks to one of multiple possible resources. The sequencing problem determines the sequence in which tasks are performed on a resource (c.f. also Fabiunke and Kock 2000). Finding the optimal scheduling (w.r.t. some performance criterion) is generally perceived as an optimization problem. The analytical intractability of many, even simple scheduling problems (Garey et al. 1976; Pinedo 2008, Appendix E; T'kindt and Billaut 2006, Ch. 4.5) means that heuristics are frequently applied to solve scheduling problems in practice (c.f. Section 2.2.1).

As they shape the value-creation process of companies, decisions on the structures and processes set out to complete PPC tasks, have profound impact on the attainment of the company's entrepreneurial targets (high quality, low cost, high delivery performance, and high flexibility (Schönsleben 2012, Ch. 1.3.1; Chryssolouris 2006, Ch. 1.3)). PPC systems are hence goal-seeking (c.f. also Mönch 2005) in such a way that they are aimed at the effective connection of markets with products and production (Olhager and Wikner

2000) and the effective allocation of current orders to resources in a manufacturing company (Enarsson 2006; Gudehus and Kotzab 2009, Ch. 20.1). Likewise, the "optimal" PPC system can (if at all) only be determined, *given* the production environment and entrepreneurial goals of a given company (Rogers and Brennan 1997). As a concretization of the above-mentioned entrepreneurial goals, production performance (and hence the quality of PPC systems) are frequently measured against the (concurrent) achievement of "logistical targets", namely (1) low levels of inventory, (2) short throughput times, (3) high capacity utilization, and (4) high due date reliability (Wiendahl 1997, Ch. 5.1; Nyhuis and Wiendahl 2009; Hopp and Spearman 2008, Ch. 6.3.2).

2.2 Architectural Styles to Production Planning & Control

With the introduction of the PPC problem and relevant terminology concerning problem decomposition and agent coordination, the thesis will now introduce hierarchical and distributed PPC approaches (in Sections 2.2.1 and 2.2.2 respectively). They represent *Architectural Styles* (c.f. Bass et al. 2003, Ch. 5.9 for a definition in the domain of software engineering), i.e. a set of features and rules, which an accordingly designed PPC system should exhibit/abide by. While the control system of any given real-world production system (if explicit and implicit PPC decisions taken by both humans and computerized decision-making entities are combined) will be unique, architectural styles refer to a combination of features[18] that are commonly found and discussed in conjunction and that combine in a logical, coherent fashion (ibid., Ch. 5.9).

2.2.1 The Hierarchical Approach

Across application domains, hierarchical control architectures are considered a convenient and popular choice for the control of complex system plants. They apply a functional decomposition of the problem into a set of sub-problems, each of which takes a subset of control decisions (Passino 2005, Ch. 1.4). From a decision making standpoint, hierarchical planning is a particular form of succession planning in which planning decisions are taken sequentially and decisions are fed-forward; however, decision-making entities in hierarchical control settings are (ideally) aware of subsequent planning steps and anticipate the consequences of their decisions upon them (Klein and Scholl 2011, Ch. 5.3), thereby creating a *leader–follower* relationship between controllers — comparable to *Stackelberg-Games* in game theory (Schneeweiss 2003a, Ch. 1).

The application of hierarchical control architectures to PPC tasks is known as Hierarchical Production Planning (HPP). This section will introduce HPP as an architectural style for PPC systems in which the control task is functionally decomposed and sub-controllers are coordinated through (hierarchical) planning where higher-level decisions form the input and constraints for lower-level decision-making.

[18] E.g. with respect to control problem decomposition and controller coordination.

Hierarchical Production Planning

The distinctly different time scales at which decisions on manufacturing design, planning, and control are taken (c.f. Section 2.1.3) and the impact of longer-term decision on short-term decisions (e.g. capacity investment decisions frame the possibility space for scheduling decisions) give rise to a natural perception of "hierarchy" within the tasks of PPC (Gershwin et al. 1986; Van Brussel et al. 1998) and have inspired researchers and practitioners in the field to develop various flavors of HPP.

In accordance with the above-mentioned definitions, the adjective "hierarchical" in HPP refers to a "decision-time hierarchy" (Schneeweiss 2003a, Ch. 1.1) where successive planning steps are implemented by a likewise hierarchical series of models (Passino 2005, Ch. 6.3.5): Hax and Meal, p. 3 explain such planning hierarchy as "each set of decisions at an aggregate level providing constraints within which more detailed decisions must be made".

HPP systems hence comprise of *multiple* decision-making entities and are hence identified as Distributed Decision Making (DDM) systems e.g. by Schneeweiss (2003a, Ch. 1.1). They are generally meant however when researchers in the domain of PPC these days talk about "centralized" production control, arguably because planning activities are usually performed by a single computer. They should not be confused with what Dilts et al. (1991) calls *monolithic* PPC architectures, where all PPC functions are collectively and simultaneously solved. These earliest approaches to computer-based PPC had to succumb long ago to increasing complexity in manufacturing operations (Dilts et al. 1991; McKay 2011), as they lacked, in particular, fault-tolerance, variability of response times, and it was generally difficult to modify the control system in the event of changes to the plant (Dilts et al. 1991; Mařík and McFarlane 2005). They bear no practical relevance in todays manufacturing environments (Dilts et al. 1991) and are thus not reviewed here.

As compared to monolithic PPC systems, Dilts et al. (1991), Graves (2011), and Miller (2002, Ch. 3.1) refer to the following main advantages of hierarchical PPC architectures: the problem decomposition lowers the planning complexity, reducing also the demand for input data at each stage. Redundancies and software development effort can be reduced and uncertainty can be addressed within a defined planning framework. The separation of decision problems along time-scales also makes it easy to align the sub-problems with the hierarchy of decision-makers (c.f. also Section 2.1.3) and to add or remove decision layers when deemed necessary.

Where sub-controllers take decisions successively, DDMs exhibit *weak information asymmetry* — i.e., an asymmetry in available information resulting from the limited amount of information typically available in early planning steps (Schneeweiss 2003a, Ch. 1.1). This asymmetry could (in principle) be removed once complete information is available; it is, therefore, weaker as compared to situations where multiple decision-making entities are active simultaneously (Schneeweiss 2003a, Ch. 1.1; Klein and Scholl 2011, Ch. 5.3), as will be the case in *distributed* PPC architectures to be introduced in Section 2.2.2. In terms of information asymmetry, HPP systems can thus be located between monolithic PPC systems, which exhibit no information asymmetry as they have only one decision-making entity, on the one hand, and distributed PPC systems, where multiple decision-making entities may take decisions simultaneously, on the other hand.

Historical Development of Hierarchical Production Planning & Control

The first widely used hierarchical PPC concept designed for computer usage was suggested by Hax and Meal (1973). Hax and Meal proposed to decompose the PPC problem functionally into four sub-problems to be executed sequentially, with one planning function building upon the results of its predecessors. Not only would this decomposition reduce planning complexity, but would allow for some planning functions to be solved at an aggregated level.[19] Aggregation (be it over time, products, processes, or people) had already previously been a staple of PPC models (ibid.) and is still an essential tool for complexity reduction in today's PPC systems (Hopp and Spearman 2008, Ch. 13.2.2; Klein and Scholl 2011, Ch. 5.3.5.2).

The basic ideas of functional decomposition and sequential decision making (partially) for aggregated planning objects, along with the advent of computers in the manufacturing environment, paved the way for the rise of computerized PPC in the 1970s, which made it possible for the first time to replace then existing approaches based on statistical order points and lot size calculations (Hopp and Spearman 2004; McKay 2011). Since then, HPP has repeatedly been integrated into ever more overarching business software products:

- MRP was introduced by Orlicky (1975) as a method to calculate dependent net demand[20].
- MRP II extended the MRP functionality to include forecasting, scheduling, and production control tasks.[21]
- Computer-Integrated Manufacturing (CIM) systems were developed since the 1970s and 80s to integrate the various IT systems related to the design, configuration, and realization of products and production system (Caridi and Sianesi 2000; Gunasekaran and Ngai 2012; Jones and McLean 1986).

All these evolutions not only maintained the hierarchical top-down planning approach (Bongaerts et al. 2000; Caridi and Sianesi 2000; Giret and Trentesaux 2015) but advanced the idea to give decision authority above an ever larger set of business decisions, to a single, centralized decision-making entity.

[19] By assigning products to families and types and run forecasting as well as initial scheduling problems on aggregated levels, before considering individual products again.
[20] Demand for sub-assemblies, parts, and individual pieces, not covered by inventory or current production orders.
[21] Hopp and Spearman 2008, Ch. 13.2.2; Chryssolouris 2006, Ch. 6.2; Herrmann 2011, Ch. 4.

Problems of Hierarchical Production Planning

In recent years, HPP systems have increasingly been criticized for not meeting the production control challenges posed by modern production environments, and their ability to cope with *future* challenges (as outlined in Section Section 1.1) is generally doubted. This section will recapitulate the most common criticism made about the previously introduced HPP implementations to better understand the renewed interest in distributed PPC.

Existing HPP approaches are commonly criticized for their rigidity and inability to adapt to changes in the production plant: Bongaerts et al. (2000, p. 125) complain that "Because of the static and deterministic nature of hierarchical control architectures, it is difficult to modify the system and to incorporate unforeseen changes into the system". Van Brussel et al. (1998) note that to incorporate plant changes in HPP systems, the entire system has to be shutdown and higher-order data structures[22] need to be re-built, e.g. after modifications to a resource. Even where technically feasible, centralized control approaches then seem ill-advised for Mařík and McFarlane (2005) in situations where frequent changes to the system plant have to be expected (c.f. also Bongaerts et al. 2000). This negative assessment also extends to CIM as an extension of MRP/MRP II today: the high hopes in integrating systems along the product life-cycle were not fulfilled (Ahrens 1996), to say the least: Ahrens (1998, p. 174) concludes that efforts to introduce CIM "ended in disaster", and Moldaschl and Weber (1998, p. 372) lament that the 1980s were a "'lost decade' from the viewpoint of organizational innovation" due to the flawed experiments with CIM. The reasons presented for this resounding failure are usually associated with an overly rigid system structure (c.f. Monostori et al. 2015)

HPP systems also perform poorly where the future is subject to uncertainty. While designing HPP systems to deal with uncertainty has been subject to wide academic interest (see Mula et al. 2006, for a review), different observations can be made in practice where (acc. to Graves 2011; Lindau and Lumsden 1995) most PPC systems do not explicitly account for uncertainty (i.e., there is no anticipation) by the planning stage and the implementation stage is assumed to be deterministic. While the notion of "hierarchical planning" (c.f. Section 2.2.1) was originally developed in the context of MRP and explicitly calls for reactive planning functions that incorporate the possibility for disturbances at lower-tier execution steps in the decision-making process (Klein and Scholl 2011, Ch. 5.3.3), most HPP systems have a "nonreactive anticipation function" where feedback from lower-level controllers is not included in higher-level decision making (Fleischmann and Meyr 2003). The interaction between sub-controllers in real-world HPP system is hence generally minted by the existence of *master-slave* or *command-response* relationships between entities (Bongaerts et al. 2000; Brennan 2000; Trentesaux 2009), more akin to classical, non-reactive succession planning (Klein and Scholl 2011, Ch. 5.3).

This rigid, successive planning approach leads to systems with little flexibility to react swiftly to unexpected changes at the shop-floor level. For Van Brussel et al. (1998, p. 262), this observation is directly linked to the idea of HPP as "hierarchical control architectures

[22] such as work-plans, product routings, etc.

almost automatically imply a top-down development methodology, which again introduces additional constraints into the solution". Similarly, Conte and Castelfranchi (1995, p. 44) note that systems "which are granted little autonomous decision-making, unfailingly but blindly execute their plans and tasks [...] are unable to react to unexpected and important events". This observation is also shared by Hatvany (1985, p. 103), who concludes that "highly centralized and hierarchically ordered systems tend to be rigid, constrained by their very formalism to follow predetermined courses of action".

Even where sub-ordinate planning steps are granted (limited) autonomy, hierarchies as a structural property in communication networks seem ill-equipped for uncertainty-plagued environments: Helbing et al. (2006a) find hierarchies to be the optimal architecture to spread information within a system (c.f. also Section 3.3.2), but highly susceptible to noise (i.e. variability) and failure of communication entities. Similarly, Malone and Smith (1988) observe that, while hierarchies can minimize coordination cost, they are vulnerable to failures of decision-making entities. For PPC, this resentment is shared e.g. by Mařík and McFarlane (2005), who see the dependence of hierarchical PPC systems upon the availability of each decision-making entity as a crucial disadvantage, while Duffie (1990) concludes that fault tolerance can only be introduced into hierarchical control systems at the expense of extra complexity.

Of course, the increase in complexity and uncertainty in production environments (c.f. also Section 1.1) is met with substantial improvements to the speed with which centralized PPC systems can analyze and solve PPC problems: Koch et al. (2011) report a speed-up of 100 million times over a 20 year timespan to 2011 in the solution speed for Mixed-Integer Linear Programs (MILPs) (such as scheduling problems), resulting from improvements in both hardware and algorithms. There are, however, reasons to doubt that even these advancements in computational power will suffice to perform in future manufacturing environments: According to Bakule (2008), the increase in the task-complexity of controlling a system with increasing size and complexity outgrows advances in computing power and memory in modern computers because of increasing (1) dimensionality (induced by the increasing number of decision variables), (2) information structure constraints, (3) uncertainty, and (4) delays (between control decision and the effect on the controlled system).

Increasing dimensionality arises from both the increasing flexibility potential within production systems and the quest of centralized computer-based planning systems to cover larger parts of the decision process in companies: Deshmukh et al. (1998) show mathematically that the shift from dedicated to multi-purpose, flexible production resources leads to an increasing complexity during production planning. Both Bertrand and Muntslag (1993) and Stevenson et al. (2005) note that MRP II performs poorly in environments with non-standard products and flexible product routings. Duffie (1990, p. 167) finds that "the complexity of computer-integrated manufacturing systems with hierarchical architectures grows rapidly with size, resulting in accompanying high costs of development, installation, operation, maintenance, and modification".

Delays in HPP arise from the two-fold communication process in hierarchical coordination processes that have to (1) inform the central decision-making instance and (2) re-distribute and implement the decision (Guetzkow and Simon 1955). Such an approach is not only time-consuming and hence bound to operate on outdated information, but also requires the centralized PPC system to render appropriate reactive measures in a short time. Given the uncertainty about the system evolution, Monostori et al. (2006, p. 698) note: "No matter how, it is next to impossible to be prepared with preprogrammed, top-down responses to abrupt changes and to complete real-time computations on sophisticated decision models before the results are invalidated". In the absence of reactive planning capability, frequent re-planning has to be performed. This does not only take time, reducing the reactiveness to unforeseen events like machine breakdowns, but also creates additional, artificial variability, system nervousness, and frustration ("schedule churn") (c.f. also Hopp and Spearman 2008, Ch. 3.1.9).

2.2.2 The Distributed Approach

This section will introduce distributed production control as the second architectural style to be considered in this research. Distributed PPC approaches are designed following a constructionist view through a bottom-up approach, generally building on a CAS perception of the manufacturing system. They will be found to combine mostly physical decomposition with coordination through various forms, ranging from data exchange to negotiation and multi-agent planning. Decision-making agents in (strictly) distributed PPC architectures render decisions based on local information and in parallel to others (observed as the essential discriminating factors of distributed PPC in Okubo et al. 2000). When compared to HPP systems, the simultaneous planning and decision-making by several decision-making entities implies that distributed PPC systems exhibit "strict information asymmetry" that cannot be removed over time Schneeweiss (2003a, Ch. 1.1).

As this section will show, the devolution of control authority has a distinct pedigree in the context of manufacturing systems, but is has gained renewed momentum with recent advances in information and communication technology. Since then (roughly the mid 1980s), the (expected) advantages of distributing control from a central planning authority[23], have remained largely identical. They usually include:[24]

- improved computational efficiency due to asynchronous/parallel computation
- increased robustness, since control does not hinge on single controller
- a more open system architecture
- improved scalability
- increased flexibility
- improved re-usability of solutions
- lower cost
- simpler programming and adaptation of control structures to structural changes

[23] A "planning office" as envisioned by Taylor (Herrmann 2006).

[24] c.f. e.g. Argoneto et al. 2008, Ch. 4.7; Baker 1998; Christensen 1994; Duffie 1990; Lesser and Corkill 1981; Mařík and Lažanský 2007; Mařík and McFarlane 2005; Monostori et al. 2015; Ouelhadj and Petrovic 2009; Sandholm 2000; Shen 2002; Shen et al. 2006a; Stone and Veloso 2000.

- better performance (especially under real-world environmental conditions)

Historical Development

This subsection seeks to trace the idea of distributing control authority in production environments over the last roughly 80 years. The review serves as a backdrop to the most recent pushes to distribute production control as part of the "fourth industrial revolution", as discussed in Section 1.1.2. It will show that (1) the general idea of control devolution and distributed parallel decision-making by multiple decision-making entities is not a new idea, but rather has (re-)surfaced repeatedly in the discussion on PPC and (2) previous attempts toward distributed production control have been thwarted by inflated claims for superiority — often marketing-driven, rather than scientifically grounded — as well as a lack of understanding for the success factors of distributed control. Explicit consideration of the history of distributed PPC ideas, will allow the remainder of this thesis to draw upon a wider range of literature streams to substantiate the research hypothesis (Section 2.3).

Early Forms of Distribution within PPC

Section 1.1.2 has already shown the increased interest in distributed PPC in the context of CPS and the technological and organizational developments surrounding the vision of a fourth industrial revolution. Despite that, it is worth noting that the idea of control devolution and distributed decision-making in manufacturing environments, spawned not long after Taylor had proposed his vision of a clear distinction between planning and doing (Hopp and Spearman 2008, Ch. 1.5; Moldaschl 1998). Three approaches will be discussed here, where "the theory is put forward that efficiency can easily be increased by partially renouncing the exercise of control (i.e., by flattening hierarchies)" (Moldaschl and Weber 1998, p. 381): dispatching rules, group work, and PULL production. All of these, as will be shown in this and the following sections, have advanced the idea of autonomous decision making in theory and practice and forestalled some of the debates on the advantages and disadvantages of distributed control that this thesis will elaborate on.

The first approach are *dispatching rules*. These are simple intuition-based solution procedures for scheduling tasks, such as routing and sequencing, by ranking decision alternatives and choosing the best-ranked option (Chryssolouris 2006, Ch. 6.3.3). Research on dispatching rules has developed since the 1950 in parallel to research in the ex ante scheduling via OR models (Panwalkar and Iskander 1977). They represent the most common form of reactive (dynamic) scheduling in industry today (Uzsoy et al. 1994; Van Dyke Parunak 1996). Dispatching rules use only temporally and spatially local information to make scheduling decisions, mostly to determine the operation sequence at a machine (Blackstone et al. 1982; Panwalkar and Iskander 1977; Pickardt et al. 2013), but similar approaches exist to make allocation (routing) decisions to determine which machine should be used (Windt et al. 2010b) and when jobs should be released (Baker 1998). Given their local information horizon, they lend themselves to heterarchical implementation (ibid.), but may not be able to explore the full solution space (Blackstone et al. 1982). In fact, "the preponderance of agent research for manufacturing has developed agent architectures that merely implement dispatching rules. It is most common to dispatch the routing decision in these architectures, assuming sequencing can then be done at each resource" (Baker 1998, pp. 302 f.). In particular, a routing decision based on the queue lengths of the machine alternatives is commonly used (Tay and Ho 2008).

From El-Bouri and Shah (2006, pp. 342 f.), we learn that the "main disadvantage of dispatching rules is that they are problem-dependent, and no one dispatching rule typically dominates others for most, if not all, performance". The combination of dispatching rules (e.g. Pickardt et al. 2013; Tay and Ho 2008) and the selection of dispatching rules given particular manufacturing system and/or job characteristics (e.g. Braglia and Petroni 1999; El-Bouri and Shah 2006) has hence been one focus of research.

The theory of *group work* developed after the landmark 'Hawthorne experiments' by Mayo (1933) and Roethlisberger and Dickson (1939) allegedly[25] showed the importance of informal (bottom-up) organization between workers on system performance (Moldaschl and Weber 1998; Hopp and Spearman 2008, Ch. 1.6.2). The experiments have had major implications on several fields of study. Thomas et al. (2005b) argue that in management theory, the perception of a duality between control and autonomy in organization science dates back to these experiments (this field of research will be picked up again in Section 4.1.2). Moreover, the group work theory evolved. Through some intermediate steps[26], the theory of (semi-) autonomous group work picked up the idea of informal self-organization (Cummings 1978). They are today considered a suitable approach to deal with the increasing flexibility requirements and variability of modern and future production environments (Wiendahl et al. 2009, Ch. 4.2); they have also been discussed as a model of semi-heterarchical PPC (Rahimifard 2004).

Finally, *PULL production* is part of the "lean production" philosophy that developed in Japan (notably at Toyota) (Arlbjørn and Freytag 2013; Holweg 2007) after World War II, and rose to prominence in the Western world with the release of the book "*The Machine that changed the world*" (Womack et al. 1990). There, the adoption of lean was euphoric after the recent failure of CIM systems (Moldaschl and Weber 1998) (Section 2.2.1). Following Hopp and Spearman (2004, p. 142), a PULL production system is defined as one "that explicitly limits the amount of work in process that can be in the system", with replenishment orders released *if and only if* final products are requested and hence removed from the system. The distributed character of PULL production principles has repeatedly been noticed (Askin and Goldberg 2002, Ch. 12.1.2; Caridi and Sianesi 2000; Theuer 2012). Baker (1998, p. 315) notes that "Pull algorithms are essentially distributed, and thus lend themselves to implementation in a multi-agent heterarchy".

Despite the resounding success in manufacturing environments, few positive implementation examples have made it into the peer-reviewed literature (so stated by Arlbjørn and Freytag 2013) and low success rates and even eventual negative results have been

[25] Despite harsh criticism of the experiment execution and result interpretation (c.f. Moldaschl and Weber 1998), the experiment and the original authors' interpretation made it to textbook platitude.
[26] Most notably the work by the *Tavistock Institute on Human Relations*.

reported (Bhasin and Burcher 2006; Browning and Heath 2009). Given its philosophical streaks[27] and diversity among "lean" implementations, it is clearly beyond the scope of this section to fully investigate the reasons for these shortcomings. However, it is worthwhile to note that the applicability of PULL techniques (especially of the KANBAN flavor) seems constrained to situations with limited static and dynamic complexity. This argument will be examined more closely in Section 3.3.1, where the impact of plant design decisions on the performance of distributed production control approaches is discussed.

In the current hype around a fourth industrial revolution (c.f. Section 1.1.2), it should also be noted that both group work and PULL have suffered from inflated expectations. As with the Hawthorne experiments, the initial publication on PULL (Womack et al. 1990) was scant on implementation details (Bhasin and Burcher 2006; Hopp and Spearman 2004) and masterfully placed as an item of science marketing (Moldaschl and Weber 1998). In the context of PULL systems, this has led to double standards: Pettersen and Segerstedt (2009) lament that in the lay literature and internet publications the "good things are pull and the bad things are push and causes and effects are not separated" (ibid., p. 200). Bonney et al. (1999, p. 53) similarly deplore "if the performance of a pull system is poor then it may be suggested that this is because the fundamentals [...] are not being observed, whereas, if the performance of a push system is poor, then that is a consequence of it being a push system".

Agent-Based Production Control

With advances in computation and communication technology in the 1980s, computers no longer had to be confined to mainframes used for production planning tasks (Herrmann 2006), but they could deployed throughout the system plant, providing the basis to realize (and investigate) computer-based, distributed control systems in manufacturing contexts (c.f. Dilts et al. 1991; Gershwin et al. 1986; Hatvany 1985).

The first concrete ideas and experiments for distributed control in production settings were published from the mid 1980s (c.f. e.g. Duffie and Piper 1986, 1987; Hatvany 1985; Van Dyke Parunak 1987; and Dilts et al. 1991, for a more comprehensive review). Explicit reference to agent-based system design was first made by Shaw (e.g. in Shaw 1987). A more comprehensive research on the possibilities of agent-based control architectures started in the early 1990s (Leitão et al. 2013; McFarlane and Bussmann 2000). Since then, a wide array of agent application has been proposed for various functions and tasks related to PPC. Reviews are provided e.g. by Bousbia and Trentesaux (2002), Leitão (2009), Leitão et al. (2013), and Shen et al. (2006a,b). While initial attempts did not allow for even communication (the "lowest" level of coordination acc. to Table 2.1) between agents (Bongaerts et al. 2000; Dilts et al. 1991), it was quickly realized that "some global information must exist in a system" (Duffie and Prabhu 1994, p. 95) and the focus shifted towards minimizing the amount of global information necessary and abolishing master-slave relationships (Duffie and Prabhu 1994; Rahimifard 2004). Going beyond information exchange, many agent-based PPC systems introduced negotiation-based coordination approaches (Shen et al. 2006a). In particular, the *Contract Net Protocol*

[27] Bhasin and Burcher (2006), for example, argue the varying success of lean as a management concept may be attributed to missing embrace of these philosophical aspects.

by Smith (1980) has become widely used. It offers a high-level coordination protocol in distributed systems. Coordination is achieved through contracting where *contractor* agents make bids to *manager* agents to fulfill tasks that the managers have announced. The managers assess the bids and award the task to the bid deemed best (Nwana et al. 1997; Smith 1980). Following extensions have mainly sought to reduce the coordination effort induced by increasing system sizes by reducing the amount of communication between agents (Shen et al. 2006a). However, market-based approaches have also attracted attention (c.f. Lee et al. 2003; McDonnell et al. 1999)

Besides the MAS paradigm and bio-inspired approaches to distributed control, the third, closely related modeling approach that has been used to establish (semi-)heterarchical control systems is Holonic Manufacturing (Sallez et al. 2010; c.f. also Tharumarajah et al. 1996). HMSs developed from an international collaborative research effort (the Intelligent Manufacturing Systems (IMS) initiative) in 1992, in which a global network of researchers investigated and tested several possible, novel control approaches (Christensen 1994). Among the six investigated approaches was the concept of HMS (ibid.). It picks up the notion of *holons* defined by Koestler (1970), who observed that in many (particularly: living) systems a clear distinction between "the part" and "the whole" was not possible and that these systems can grow to create stable "intermediate forms" in sight of challenging, evolving environmental conditions (McFarlane and Bussmann 2003; Tharumarajah et al. 1996). Holons are then recursively defined as comprising multiple sub-holons (Babiceanu and Chen 2006). At each level of abstraction, holons are self-reliant (capable of surviving disturbances) and yet sub-ordinated to higher-level holons that ensure efficiency at the level of the whole (Botti and Giret 2008, Ch. 2.1). Holons can be composed into larger systems, called *holarchies*, which Leitão et al. (2013, p. 2361) define as "a hierarchy of self-regulating holons that are simultaneously autonomous wholes for their lower parts and dependent parts for higher control levels". The HMS concept set the starting point for a series of frameworks for agent-based distributed manufacturing control developed from the mid 1990s (ibid.). Important contributions include the Product Resource Order Staff Reference Architecture (PROSA) in Van Brussel et al. (1998) that has become the central reference point for most recent HMS-design frameworks and is still among the frameworks deemed best equipped for the already mentioned challenges faced by modern manufacturing systems (Giret and Trentesaux 2015). The progress has repeatedly been subject to review papers (Babiceanu and Chen 2006; Giret and Trentesaux 2015; McFarlane and Bussmann 2003) that allow to follow the development. Holonic thinking offered a good starting point for the discussion of agent-based production control approaches, since the concept of the holon and the concept of the agent share a range of similar or identical features (such as autonomy, cooperation, reactive, c.f. Botti and Giret 2008; Glanzer et al. 2001). The self-referencing recursive nature of systems that is appealed to here, is also an important property of CAS (Castellanos 2012, Ch. 2.4.3.6) and generally essential to systems theory (Klir 1991, Ch. 3).

Problems of Distributed Production Planning & Control

Despite the sustained interest in academic research, industrial applications of agent-based manufacturing control systems have so far been limited.[28] A number of obstacles to wider adoption have been identified, including the lack of interoperability between existing commercial computing systems and agent-based systems, insufficiently flexible production resources, unavailability of development tools and platforms, high investment cost, and the lack of industry standards (Dilts et al. 1991; Leitão 2009; Mařík and Lažanský 2007; McFarlane and Bussmann 2003; Shen et al. 2006a). Mařík and Lažanský (2007), Mařík and McFarlane (2005), and McFarlane and Bussmann (2000) also point to enduring impact of reductionist thinking on the education and training of production system engineers, making them less at ease (or outright unwilling) to design or support distributed agent-based control approaches. Gorecky et al. (2014) finally point to unsolved issues concerning the interaction between "intelligent" products and human operators in the production environment which hinder the adoption of novel PPC approaches.

Another major concern that has already significantly hampered the acceptance of distributed control architectures by industry is the inability to predict (or bound) the performance of manufacturing systems under distributed PPC (Mařík and Lažanský 2007; Mařík and McFarlane 2005; Trentesaux 2009; Van Brussel et al. 1998). De Carolis et al. (2017) aggregated feedback from researchers and industry experts, concluding that uncertainty about performance was the major concern with respect to the application of CPS in production operations. While it is deemed impossible to predict performance statistics at the individual product/order level, even at the aggregate level only average performance levels can be predicted (Van Brussel et al. 1998).

These problems can largely be traced back to the constructionist, bottom-up design approach. Jennings and Wooldridge say about MAS in general:

> *"The first major problem is that the overall system is unpredictable and nondeterministic: which agents will interact with which others in which ways to achieve what cannot be predicted in advance. Even worse, there is no guarantee that dependencies between the agents can be managed effectively, since the agents are autonomous and free to make their own decisions. [...] The second main disadvantage is that the behavior and properties of the overall system cannot be fixed at design time. While a specification of the behavior of an individual agent can be given, a corresponding specification of the system in its entirety cannot, since global behavior necessarily emerges at run time."*
>
> — Jennings and Wooldridge (1995, p. 364)

The (potential) lack of global performance has been attributed to the lack of a global perspective, the absence of an overall system controller (c.f. Jennings and Wooldridge 1998), and the very concept of leaving agents autonomy (c.f. Section 3.5.3). As Duffie and Piper (1987, p. 179) note: "the fundamental objective of maintaining local autonomy contradicts objectives of optimizing over-all system performance". Good performance is assumed to be in reach for strictly heterarchical systems, primarily where demand is

[28] Leitão 2009; Mařík and Lažanský 2007; Mařík and McFarlane 2005; Shen 2002; Trentesaux 2009; Van Dyke Parunak 1996; c.f. Leitão et al. 2013, for a timeline of examples.

stable and predictable, order throughput times are short and stable, standard products and standard procedures are requested by customers, and local scheduling rules can be harmonized (c.f. Gudehus and Kotzab 2009, Ch. 8.8.7), i.e. a stable system environment, already discussed in the context of PULL systems, can be reached. Jones and Saleh (1990) hence doubt that (purely) heterarchical control architectures can work in highly structured and unpredictable shop floor settings, stating that "although theoretical foundations for this approach do exist, they have not been shown to be practical for controlling the dynamic evolution of a stochastic system" (ibid., p. 62). Ouelhadj and Petrovic (2009) review comparative studies between centralized and distributed control architectures to find that distributed control architectures yield high performance especially in small systems (few agents). Where this is not given, even proponents of distributed control architectures, such as Mařík and McFarlane, p. 32, concede: "Almost universally, where a centralized solution can be simply implemented, maintained, and changed, it will surpass a distributed solution in terms of conventional performance. So, classical centralized solutions might be more efficient in many situations".

Striking analytical examples have been reported for the lack of performance of systems under distributed routing or sequencing decisions, including settings in which system-performance may *decrease* as additional capacity is added. Examples include the famous *Braess' Paradox*, a phenomenon from traffic research (c.f. Sections 3.3.1 and 5.3.4) and the *Graham Anomalies* (further discussed in Section 3.4) that describe performance drops when additional servers are added to parallel server systems processing a priority ordered list of tasks of different work contents. In the domain of queuing networks, re-entrant lines or other system structures exhibiting cyclic dependencies may be prone to becoming unstable system behavior when increasing capacity (Dai et al. 1999; Kumar and Seidman 1990). In particular, Kumar and Seidman (1990) showed that all priority rules are susceptible to instability in the presence of cyclical material flow networks. Baron et al. (2009) find a similar effect, called the *Capacity Allocation Paradoxon* in much simpler queuing networks with limited buffers, where adding additional capacity to a (originally stable) queuing system can result in instability. However, they also prove that this instability can be prevented through adjustments to the priority rules that can likewise be distributed to all machines (do not require inter-machine communication).

In addition, Van Brussel et al. (1998) observe that the performance of distributed PPC systems is heavily dependent on parameter tuning — especially parameters of the coordination mechanisms. In particular:

> *"1.) The global system performance, e.g., throughput, is very sensitive to the definition of the market rules, and to the fine tuning of the rules (e.g., relative importance of transport times). 2.) The control system cannot guarantee a minimum performance level in case the system goes outside the working scope for which the rules were tuned. Only an average overall performance can be predicted when the system is inside its nominal working scope. 3.) Prediction of the behaviour of individual orders is impossible. The flow time of an order highly depends on the nature and status of other orders in the system."*
>
> — Van Brussel et al. (1998, pp. 262 f.)

The resulting heavily path-dependent and potentially chaotic system behavior also impedes system analysis and "debugging", adding to the undesirable traits of distributed PPC systems (Monostori et al. 2015).

As with the overall performance, the lack of predictability is blamed on the absence of a central supervisory entity (Mařík and McFarlane 2005). In particular, graceful handling of unexpected situations cannot be guaranteed, making issues like safety, fault-tolerance, and deadlock avoidance major concerns in the development of distributed PPC systems beyond issues of performance (Trentesaux 2009).

2.3 Hypothesis: Mix of Hierarchical and Distributed Control Attains Best Performance

It has become apparent from the previous discussion that while "hierarchical and heterarchical models have limitations, they also have several desirable characteristics" (Heragu et al. 2002, p. 562). One might be tempted then to consider whether a *combination* of both hierarchical and distributed traits in control architectures may exhibit properties superior to either of the extremes. As this section will show, this hypothesis has indeed been pondered in numerous publications across many of the scientific domains already discussed in this thesis. Quantitative evidence for and structural insights into such optimal combination have, however, been largely missing.

2.3.1 A Broadly Assumed Hypothesis

In Organization Theory

As Thomas et al. (2005b) recall, the aforementioned Hawthorne experiments have laid the foundation for what has become known as the "Control vs. Autonomy Duality" in management research. The duality was expressed prominently in the discussion of multinational companies by March (1991). It has more recently been extended into a broader set of dualities that touch various aspects of control and autonomy (e.g. hierarchies vs. networks, standardizing vs. customizing, etc.) (Evans and Doz 1999; Pettigrew and Fenton 2000; Thomas et al. 2005a) as well as to the idea that a cyclical change between control and autonomy should be exercised Thomas et al. (2005b). It is worthwhile to mention at this point the definition of "duality" in organizational research that implies the existence of "opposing forces that must be balanced" (Janssens and Steyaert 1999, p. 122) as they are complementary and not contradictory (c.f. Evans and Doz 1999; Graetz and Smith 2008). In that, dualities are different from "Trade-Offs". Thinking in terms of such dualities has, in fact, become common-place in management theory (c.f. e.g. Evans and Doz 1999; Janssens and Steyaert 1999; Pettigrew and Fenton 2000). Vicari et al. (1996) develop in the context of control vs. autonomy duality three *organizational states* that describe the "relationship between survival and advancement activities" (ibid., p. 193). In the organizational state of *inertia*, the focus is on knowledge exploitation, stressing the need for order and stability. The rate of knowledge generation in such organizations is believed to be low due to the lack of experimentation. The state of *dissolution*, on the other hand, is reached when "contradictory forces and alternative courses of strategic action overbalance order in a bewildering manner" (ibid., p. 195). Between the two is the state of *extension*, where experimentation is permitted and the organization is held at the (previously mentioned) *edge of chaos*. To foster experimentation, "authoritative structures are substituted by self-organization" (ibid., p. 194). According to Vicari et al. (ibid.), this state allows to combine high levels of knowledge exploitation with potentially high levels of knowledge generation.

The academic consensus in organization theory today promotes such dual organizational forms that combine the controllability advantages of classical hierarchical forms of organization with the responsiveness of more distributed approaches (Graetz and Smith 2008; McKelvey 2004). Stacey (1993) builds upon the "Edge of Chaos" theory developed in computer science (Langton 1990) to suggest that balancing centralization and decentralization can maximize creativity and adaptiveness. More generally, research on the design of socio-technological systems assumes the existence of a "best match" between performance and control that makes the search for such fit worthwhile, though the search process for such optimum is not addressed directly (Moldaschl and Weber 1998). Also CLT, as a leadership framework for the management of CAS, "seeks to foster CAS dynamics while at the same time enabling control structures appropriate for coordinating formal organizations and producing outcomes appropriate to the vision and mission of the system" (Uhl-Bien et al. 2007, p. 304).

It is generally assumed, however, that the search for a "best match" between control and autonomy is hard, if not impossible (Thomas et al. 2005b). Pettigrew and Fenton (2000, p. 295) argue (summarizing March 1999, Ch. 1) that "defining an optimal mix of exploration and exploitation is difficult or impossible. It involves trade-offs across space, time, people, and levels in a system. In other words, the experience of dualities and their management is likely to be highly context sensitive."

In Early Forms of Distributed Production Control

Also in the context of PULL production, researchers have sought to address previously problems with the application of pure KANBAN systems by means of integrating traits of hierarchical PPC into the controller design. Gudehus and Kotzab (2009, p. 207) state that in situations where the conditions for purely local scheduling are not given, "a *combination of local and central scheduling* may achieve most of their advantages and avoid the disadvantages" (highlighting by Gudehus and Kotzab). Cochran and Kaylani (2008) report on research (and review literature) on "hybrid control systems" where PULL systems are either horizontally coupled with PUSH (i.e. centrally controlled) production systems or higher level planning functions are conducted in a centralized fashion (vertical integration). The idea to use the customer order decoupling point (CODP) as the separating point between PUSH and PULL production has become a popular design recommendation (c.f. Olhager and Östlund 1990).

The production control approach Constant Work-in-Process (ConWIP) was developed explicitly to address the aforementioned (Section 2.2.2) problems with KANBAN systems

(Spearman et al. 1990). In ConWIP systems, PULL is only applied at the front of the line, whereas machines within the line PUSH the product forward (Spearman and Zazanis 1992). While in KANBAN systems the circulating cards prescribe the product to be produced, in ConWIP the product to be produced with the next arriving card is determined by a "backlog list" maintained by the planning department and likely the result of some overlaying centralized production scheduling (Spearman et al. 1990). Ryan et al. (2000) show that with the added element of centralized planning, assumptions on the plant layout for KANBAN can be relaxed and ConWIP can control more complex manufacturing systems, even job shop environments. Puchkova et al. (2016) find that ConWIP provides better protection against disruptions and demand peaks then KANBAN, while keeping line costs lower then PUSH production, thereby establishing a good trade-off between the virtues of PUSH and PULL.

In Agent-Based Production Control

The hypothesis that some combination of distributed and hierarchical traits is necessary to avoid undesirable traits in distributed PPC systems has a long tradition within the research on agent-based distributed production control, dating back to its very beginnings (c.f. e.g. Hatvany 1985).

The effect of a combination of hierarchical and distributed control traits on performance has frequently been the subject of hypotheses. In related publications, Philipp et al. (2006), Philipp et al. (2007), and Windt et al. (2008a) suggest a curvilinear relationship between the degree of "autonomous" (i.e. distributed) control and the logistics target achievement and the maximum performance attained "in between". Windt et al. (2008a) propose to measure the degree of autonomous control through a catalog, without giving hints how to translate the answers into a numerical (percentage) value. Unlike the other two mentioned papers, Philipp et al. (2007) suggest that the effect of autonomous control on production system target achievement is modulated by the system complexity. Low-complexity production systems are believed to reach maximum performance with little distributed control. As the complexity increases, the aforementioned curvilinear relationship between the level of distributed control and the achieved performance (c.f. Fig. 1.4) re-emerges. A similar (two-dimensional) relationship is suggested by Zambrano Rey et al. (2014), who seek to operationalize the "degree of autonomous control" by referring to the complexity of the decision function applied at the agent level.

It is further believed that the time-horizon in which a control system has to render decisions (be it offline planning tasks or online control with near real-time requirements) influences which control architecture would perform the "best". The general aptitude of self-organization processes to deal with dynamic, fast-changing situations (De Wolf and Holvoet 2005) has led the logistics and PPC community to believe that distributed control systems are more prone towards short-term optimization and robustness in the face of variability and disturbances. On the other hand, hierarchical centralized control systems are believed to offer the prospect of long-term global optimization under the premise of relative stability in the network structure and performance requirements (c.f. e.g. Caridi and Sianesi 2000; Trentesaux 2009). It has repeatedly been suggested that there is a correlation between the desired control behavior and the optimal layout of the control

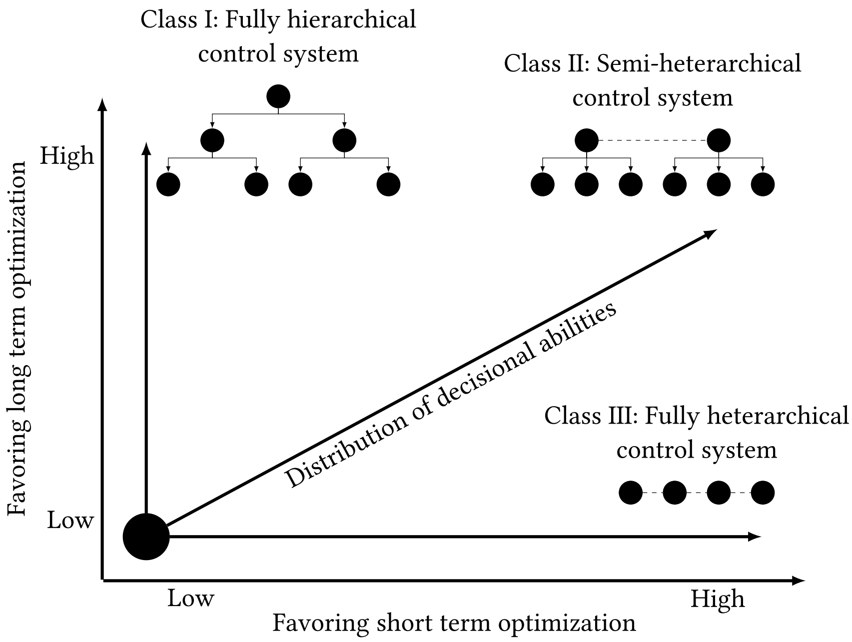

Figure 2.2: Degree of hierarchical coordination in the trade-off between short- and long-term optimization. Adapted from Trentesaux (2009). Every controller at the lowest hierarchy level is connected to a controlled subsystem. A monolithic control system (only one controller) would be on the lower left corner of the plane.

system (Mařík and McFarlane 2005; Rogers and Brennan 1997; Trentesaux 2009). To this end, Trentesaux (2009) extends the aforementioned dualism by arguing that the optimal control design is a function of the desired degree of long-term performance vs. short-term optimization and that strictly distributed, hybrid, and hierarchical control architectures could be placed on the efficient frontier, serving different design aims. His perception of the dualism is shown in Fig. 2.2. The semi-heterarchical control systems (Class II) are believed to show two advantages over purely distributed approaches (Class III): They avoid unpredictable behavior (impeding industrial adoption) and offer a viable transformation path from the current highly centralized control structures to more heterarchical ones. Without hypothesizing about a particular shape of the trade-off, a mix of hierarchical (top-down) and distributed (bottom-up) designs is also assumed to be optimal for PPC system design by Brennan and Norrie (2001), Giret and Trentesaux (2015), and Leitão and Restivo (2008).

The apparent disadvantages of both hierarchical and distributed control approaches also reflect on the strengths and weaknesses of the reductionist and constructionist design approaches, touched upon in Section 1.3 (specifically, Table 1.1): Jennings and Campos (1997, p. 12) state about reductionist and constructionist views: "However neither of these approaches is very satisfactory: the former erodes the agent's autonomy, and the latter is something of a black art". In the domain of PPC, this implies that a hierarchical PPC approach can be designed according to known and established design rules, since the system behavior is deduced top-down in a reductionist fashion (c.f. Section 1.4), these approaches cannot account for autonomy on the side of low-level decision-making entities (agents), as this contradicts the necessity to plan behavior top-down. Distributed PPC, on the other hand, builds upon these autonomous agents, but the logical conclusion from individual to aggregate behavior is inherently difficult and subject to experience and guess work, which also rules out performance guarantees. This understanding is shared by Kalenka and Jennings (1999, p. 136), who discuss requirements for the application of agent-based architectures in e.g. manufacturing, to find: "In such systems, what is required is the ability to exploit the conceptual power of autonomous agents (as in the constructionist view), but to ensure the overall system performs in a coherent manner (as in the reductionist view)".

2.3.2 Experimental Validation and Design Recommendations are Missing

The above-stated hypothesis has implicitly or explicitly motivated many research attempts in the design of distributed (production) control systems. As will be discussed in Chapter 3, there is, in fact, a large number of previous scientific work that can inspire attempts to navigate the design space of (semi-)distributed PPC systems in search for an optimal mix between hierarchical and heterarchical features. This subsection will review existing quantitative research that explicitly seeks to investigate the relationship between system control and achieved performance. It will be concluded that the existing literature can neither explain nor fully reproduce the hypothesized impact of distributed control on production performance. While the experiments reported in this section provide empirical evidence, they cannot be translated into guidelines for production system designers to develop manufacturing systems that combine the benefits of both control styles.

Indications that neither purely hierarchical nor purely distributed PPC architectures were a promising design guideline for PPC systems were found fairly early in both laboratory systems (Paul et al. 1997) and simulation studies (Rogers and Brennan 1997), where a hybrid control architecture (in the sense of Class II systems in Fig. 2.2) outperforms strictly hierarchical and strictly heterarchical competitors with respect to both flowtime and tardiness. In Brennan (2000) and Brennan and Norrie (2001), the experiment setup is extended to show that the advantages of the hybrid control architecture over a purely distributed control architecture require a minimum time given to the "bargaining agents" to coordinate (through negotiation). The required time increases when a stochastic production environment is assumed.

More recently, Scholz-Reiter et al. (2009a) seek to experimentally investigate and confirm their previous hypothesis about a curvilinear development of performance as a function of the "degree of autonomous control" (Fig. 1.4). In their simulation study, products visiting one of multiple identical and parallel machines are increasingly allowed to choose the machine freely and deviate from an offline-created schedule. While the study finds the performance increases with increased such autonomy, it cannot reproduce a decrease in performance for high levels of autonomy. An optimal, medium degree of distributed control is, in fact, found in the work of Gronau and Theuer (2016), investigating the ideal degree of autonomy in the application of CPS for PPC. Their conclusion, however, is based on the hypothesized nature of performance curves and therefore cannot be considered an experimental proof of the assumption.

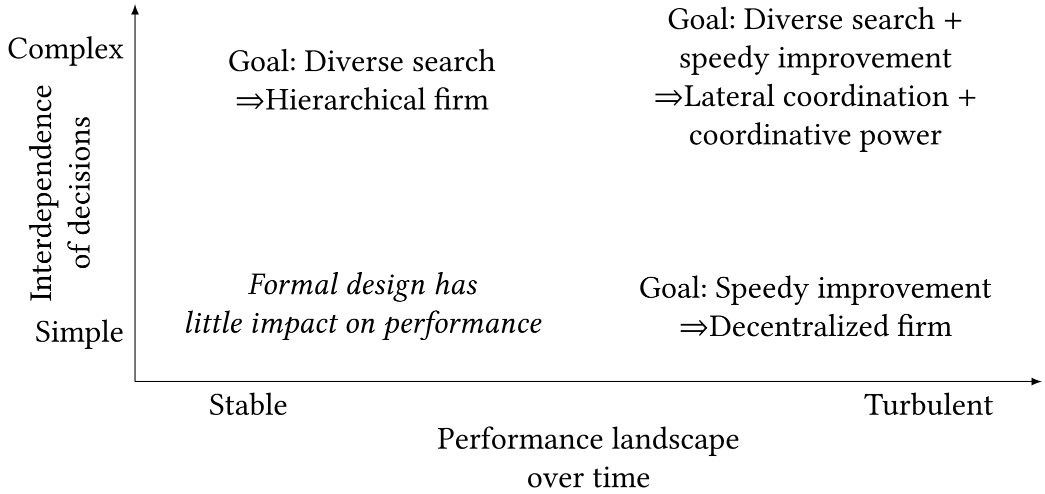

Figure 2.3: Appropriate goals and formal designs for organizations as a function of the structural and dynamic complexity of the environment, adapted from Siggelkow and Rivkin (2005).

In the context of organization theory, Siggelkow and Levinthal (2003) and Siggelkow and Rivkin (2005) apply minimal models to investigate the impact of organization structure (among other things) on company performance. The results of Siggelkow and Rivkin (2005) suggest that in simple and stable environmental conditions, more hierarchical organization structures excel, whereas more turbulent (yet still simple) environmental conditions call for more decentralization and heterarchy among organization entities. They formulate a set of recommendations that are depicted in Fig. 2.3. A similar nod to the relationship between organization design traits and performance characteristics was found by Dechaume-Moncharmont et al. (2005). They show, using a stylized model of bee foraging, that accumulation of more information does *not* automatically improve system performance where the environment is subject to change (food sources appear and disappear), as information takes time to be transmitted and may be outdated. They find that over a large parameter space a combination of exploiting bees (those who follow accumulated global information) and exploring ones (those who test new routes independent of available information) achieves the best system performance. The incongruity of high performance and quick adaptation has been reported as the result of multiple experiments: Carley and Ren (2001, Sec. 4) simulated operations of armed forces, concluding that "it is difficult, and perhaps impossible to design for both adaptability and high performance". Mason and Watts (2012) summarize existing experimental research on the relationship between network structure and organization performance (e.g. Fang et al. 2010; Lazer and Friedman 2007) by stating that hindering the speedy spread and adoption of "trial solutions" throughout the network (as a centralized, hierarchical organization structure would facilitate, c.f. Section 3.3.2), agents can be forced to spend more time on exploration of decision alternatives, instead of converging "prematurely on a suboptimal decision" (Mason and Watts 2012, p. 764). Solow and Szmerekovsky (2006) set out specifically to ask "how much control should be exercised to achieve optimal system performance;

or, in other words, under what conditions do systems benefit from different amounts of central control?" (Solow and Szmerekovsky 2006, pp. 52 f.; c.f. also Hazy et al. 2007). To this end, they formulate an abstract optimization model that depends on a parameter $\lambda \in [0, 1]$ termed the "amount of control" and conjecture conditions for the effectiveness of both very low and very high levels of control. However, they are not able to make statements about the medium range.

Section 2.3 has substantiated the research hypothesis first formulated in Section 1.4. It can be concluded that existing literature does not provide a sufficient answer to research question Q_1, and, while being widely hypothesized, experimental results for research question Q_2 that manage to reproduce the assumed curvilinear relationship between the "degree of autonomous control" and manufacturing performance is scarce, analytical evidence is missing. Chapter 3 will build upon this result by trying to establish a design space classification that aggregates design decisions allowing a designer to position the production system in the trade-off between hierarchical and distributed control. Research question Q_2 will be picked up in Chapter 4, where a minimal model will be deployed to deliver the first-ever analytical evidence for the shape of the performance curve as a function of (one possible feature of) hierarchical PPC.

Chapter Three

Navigating the Design Space between Hierarchical and Distributed Production Control — A Classification of Design Choices

> *"how far can we go in analyzing alternative coordination processes for problems such as resource allocation? Can we characterize an entire 'design space' for solutions to this problem and analyze the major factors that would favor one solution over another [...]?"*
>
> Malone and Crowston (1994, p. 110)

Previous Publications

The classification model presented as part of this chapter (Fig. 3.7) has been published in a more condensed form in:

Henning Blunck and Julia Bendul (2016). "Controlling Myopic Behavior in Distributed Production Systems — A Classification of Design Choices." In: *Procedia CIRP* 57.Supplement C. Factories of the Future in the digital environment - Proceedings of the 49^{th} CIRP Conference on Manufacturing Systems, pp. 158–163. issn: 2212-8271. doi: 10.1016/j.procir.2016.11.028

An initial attempt to validate the model was made in:

Tianyi Wang, Henning Blunck, and Julia Bendul (2017). "Exploring the Design Space for Myopia-Avoiding Distributed Control Systems Using a Classification Model." In: *Service Orientation in Holonic and Multi-Agent Manufacturing, Proceedings of SOHOMA 2016*. Ed. by Theodor Borangiu, Damien Trentesaux, André Thomas, Paulo Leitão, and José A. Barata Oliveira. Vol. 694. Studies in Computational Intelligence. Springer International Publishing. Chap. 26, pp. 295–304. isbn: 978-3-319-51099-6. doi: 10.1007/978-3-319-51100-9_26

> Both publications, along with ideas developed in this chapter were submitted *after* the submission and defence of this thesis for peer-review in:
>
> Julia C. Bendul and Henning Blunck (2018). "The Design Space of Production Planning and Control for Industry 4.0." In: *Computers in Industry*. Under review

To operationalize the research hypothesis developed in the previous chapter, this chapter develops a classification model to describe the design space of such hybrid PPC systems. To this end, the concept of *myopia* as a feature of agents' decision making is introduced as a gauge to assess the "degree" to which a given PPC approach aligns with the idealized concept of distributed control (Section 3.1). Section 3.2 presents the proposed model in demarcation of existing attempts and provides an overview of the literature streams reviewed for this contribution. The following sections (Sections 3.3 to 3.5) then classify the literature along decisions made in the design and operation of manufacturing systems. Collectively, they describe a design space for hybrid PPC approaches that can be used to navigate the trade-off between the two "poles" of fully hierarchical and completely distributed control architectures (Section 3.6).

In the context of the thesis, this chapter's role is twofold. First, the developed classification model is the result of an interdisciplinary literature review reported here. Second, it will provide the mental framework for the quantitative investigations in the following chapters (c.f. Section 1.7), which will look more closely at the relationship between some of the presented measures to reduce the impact of myopic decision-making.

The work toward the reduction of myopia in agent-based PPC in this chapter builds upon and extends previous works by Zambrano Rey et al. (in particular in Zambrano Rey et al. 2013, c.f. also Section 3.2.1), by offering a new and more elaborate classification of myopia control methods and reviewing literature beyond the scope of agent-based production control.

3.1 Myopia

To conceptualize why distributed control can lead to a loss of performance in PPC settings, this thesis uses the concept of *myopia* or the degree of myopic behavior in agent's decision-making. First, Sections 3.1.1 and 3.1.2 will introduce the term and point to properties of distributed PPC systems that induce myopia. Section 3.1.3 reviews the possible approaches to measure the impact of myopic decision-making on system performance, many of which will be used in the reviewed literature as well as in the later chapters of this thesis. With this background, Section 3.1.4 discusses the aptitude of myopia as a reference to develop a design space for hybrid PPC systems.

3.1.1 Definition and Evolution of the Term

As a medical condition, myopia (also known as near- or short-sightedness) describes the affected person's visual inability to focus on objects further away (Merriam-Webster 2016).

The term, however, quickly entered the colloquial sphere as well as various scientific disciplines to describe decision-making processes that overemphasize the short-term (near) effects at the expense of more long-term (far) effects. In decision-making theory, the conceptualization of myopia goes back to Pearl (1988, Ch. 6.3.2), who defines myopic decision-making strategies as (1) analyzing each source of information separately and (2) taking a decision after each observed information source after which a new state will be terminated externally (Becker et al. 2009; Russell and Wefald 1991). Becker et al. (2009)[29] define the *myopic-greedy assumption* as comprising the following two elements: (1) each source of information is evaluated separately and (2) a one-step-horizon is applied in sequential decision-making.

In economic theory, the term has come to describe decision-making policies that predominantly focus on short-term gains while failing to recognize and consider more long-term effects. Benartzi and Thaler (1993) introduced the *Myopic Loss Aversion* (MLA) theory to explain the difference in returns between equities (stocks) and risk-free assets (treasury bills). In their argument, a combination of *loss aversion* (the higher sensitivity of individuals to losses than to gains) and *mental accounting* (the tendency of humans to perceive economic outcomes differently, based on the frequency and circumstances of evaluation) can explain the otherwise implausibly large difference in premiums. Myopia has been used in management theory[30] to describe a 30 year-long debate over the (alleged) undervaluation of long term growth perspectives by managers — in particular in the Anglo-Saxonian world — in favor of short-term gains (c.f. Felstead 2016; Laverty 2004). Laverty (2004, p. 950) defines myopia as “a characteristic of a decision that overvalues short-term rewards and undervalues long term consequences”. Pressure from capital markets is the most commonly assumed reason for such behavior (Marginson and McAulay 2008). A similar notion is also used in supply chain management theory, e.g. in the discussion of inventory models. Here *myopic decision policies* maximize the target function only for the next period while ignoring the development of planning periods further ahead (Lovejoy 1992; Wang 2001).

So far, notions of “near” and “far” have had a distinctly temporal dimension. However, it may also be interpreted in a more spatial fashion: In robot control, for example, myopic behavior is associated with insufficient information to make optimal decisions either as a deliberate decision to prune the search space during decision-making Mataric (1992) or due to a limited sensing range (Bajracharya et al. 2009). Such conception is also in line with March (1999, Ch. 11), who identifies myopia in the domain of organizational learning processes as a condition when long-term effects are ignored and the near neighborhood is preferred over the “larger picture”.

Myopia was first used in the context of distributed PPC by Trentesaux (2009) and said to be related to the “uncertainty of knowledge about the future states of both the control system and the controlled system, uncertainty that increases rapidly over time” (ibid., p. 975). Much of the research explicitly addressing myopic behavior in distributed production control settings was carried under his auspice as the PhD work of Zambrano Rey (Zambrano Rey 2014; Zambrano Rey et al. 2011, 2012, 2013, 2014). Zambrano Rey et al.

[29]Based on the *Meta-Greedy Assumption*, defined by Russell and Wefald (1991), which are fulfilled by algorithms that opt for the decision alternative with the highest immediate benefit.

[30]along with similar terms like “short-termism”.

(2013, p. 801) define that "myopic behavior appears when decision-makers overemphasize current-term results at the expense of long-term performance" and define further the difference between *temporal* and *social* myopia, describing an agent's inability to project the consequences of its own actions into the future (temporal) or with respect to other agents (social). Both temporal and spatial dimensions of "near" have been used in the discussion of myopia in distributed PPC: The term "spatial myopia" is explicitly used e.g. by Trentesaux (2009) to refer to situations where "local" decision criteria (the performance of one intelligent product) are applied to evaluate decision alternatives, instead of more global ones (e.g. overall schedule performance). Since the information considered and the decision-making approach taken by individual agents is subject to the manufacturing system designer's discretion, myopic decision-making can be associated with decisions of the manufacturing system designer and/or planner. This assumption is shared by Zambrano Rey et al. (2013, p. 802), who note that the "degree of myopic behavior present in heterarchical control architectures depends on their structural and operational design".

3.1.2 Triggers for Myopic Decision-Making in Distributed Production Planning & Control

Where myopia in managerial decision-making might be explained with limitations of the human mind and/or flaws in capital markets, it is not immediately clear why computerized agents should likewise fall victim to myopic decision-making. This section discusses properties of agent decision-making that can contribute to myopia.

Agents, we learn from Carlin and Zilberstein (2009), are subject to uncertainty about their own state, the state of other agents, and the consequences of their actions. They also act under *bounded rationality*, in that they need to render behavioral decisions using bounded resources (time, computational power) (Tokoro 1996; c.f. also Zambrano Rey et al. 2013). Marzo Serugendo et al. (2004), in addition, point to possible temporal (the agent's lifetime may be insufficient to achieve the agent's goal) and institutional (agents are limited through policies) limitations of agents.

Within the HMS community, bounded rationality is the key argument to assume myopic decision-making (and by extension to call for holarchic rather than fully decentralized, control system structures) (c.f. Valckenaers et al. 2008; Zambrano Rey et al. 2013). The argument is based on the science-philosophic discussions made in Simon (1996) that limited computation and communication capacities *have* to lead to bounded rationality in decision-making, since even the addition of more agents would only increase the necessary communication and hence could not "outpace" the rapid increase of uncertainty over time (Valckenaers et al. 2008). The existence of bounded rationality within a demanding and complex environment (c.f. discussion in Section 1.1) lead Valckenaers et al. (ibid.) to the conclusion that HMS cannot attain optimal solutions.

While sub-optimality is deemed inevitable om the above-mentioned case, there is strong evidence that the *degree* of myopic decision-making on part of the agents is also a direct result of decisions made by the control system designer:

First, the amount of information analyzed before decision-making[31] can influence myopia regardless of the decision-making function applied: Sengupta and Ackoff (1965) show and quantify the efficiency loss due to separated information silos as the difference between the optimum attained in a global optimization and the optimum attained when each agent has the responsibility of changing a subset of the decision variables, but its cost function (as part of the global cost function) also includes other decision variables, controlled by other agents.

Another design decision through which the local information horizon (the subset of information available to one agent for decision-making) is influenced is problem decomposition (Tharumarajah 2001, c.f. also Section 2.1), which contributes to myopic decision-making by creating artificial information silos and withholding information from decision-making agents. Pochet and Wolsey (2006, Ch. 2.2.3) show how decomposition leads to performance and flexibility losses in MRP environments. Since agent-based production control approaches generally follow a physical decomposition approach (c.f. Section 2.1), the available information in agent-based PPC systems will often be limited to those machines/products/etc. "closest" to the intelligent object making a decision.

A particular trigger for social myopia is selfish decision-making. As already discussed e.g. in Section 2.1.2, it implies the non-consideration of other agents in the decision-making process. There is strong evidence that within the limited lifespan and sequential decision-making of PPC agents, increasing the agents selfishness can lead to losses in overall performance:[32] Through a series of papers in the late 1970s and 1980s, Dubey showed that the coordination result achieved by selfish agents can have a significantly worse societal price (or benefit) associated with it in comparison with the (possible) outcome under centralized optimization (c.f. Dubey 1986). March (1999, Ch. 11) stresses that the solution deemed optimal for the survival of a subsystem may not be optimal for the survival of the system as a whole, thus leading to myopic behavior (from the standpoint of the overall system) when parts are allowed to follow their own "agenda". A similar assumption is made in the literature on agent-based production control when the reduction of autonomy is discussed as a way to reduce myopia (c.f. van der Vecht et al. 2007; Zambrano Rey et al. 2013, this argument is extended in Section 3.5.3).

3.1.3 Measuring the Impact of Myopia

So far, the definitions given in Section 3.1.1 as well as the triggers discussed in Section 3.1.2 allow to judge if a given decision *strategy* exhibits myopic behavior. On a *system level*, the *impact of* myopic behavior is an emergent property in the sense of Section 1.5.1 that arises from the interplay of system elements within the controller and in the plant. Measuring the degree of myopic decision-making on the agent's part does *not* allow to quantify its effect on performance metrics at the system level.

[31] Decision-making based on limited information was at the core of the *meta-greedy assumption*, c.f. Footnote 29 (p. 49).

[32] The question to what extent pre-occupation with one's own self-interest generally boosts or impedes society is a question for economic theory (c.f. James Jr. and Rassekh 2000) that this thesis is not interested in.

Where changes in system performance are associated with myopic decision-making, the existence of a "benchmark scenario", a comparable decision strategy where agents do not (or to a lesser extent) comply with the definition of myopic decision-making is implicitly assumed. The comparison of some performance measure across the two allows the quantification of the impact of myopia on this performance measure.

Let h_m and h_{ref} be two control approaches, where h_{ref} is known or believed to exhibit a lesser degree of myopia, one may calculate the (relative) *impact* of myopic behavior on a performance metric f of a system (δ_f) as

$$\delta_f(h_m, h_{ref}) = \begin{cases} \frac{f(h_m)}{f(h_{ref})} & \text{where a } \textit{low} \text{ value of } f \text{ is desirable} \\ \frac{f(h_{ref})}{f(h_m)} & \text{where a } \textit{high} \text{ value of } f \text{ is desirable.} \end{cases} \tag{3.1}$$

A very similar definition is, in fact, well-known in the field of game-theory: Koutsoupias and Papadimitriou (1999) first used the term *coordination ratio* to describe the ratio between the overall (or *social*) cost associated with a Nash Equilibrium (NE) and the social cost associated with the socially optimal solution. The concept is today usually referred to as the *cost of anarchy*, a term first used by Papadimitriou (2001) (c.f. Koutsoupias and Papadimitriou 2009). Above given definition is less strict in allowing h_{ref} to be not socially optimal (a fair assumption, given the mathematical intractability of many decision problems related to manufacturing) and accounting for situations where a higher performance metric is deemed advantageous. As with the cost of anarchy, this measure of myopia takes a value > 1, if and only if the less myopia-affected approach improves the observed performance metric.

The literature provides three approaches to attain a relative assessment of the impact of myopic decision-making by comparing performance metrics.

First, where the agent equilibrium can be computed analytically, we can compare the social cost or the overall benefit accrued by all agents compared to a reference setting that is known to minimize social cost (be socially optimal). Such an approach is naturally applicable to general equilibrium models, as studied by Caliendo and Gahramanov (2013), who use the approach to quantify myopia in saving decisions by consumers or Namatame and Sasaki (1998) who study the total utility gained by competitive and cooperative agents over increasing population sizes. Bartholdi et al. (1993) exploit analytically determinable optima, to find that resource-allocations based on the honey-bee algorithm accrue at worst twice the cost of the optimal allocation. Another stream of research that supports the analytical derivation of the impact of myopic decision-making to which this chapter, as well as Chapter 5, will make repeated references, is algorithmic game theory with applications to selfish routing (c.f. Section 1.5.4). There, the socially optimal solution can be calculated and compared to some equilibrium reached by myopically-acting agents (Beckmann et al. 1955; Roughgarden 2005, Ch. 2, and Section 5.3.4 of this thesis).

The impact of myopic behavior may also be assessed when the decision function can be analytically defined and parametrized and when empirical data is available. This approach is chosen by Himarios (2000) in an attempt to assess consumer behavior in economics.

He defines a general decision function class that includes a parameter for myopic decision-making (weighing short-term gains over long-term gains) and determines the degree of myopic behavior by this function to empirical data.

Finally, where aggregate agent behavior can only be studied numerically (e.g. through a simulation model), the degree of myopia may be estimated by measuring some performance metric in *ceteris paribus* comparisons of two control approaches. An example is given in Zambrano Rey et al. (2014), who use the change in completion time variance in a HMS to assess the relative impact of myopia, comparing two control approaches. Given the prevalence of agent-based and other forms of simulation models in research on manufacturing system control (c.f. Section 3.2.2), it is not surprising that most of the PPC literature reviewed in the following sections applies such simulation-based myopia assessment.

Such empirical approach to assess the presence and the effect of myopia constitutes a problem: In any larger-scale manufacturing system, the comparison of two control approaches will measure a net difference, a net effect of shifting the balance between more and less myopia-prone control approaches. The direction of this net effect, however, is unclear. One cannot know if a decrease in the degree of myopia will lead to an improvement or deterioration of the performance measure. Myopia, Section 3.1.2 has shown, is a necessary and inevitable by-product of the key traits of distributed production control as defined in Section 2.2.2: selfish decision-making of agents under bounded rationality based on limited local information. The same features that cause myopic behavior also *enable* the desirable traits of distributed PPC such as responsiveness, adaptability, flexibility, and so on (c.f. Section 2.2.2). These traits are at risk where the designer seeks to limit or even rule out myopic behavior. It should be emphasized that minimizing myopic behavior in a production control architecture is *not* a goal in itself and the absence of myopia is not suitable for assessing the quality of a production control architecture.

3.1.4 Myopia as a Reference for the Design Space Classification of Hybrid PPC Systems

Changing the level of myopia, however, *can* guide the system designer to move the systems closer to the architectural poles of hierarchical PPC (low levels of myopia) or strictly distributed PPC (high levels of myopia), allowing a target-driven exploration of the design space for hybrid PPC systems that falls in between the two poles.

Research questions Q_0 and Q_1 pose in broad terms the question as to how manufacturing systems should be designed to achieve optimal performance under distributed control. In the light of Sections 1.2 and 2.3, research question Q_1 is concerned, in particular, with design decisions that can place a manufacturing system (plant and controller) between purely hierarchical and strictly heterarchical control, i.e. define a hybrid control system, as these are hypothesized to provide the desired, good performance characteristics. To understand and purposefully design hybrid control systems, it is necessary to characterize and explore a design space between HPP, on the one hand, and (strictly) heterarchical, distributed control systems on the other hand. Unfortunately, already the brief descriptions in Section 2.2 suggest multiple possible dimensions along which both could be compared:

from the form of agent coordination to control network structure, the information horizon of decision-making entities, etc.

To integrate these various design dimensions within one conceptual model, this thesis hence the degree of myopic decision making to conceptualize a design space for hybrid PPC systems, i.e. a continuous design domain between the previously introduced architectural styles of hierarchical and distributed PPC (Section 2.2). Likewise, for any given PPC approach, the degree of myopia present can be used as a gauge to assess the position of this control approach within this design space.

Based on the previous discussion, three key considerations support the choice of myopia:

First, by referring to the properties of agent decision-making as the landmark property of distributed PPC — not the presence of agents in the first place — it is acknowledged that also hierarchical PPC are DDMs that can be thought of and implemented as MAS. Notably though, agents in a hierarchical PPC do *not* exhibit (as much) the triggers of myopic behavior, as collected in Section 3.1.2.

Second, the close relationship between the drivers of myopic behavior, identified in Section 3.1.2, and the defining characteristics of distributed control, as derived earlier in this thesis (c.f. Section 2.2.2), support the hypothesis that many design decisions made to "distribute" PPC, can also be interpreted as decisions that enhance myopic decision-making.

Finally, myopia has been used before to conceptualize the negative traits of (strictly) distributed PPC and systems in more general terms. Previous sections have already pointed to the work of Zambrano Rey (2014) and Zambrano Rey et al. (2013). Even closer to the idea pursued here, yet conducted in a different domain (software-defined networks), is the work of Matni et al. (2015), who use (the absence of) myopia to characterize fully centralized and fully distributed myopic algorithms that then serve as the corner stones to define a more continuous hybrid design space for software-defined networks.

3.2 A Classification Model for the Design Space of Hybrid PPC Systems

3.2.1 Existing Literature on PPC Design Space Classification

In order to answer research question Q_1, this chapter develops a classification system for design decisions that can limit the impact of myopic behavior and hence position a particular control approach in the design space of hybrid PPC systems. Closest to this work is a similar classification of design choices presented by Zambrano Rey et al. (2013). There, the authors pick up earlier work from MAS-theory[33] to assign design ideas in HMS research that aim to reduce *myopic behavior* into one of three basic mechanisms,

[33] Namely van der Vecht et al. (2007) who, in turn, build upon Barber et al. (2001) and Falcone and Castelfranchi (2001)

through which agents (at run time) can be influenced to reduce myopic behavior. The three mechanisms are (c.f. also van der Vecht et al. 2007):

Influence by environmental modification occurs when the environment within the agent's information horizon is modified, changing the agent's (perception of its) environment.

Influence by belief alteration describes the exchange of messages between agents, providing additional information that agents did not observe themselves.

Influence of goal/task determination is the process of one agent influencing the decision process of another agent. The influence can range from a suggestion to a binding command.

Through the three influence mechanisms, different stages in the agent's decision-making process can be affected (ibid.). This sentiment is in line with the conclusions of Schneeweiss (2003a, Ch. 13), who finds that coordination schemes between agents may influence the decision criterion, or the "decision field" (the set of possible decision alternatives), or both. In his doctoral thesis, Zambrano Rey uses a more elaborate framework that orders myopia reduction approaches by the degree of *interference* with the underlying concept of agent autonomy, ranging from approaches that improve how local information is perceived and shared to examples in which scheduling conflicts are actively resolved for the agents by a centralized authoritative instance (Zambrano Rey 2014). The work focuses only on online (during run time) applications of simulation and optimization techniques to improve agent-based planning and disregards e.g. the effect of changes to the system plant (c.f. Section 3.3). While Zambrano Rey (2014) and Zambrano Rey et al. (2013) focus on HMS literature for evidence of myopia avoiding design decisions, a more interdisciplinary review of cooperative control with applications to distributed PPC systems is presented by Monostori et al. (2015), who focus on introducing key concepts and terminology from related fields of research without providing a classification or giving design recommendations.

Pach et al. (2014) review existing ideas for hybrid (or semi-heterarchical) PPC systems and identify that the two fundamental approaches employed for control systems that seek to combine hierarchical and heterarchical traits are structural dynamics (the system architecture may evolve over time) and the "control homogeneity" that determines, whether changes in control should apply to the whole system or only certain parts. This concept was used for a comprehensive review of hybrid PPC approaches by Cardin et al. (2015). However, the approach does not go beyond these two design dimensions, and the literature reviewed stems solely from the domain of distributed PPC.

In other related literature, Böse and Windt (2007) introduce a morphological box to assess a system's degree of autonomous control. The evaluation assigns points when particular criteria of decision-making, information processing, and decision execution are met (e.g. local storage of information, ability to interact with other agents, heterarchical organization structure). Similarly, Shen et al. (2006b) try to classify existing agent-based production control architectures according to different design criteria like agent-encapsulation, -organization, and system dynamics. Both contributions, however, do not address myopic behavior or the design space between distributed and hierarchical PPC.

Such trade-offs are discussed in organization theory where Pettigrew and Fenton (2000, Ch. 10) investigate dualities arising in innovative organization structures. Their findings are aggregated by Thomas et al. (2005a) to extend the classical "control vs. autonomy duality" with additional dualities such as "control vs. self-organization", "hierarchies vs. networks", and "stability vs innovation". Their work, however, is outside of the domain of PPC and lacks recommendations to navigate the trade-off that can be operationalized in manufacturing system design.

For the more abstract class of Distributed Artificial Intelligence (DAI) systems, Decker (1987) provide a taxonomy of design-decisions, mentioning the following four dimensions (c.f. also Stone and Veloso 2000): (1) agent granularity (coarse vs. fine) (2) heterogeneity of agent knowledge (redundant vs. specialized) (3) method of distributing control (benevolent vs. competitive, team vs. hierarchy, static vs. shifting roles) (4) communication possibilities (blackboard vs. messages, level of content). In the same publication, they provide a taxonomy for distributed problem solving, that distinguishes along the two dimensions of "control" and "communication". Nevertheless, design characteristics are not connected with expected emergent properties in both taxonomies.

The work presented in this chapter goes beyond the contributions by Zambrano Rey (2014) and Zambrano Rey et al. (2013), by (1) drawing from a broader set of research disciplines (not only HMS literature, c.f. Section 3.2.2), (2) incorporating decisions affecting the system plant and not only the controller, and (3) connecting the found evidence with decisions in the design and operation of manufacturing systems, thereby facilitating the application of the results during the design of future production (control) systems.

3.2.2 Research Disciplines Reviewed

When proposing a heuristic framework, it is important for the authors to state their own backgrounds (c.f. Kubicek 1977, Sec. 4). So, before discussing in more detail the proposed classification model, it should be outlined which research streams were considered during its creation.

This thesis will mainly focus on the research fields introduced in Section 1.5. Exhaustiveness, however, can *not* be claimed, as the abstract question of how to model, predict, and shape collective emergent behavior interfuses more research disciplines then could reasonably be covered here.[34] Instead, a positive argument can be made for the research streams considered here:

In the wide field of CAS theory, this research will particularly draw upon findings in statistical physics. As Lesne (2006, p. 248) note: "The determination of 'macroscopic'

[34]For further evidence, here are some of the research streams recommended for further attention to understand self-organized, complex systems: Monostori et al. (2015) cites ideas cooperative control theory, cooperative game theory, and cooperative learning. Schneeweiss (2003c) recommends to review results in applied mathematics and OR, micro economics (particularly: game theory), accounting, and artificial intelligence to understand DDMs in the context of supply chains. For the design of MAS (as an application of harnessing the power of emergence in a goal-directed manner), Tokoro (1996) recommend borrowing from game theory, molecular biology, ethology, cerebrum physiology and complex systems theory.

collective behavior of a large number of 'microscopic' elements lies at the very core of statistical mechanics". The thesis will focus in especially on research surrounding CAs (c.f. Section 4.2.3) (or comparable discrete dynamics) on networks. CAs and concepts that have been developed/investigated through them (such as *artificial life* and *self-organized criticality*, c.f. Bak et al. 1987) are established tools from statistical physics for the investigation of emergent system behavior and have been mentioned by both Castellani and Hafferty (2009) and Goldstein (1999) as relevant research streams. By discussing them in the context of the network structure, the descriptive power and available quantitative measures developed in the context of network science (Barabási 2012; Costa et al. 2007), which have successfully been applied to describe and investigate manufacturing systems and supply networks at different aggregation levels (Becker 2016), can also be used.

ABM have themselves advanced our understanding of complex systems (Castellani and Hafferty 2009). Agent-based Discrete Event Simulations (DESs), in particular, have been the primarily used modeling method to investigate the performance of self-organization-based PPC systems in production logistics and to model distributed PPC systems. Reasons for this include:

- Discrete event simulation is the method of choice in modeling manufacturing systems regardless of control architecture (Cassandras and Lafortune 2008, Ch. 10; Kuhn and Wenzel 2008; Wenzel et al. 2008, Ch. 1) and their scheduling (Hopp and Spearman 2008, Ch. 15.6.2).
- The physical decomposition approach of production systems results in a modular system concept, agents are adept to represent (c.f. Mönch 2005, Ch. 2.6.5.2; Okubo et al. 2000).
- MAS bear close resemblance to the reactive, geographically distributed nature of production control where decisions are taken in response to discrete events from the shop-floor (Mönch 2005, Ch. 2.6.5.2).
- Alternative control methods of any architecture are easy to implement (Ringhofer 2012).

With graph dynamics on the one hand and DES on the other hand, the thesis also covers the two extreme ends of the range of models suitable to represent discrete event systems in terms of model complexity[35].

The following sections will also cite supporting evidence from the domain of organization science. The field has long investigated the evolution and performance of organizations arising from individuals. More recently, it has changed its perspective toward a bottom-up, constructionist view of organizations (Section 1.5.3). It has also, as will become apparent in the remainder of this chapter and Section 4.1.2, produced work that seek to explain and understand collective behavior in organizations through minimal model experiments.

Finally, this thesis will extend the review, beyond research streams mentioned in Section 1.5, by looking at results from OR, the research domain traditionally concerned with questions of scheduling.

[35] c.f. Cassandras and Lafortune 2008, Ch. 1.3.3 and Ch. 10.3; Reveliotis 2005, Ch. 1; Zimmermann 2008.

3.2.3 Goal & Structure of the Classification Model

The classification model to be developed in the following sections addresses research question Q_1 of this thesis:

> Q_1: *"Which design decisions concerning both controller and plant impact the duality between hierarchical and distributed control in PPC?"*

An answer is formulated by (1) providing a set of dimensions through which the degree of myopia can be adjusted and (2) giving implementation hints as to how the degree of myopic behavior can be adjusted along each of the dimensions. The goal of the model is to allow manufacturing system designers to purposefully adjust the degree of myopic behavior in their systems. Given the widely assumed preference for hybrid control systems (c.f. Sections 1.2 and 2.3), the classification models gives designers concrete design suggestions, how to design the system plant and controller for optimal performance under distributed control (research question Q_0).

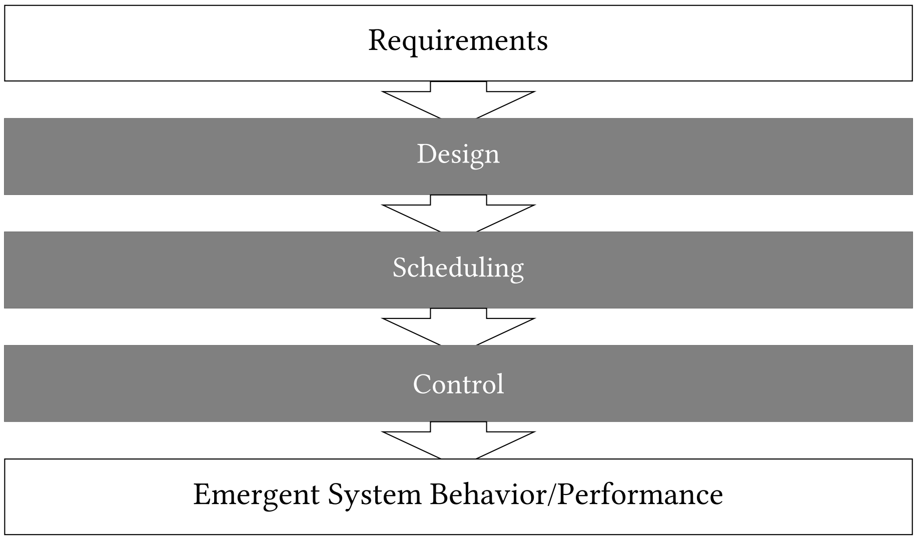

Figure 3.1: Basic structure of the proposed Myopia control classification model.

As an initial structure for the model, the three hierarchical decision processes of (system) design, (production) scheduling, and (production) control are assumed (Fig. 3.1). The process is hierarchical in that each step has to take into consideration decisions made in the preceding steps, as well as new information and (random) effects previously unknown. The model picks up the processes within PPC (c.f. Section 2.1.3) in which decisions with *concrete* ramifications for individual orders are taken, while not considering (aggregate) planning steps, that largely focus on anticipated demand (Gudehus and Kotzab 2009, Ch. 2).[36]

As a collective process, the three decision steps take an input and produce a resulting system behavior as output. The input is given by the *requirements* defined for the manufacturing system. Following Bonney (2000), these include *market* requirements (notably

[36] System design decisions are also taken based on assumed aggregate demand. However, owing to the strategic nature of the decisions, they still have concrete ramifications for individual orders.

the production mix and volume) (c.f. also Chryssolouris 2006, Ch. 6.2) as well as the *company* requirements and philosophy. Those include, in particular, the set and relative importance of manufacturing performance measures (Bonney 2000). The output is constituted by the observed system behavior, and the resulting performance metrics which are assumed to be — following Section 1.3 — at least partially emergent in nature.

Sections 3.3 to 3.5, pick up the three stages of Fig. 3.1 and introduce the design dimensions as well as the literature found to be related to each of them. The resulting classification model is presented toward the end of this chapter in Figure 3.7.

3.3 Myopia-Control during System Design

Following the common decomposition of systems into a plant and a controller (c.f. Section 2.1 and Fig. 2.1), this thesis breaks down the design stage into a discussion of design decisions w.r.t. the plant and the controller.

3.3.1 Plant Design

By addressing questions of physical systems design as part of the classification model, this thesis considers the impact that the design of the physical material flow has (1) on the performance of any production architecture and (2) on the selection of the "best" control architecture. As McKay (2011, p. 29) paraphrases, it "has become a historical and sometimes hysterical truth — if the various components of manufacturing planning and control are well matched to the environment, things will go well, and if they do not make sense, chaos and mayhem will dominate the milieu". In this subsection, the impact of both (operation) flexibility and system complexity on the impact of myopic behavior will be discussed. These (abstract) system properties are chosen, since they are frequently discussed in the context of (distributed) PPC and its future challenges (Section 1.1.1). The separate discussion should not obfuscate the fact that the two are intrinsically interlinked (Chryssolouris et al. 2013).

Notably, changes to the system plant do *not* change the degree of myopic behavior associated with the agent's decision making (as do design decisions discussed from Section 3.3.2 onward). However, changes in the system plant can be expected to lower the degree to which myopic behavior affects the system performance negatively (c.f. also Section 3.1.3). Through system plant design decisions that reduce the negative impact of myopic decision-making, the designer may tolerate more myopic behavior of the agents (which is shaped by subsequent design decisions, as discussed in Sections 3.3.2, 3.4 and 3.5). It is, therefore, important to consider also design considerations concerning the system plant to fully answer research question Q_1.

Operation Flexibility

Agent-based production control approaches have been framed as a quest to exploit the flexibility potential in manufacturing systems (c.f. e.g. Brennan and Norrie 2003; Van Brussel et al. 1998; Windt et al. 2010a). Leitão (2009) see flexibility as the gateway to efficiency in distributed production control approaches with medium to high demand variability[37]. Flexibility in manufacturing systems gives intelligent products under distributed control the ability to react swiftly to disruptions (Brennan and Norrie 2003; Van Brussel et al. 1998).

There are various classifications and nomenclatures to describe forms of flexibility in manufacturing systems (c.f. e.g. Chryssolouris 2006, Ch. 1.3; Sethi and Sethi 1990). The thesis will focus here especially on what Chryssolouris et al. (2013, p. 6788) call *operation flexibility*, the ability "to produce a set of products using different machines, materials, operations and sequences of operations". It is the result of flexibility on the side of machines and processes, and the manufacturing system itself (material flows, etc.) (Chryssolouris 2006, Ch. 1.3; Chryssolouris et al. 2013).

While flexibility is deemed necessary to reap the benefits of distributed PPC architectures, its presence obviously increases the scope for negative impact of myopic behavior, as agents are given more decision alternatives. This effect can be observed in results off algorithmic game theory, in particular "selfish routing games" (c.f. Chapter 5). In this model setting, agents given flexibility w.r.t. their route choice and making that decision in a purely self-interested fashion may induce a price of anarchy — an increase in social cost compared to the centrally coordinated, socially optimal flow distribution. In particular, Braess (1968) showed that adding flexibility (in the form of path alternatives) may actually increase the overall accrued cost. After the author, this phenomenon is commonly called *Braess' Paradox*. Christodoulou and Koutsoupias (2005), Correa et al. (2004), Czumaj et al. (2010), and Roughgarden and Tardos (2002) all find strict maxima on this particularly striking example for the cost of anarchy.

Despite the analytical nature of the results, they do not establish a clear functional relationship between the degree of operation flexibility and the cost of anarchy. Roughgarden (2003) showed that very simple networks are sufficient to produce a high cost of anarchy. In addition, Papadimitriou and Valiant (2010) note that the cost of anarchy also depends on the overall flow on the network and is the most distinct for a rather small band of flow rates (a result confirmed by the experiments of Youn et al. 2008).

Some works have been reported where material flow networks are carefully extended or reduced to avoid the cost of anarchy, i.e. achieve high flexibility without suffering the negative impact of myopic behavior. Roughgarden (2001) considers the task of selecting from a set of candidate edges in a flow network the subset that exhibits the best performance under selfish routing, finding that the optimal solution to the problem is $\mathcal{NP}$ hard. In Chapter 5, the properties of a particular class of flow distribution, *utilization-attaining* flow distributions, are discussed, which seem desirable as they allow the system designer to implement "tactical underutilization" (Schönsleben 2012, Ch. 13.1.1) of capacity. This

[37] whereas for traditional hierarchical approaches, Leitão (2009) claim that specialization is the key to efficiency in settings with low variability

thesis establishes that these flow distributions can (only) be obtained for very material flow configurations of very limited flexibility. It is also concluded that focusing only on throughput time minimization on the side of the agents' decision-making function lies at the core of the unsatisfactory behavior in situations with higher levels of flexibility.

Complexity of the Material Flow

Like flexibility, *complexity* is a notoriously difficult concept to capture and a number of alternative measures have been proposed (Chryssolouris et al. 2013; Ebeling et al. 1998; Gell-Mann 2002; Philipp et al. 2007). In the context of manufacturing systems, it is a common practice to differentiate between static and dynamic complexity (c.f. Chryssolouris et al. 2013; Frizelle and Woodcock 1995; Vrabič and Butala 2012; Windt et al. 2008b). The overall impact of complexity on the manufacturing system and supply chain performance has been the subject of sustained academic interest and shall not be rephrased here (the reader is referred to e.g. Bozarth et al. 2009; Deshmukh et al. 1998). This thesis focuses instead on results that indicate the particular susceptibility of distributed control approaches for increasing complexity.

One reason to assume that distributed control approaches may face difficulties in the eye of increasing complexity is the need for more information and more complex decision routines required for more complex systems. Jones et al. (2002) conclude that the amount of information needed to make good decisions is directly proportional to system complexity. A similar statement is made by the "Law of Requisite Variety" (Ashby 1961, Ch. 11).

Further evidence comes from the strict focus of previous attempts of distributed production control — in particular PULL production — to rather pristine environmental conditions: Im and Schonberger (1988), Spearman and Zazanis (1992), and Spearman et al. (1990) hint at a number of requirements for a successful implementation of KANBAN principles, including a stable production environment with only a few part-numbers on the line, standardized jobs, and short setup times. Stevenson et al. (2005) accordingly fundamentally rules out the application of KANBAN in make-to-order environments. It was experimentally shown as early as 1987 that the success of KANBAN depends on these rather pristine environmental conditions, with performance dropping quickly as complexity levels increase (Krajewski et al. 1987). Moreover, Gudehus and Kotzab (2009, Ch. 8.8.7) state that "local scheduling" approaches (like PULL systems) work well under stable conditions, allowing to eliminate PPC overhead and distribution of responsibility. However, problems may occur from unexpected variations in work content, new products, and changing requirements. More broadly, Browning and Heath (2009) conceptualize that negative results of lean implementations are driven by increasing uncertainty and instability. So strong is the connection between the application of PULL concepts and the reduction of system complexity in the introduction process that authors like Krajewski et al. (1987) and Spearman et al. (1990) see a reverse effect: For them, the introduction of KANBAN control approaches can be understood as an effort to actively shape the structure of the plant. Spearman et al. (1990, p. 881) suggest that a "significant portion of the reason for kanban's apparent superiority over push systems may be its requirement for, and facilitation of, environmental improvement".

In the realm of agent-based production control, some authors have investigated the impact of increasing static and dynamic complexity. Philipp et al. (2007) and Scholz-Reiter et al. (2006) compare two distributed control approaches — one that makes decisions based on queue length and another, pheromone-based approach that makes decisions on the basis of past system behavior, finding that, in particular, the queue length-based approach showed only little negative response in performance to increased complexity (be it in the form of system size, number of product classes, or standard deviation of processing times). It should be noted that the investigated system architecture, a linear sequence of fully connected sets of identical, parallel servers would show low complexity (when calculated as the entropy of connections between subsequent layers). Caridi and Sianesi (2000) show that negotiation-based distributed control performs well when demand is homogeneous (low dynamic complexity), but deteriorates with the complexity of product structure (static complexity), as the demand time series become more diverse.

Stylized Model Experiment

To further explore the impact of plant complexity on the performance of distributed PPC systems, this thesis investigates a stylized example of a flow-shop environment.

In particular, consider a flow shop with 30 consecutive *levels*, each comprising $N = 5$ parallel, identical servers. Operations have to be processed by one of the servers at each level, before advancing to the next level. This network is an extension of the one studied, e.g. by Scholz-Reiter et al. (2006), who use a configuration with three levels and three servers per level.

By default, every server will be connected to exactly one server in the next stage. However, additional links are possible. In particular, let p denote the probability of a connection between a server on level i and any other server on $i + 1$. To avoid artefacts at the boundaries of the system, the last level is connected to the first level, creating a circular setup, where agents can progress indefinitely through the levels.

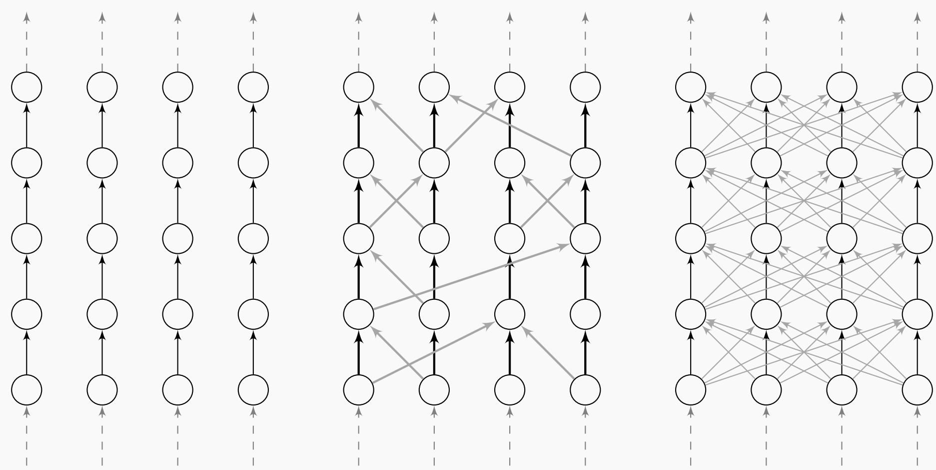

(a) $p = 0$. (b) $p = 0.25$. (c) $p = 1$.

Figure 3.3: Exemplary network configurations for different values of p. Randomly added edges are shown in gray. In the experiments, the last level is connected tp the first level, to form a ring structure.

Using the above-discussed entropy-driven measures of complexity, the complexity of the material flow network may now be assessed through the complexity of the connections between two subsequent levels. In particular, the probability that two random nodes on successive levels are connected is

$$p_e = \frac{1}{N} + \frac{N-1}{N} \cdot p \tag{3.2}$$

and the associated complexity entropy in bits is calculated as

$$H = -\log_2 p_e \cdot p_e. \tag{3.3}$$

Figure 3.4 shows the entropy as a function of p.

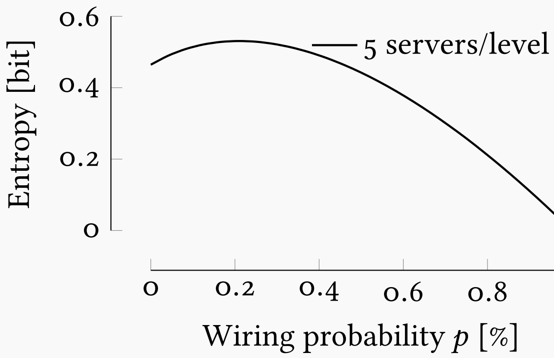

Figure 3.4: Complexity (measured as Entropy) of the material flow network as a function of the number of wiring probability p.

The performed experiment measures the number of levels that a product completes per unit of time. The number of products is identical to the number of servers. As the mean processing time on every machine is 1, agents should be able to clear one level per unit of time under ideal circumstances. As the actual processing times follow $\sim \mathcal{N}(1, 0.2^2)$, a performance < 1 can be attributed to waiting times.

In this setup, agents apply the following decision-making approach: after *every* level, the agents consider every possible path (sequence of servers) across the next four levels. For every machine on one of these paths, the current Work-in-Process (WIP) is observed and summed up to calculate the total expected waiting time per path. The path with the lowest WIP (and hence: expected waiting time) is chosen. Note that only the first decision, the server on the next level, is actually implemented, since a new planning is initiated after the next level. A similar, spatially defined information horizon in the context of distributed PPC was used previously e.g. by Becker et al. (2016) and Windt et al. (2010c).

For the experiments reported here, the setup comprises of 30 levels with $N = 5$ servers per level. The system is filled with 40% more products then servers in total (so $1.4 \cdot 5 \cdot 30 = 210$ products). The simulation is run for 500 time steps, with the first 250 being omitted as warm-up time.[a] The number of levels advanced per unit of time is reported as the performance measure.

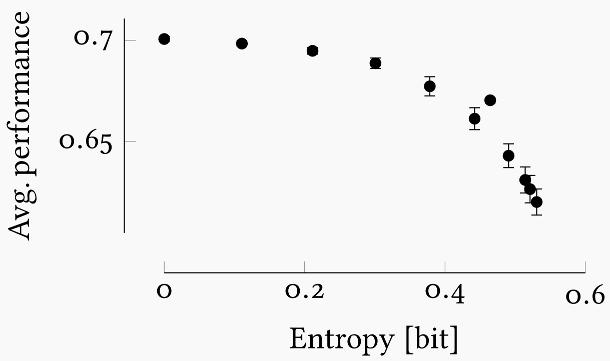

Figure 3.5: Measured performance as function of network entropy. Statistics calculated over 100 simulation runs.

The measured performance as a function of above-defined entropy, shown in Fig. 3.5, indicate a clear negative relationship between complexity and performance. It validates the hypothesis that changing the degree of complexity of the underlying material flow network can affect the performance of distributed PPC approaches. It also questions the validity of grids with fully connected levels (as used by e.g. Scholz-Reiter et al. 2006) as testing grounds for distributed PPC as these systems — despite their "complex" appearance (c.f. Fig. 3.2(c)) — have low entropy and are therefore prone to show good performance under distributed myopic decision-making strategies.
The discussion of this model is continued in Section 3.5.1, in the context of other design dimensions.

[a] The products are initially distributed equally across all servers, so a warm-start of the simulation can be assumed.

3.3.2 Controller Design

The system controller is made up of all decision-making entities. In a distributed PPC setting, these are the agents associated with entities on the shop floor.

The organization of these agents is one the most frequently discussed design decisions in the domain of agent-based PPC, and the design of MAS in more general and extensive discussions can be found in the literature[38]. Usually, a complete heterarchy, where communication is strictly between peers, is contrasted with a hierarchy which is purely minted by a command and response relationship between agents (c.f. e.g. Dilts et al. 1991). This points (again, c.f. Section 3.1.4) to the problem of discussing "hierarchy" in the context of PPC architectures: The term usually describes a combination of design

[38] e.g. Baker 1998; Bousbia and Trentesaux 2002; Ge et al. 2017; Horling and Lesser 2004; Isern et al. 2011; Leitão 2009; Ouelhadj and Petrovic 2009; Shen 2002; Shen et al. 2006a,b; Wong et al. 2006.

decisions, e.g. on the coordination between agents. One example here is Brennan (2000), who test three control architectures with an increasing degree of hierarchy (inspired by Dilts et al. (1991)). However, in that contribution, agents at the top of the hierarchy also possess planning capacity and thus it is difficult to attribute the results to either myopia reduction mechanism. Staff holons in PROSA can similarly provide centralized algorithms like scheduling, although basic holons may not be forced to obey their recommendations (Heragu et al. 2002).

In an effort to decompose the various measures of myopia reduction associated with "hierarchical" control, the focus of *this section* will be on the role of structural hierarchy within the network of agents as an aid to communication and peer-to-peer coordination. The reader is referred to other sections (in particular Sections 3.4 and 3.5.1) for the treatment of aspects, where agents assume elevated positions in the agent society owing to their capacities and authority over other agents.

Degree of Hierarchy in Control Network Structure

Already Simon (1962) observes that hierarchy was a widely found feature within many complex systems across disciplines. He assumed that hierarchy allowed a system to advance more quickly in its evolution (c.f. also Koestler 1970; Thomas et al. 2005a). Several properties of the underlying task can give rise to hierarchical control structures: Passino (2005, Ch. 1.4.1) points, e.g. to problem decomposition, interdependencies between tasks, and physical distribution of system components. Galbraith (1974) sees increases in uncertainty as a driver to change from rule and program-based coordination to coordination through hierarchy. Hierarchical features are present in many networks beyond control, identifiable as the parallel existence of scale-free degree distribution and high clustering (Ravasz and Barabási 2003). Ravasz et al. (2002) show that the hierarchically organized clusters are a frequently organized organization scheme in biological systems, especially in metabolic networks.

It is widely assumed that horizontal (or heterarchical) and vertical (or hierarchical) agent architectures represent the two poles of MAS system design (c.f. e.g. Bond 1990). To further flesh out the decision space between these poles, Dilts et al. (1991) propose a discrimination between heterarchical, modified hierarchical (lower-level control agents are given some autonomy and the ability to exchange information amongst each other), and proper hierarchical.[39] Duffie and Piper (1987) and Duffie and Prabhu (1996) go one step further, in suggesting the existence of a continuous range of "degrees of hierarchy".

Horling and Lesser (2004) and Isern et al. (2011) point to a number of organization schemes applied in MAS design, some of which such as *federations* and *holarchies*[40] can be interpreted as forming a "middle ground" between heterarchy and strict hierarchy. Shen et al. (2006b), for example, see federations as the third major architecture blueprint for agent structures in distributed PPC between hierarchical and strictly distributed approaches.[41]. Vrabič et al. (2012) discuss the idea of using clustering algorithms as a

[39] They also included a purely centralized architecture with a single controller, which is omitted here for its lack of modern-day significance (c.f. Section 2.2.1 and Dilts et al. 1991).

[40] which pick up the concept of *holons* (c.f. Section 2.2.2).

[41] Called "autonomous agent approaches" by Shen et al. (2006b)

vehicle to find suitable groups of agents within a PPC setting to form a team, although without providing any experimental evidence for their hypothesis.

Beyond its ubiquity and popularity, there are strong analytical arguments in favor of hierarchy: Helbing et al. (2006a) show analytically that hierarchical networks provide the fastest information transmission across networks. They are, however, vulnerable to failures of decision-making entities. A similar conclusion was reached by Malone and Smith (1988). Also in the domain of MAS and PPC design, hierarchies are generally associated with more controlled and coherent behavior (c.f. e.g. Bond 1990; Duffie and Prabhu 1996; Horling and Lesser 2004; Trentesaux 2009). The often-observed high cost of centralized coordination, on the one hand, is explained by Crutchfield and Mitchell (1995) with (1) the potential limitation that central entities can place on information processing, (2) lack of robustness to faults, and (3) the equitable resource allocation necessary for a central entity. Decentralized structures, on the other hand, are found to be more robust at the expense of increased communication effort (summary from Malone and Crowston 1994). In minimal model experiments of organization structures, Siggelkow and Levinthal (2003) find that centralized decision-making may be indicated where cross-divisional interdependencies complicate a “clean” problem decomposition into departments.

Within the research community of distributed PPC, the belief in desirable properties of hierarchical control networks is so widespread that researchers, starting with Brennan and Norrie (1999) and Maturana and Norrie (1996), have considered the idea of dynamically changing the control network architecture of distributed PPC systems to temporarily include hierarchy (called “partial dynamic hierarchies” by Brennan and Norrie (1999)), should the need for increased global coordination be detected. This idea was later picked up in the ADACOR/ADACOR^2 reference architectures for distributed PPC (Barbosa et al. 2015; Leitão and Restivo 2006). ADACOR makes use of the ability to organize holons in both hierarchical and heterarchical fashions (Mařík and Lažanský 2007): A hierarchical control approach is followed during “normal” operation where global optimization is favored over high flexibility and reactiveness. In this setting, subordinate agents closely follow the decisions of higher-level “supervisor” holons. If an “emergency” is detected, the degree of autonomy is increased to allow a quicker adaptation to the changed environmental conditions (Leitão and Restivo 2006). This so-called *reactive control* has been studied as a way to combine the advantages of high and low degrees of hierarchy constructively by changing between them when the necessity is detected (c.f. Cardin et al. 2015; Pach et al. 2014, for reviews). Bongaerts et al. (2000) note that the concept of holons leaves some degree of hierarchy in the system (by virtue of the recursive definition of holon), thus lending itself as a design paradigm for semi-heterarchical systems.

While in ADACOR hierarchy in agent architecture and master-slave relationships between agents are again used simultaneously, there are also contributions to PPC literature where hierarchy is only used to further information sharing: Cardin and Castagna (2012) suggest through a simulation experiment that adding a central information gathering and providing resource holon can reduce myopic behavior not by making decisions itself, but by coordinating and facilitating the decision-making process among the lower-level agents. Such an approach is called a “Mediator Approach” by Ouelhadj and Petrovic (2009) and Shen et al. (2006a). Shen et al. (2006a, p. 572) say that the “Mediator approach is another type of federated architecture. In addition to the functions of a facilitator and a broker, a mediator assumes the role of system coordinator by promoting cooperation and learning among intelligent agents".

Outside the domain of PPC, the impact of network architecture on the performance of distributed decision-making systems to solve differentiation games on networks has been investigated by means of human subject networks in Kearns et al. (2006) and McCubbins et al. (2009). The authors found a significant impact of the number of connections between agents on the collective performance of human subjects (this result was confirmed for computerized agents in Hadzhiev et al. (2009)). The result is more differentiated where additional connections also impose additional constraints: Here experiments reported in Enemark et al. (2011) and Judd et al. (2010) indicate that lateral connections that impose additional restrictions on the solution reduce performance, whereas connections that do not lead to additional constraints continue to be beneficial. Kearns et al. (2006), in particular, showed that the existence of "leader nodes", which are connected to many other nodes in a non-random fashion and could be interpreted as a kind of central system control, lead to a significant increase in the solution-finding performance (ibid.). Chapter 4 will pick up this idea and model language to investigate in more detail the effect of "leader node" on the problem-solving capacity of MAS systems.

3.4 Myopia-Control during Offline-Scheduling

Before the start of the production, scheduling decisions (machine allocation and/or sequencing decisions) can be made *offline* without the necessity to respond quickly. Classically, OR techniques are used to find the optimal schedule for the planning period (c.f. e.g. Pinedo 2008). While such ex-ante planning through stylized models usually underestimate the complexity and dynamics that can occur in a real-world system (c.f. ibid., Ch. 16.1), the mathematical formulation still holds the promise of finding a global optima. Hence, even many proponents of distributed control find that some degree of ex-ante scheduling is desirable: Brennan and Norrie (2001, p. 244) conclude that the "localised, reactive approach to control, characteristic of heterarchical architectures, cannot result in globally satisfactory performance. Because of the highly structured nature of shop floor control some degree of planning is useful, even in an unpredictable environment". This section reviews ideas that seek to carry over some scheduling decisions from offline scheduling to online distributed production control.

Partial Schedule Flexibility

Under *partial scheduling*, static scheduling is used to pre-determine a *subset* of all scheduling decisions before the start of the production, with the remainder being determined *online* (Herroelen and Leus 2005; Ouelhadj and Petrovic 2009; Perez-Gonzalez and Framinan 2015). The underlying hypothesis is that "the global scheduling performance is determined primarily by a subset of the scheduling decisions to be made" (Wu et al. 1999, p. 113). These pre-determined scheduling decisions can then either *constrain* the set of decision alternatives given to agents during the production (Agre and Chapman (1990, Appendix) discuss this perception of planning as "plans-as-constraints"), or *guide* the decision-making of agents during runtime to more promising regions of the solution space, i.e. solution alternatives that are close to the statically calculated initial solution. An example of the latter approach is provided by Verstraete et al. (2008b).

Bongaerts et al. (2000) put forward the idea that partial planning can be used to position a PPC system in the design space between distributed and hierarchical PPC. They argue that a PPC system may be called fully autonomous if its agents are only bound by the technologically-induced scheduling constraints. Every additional scheduling decision made upfront, shifts the system toward the hierarchical perception of PPC. Grundstein et al. (2015) use a similar rationale when using the order release times determined hierarchically through an HPP system to control the order release into an otherwise heterarchically-controlled manufacturing system. There is even analytical proof that the absence of pre-determined scheduling decisions can lead to myopic behavior: Graham has shown the existence of what eventually became known as the "Graham Anomalies", the counterintuitive increase in throughput time after adding resources, when sequencing decisions are made purely on the basis of local priority rules (Graham 1966).

Partial scheduling can be classified as a special form of predictive-reactive scheduling where the baseline schedule is created largely without the anticipation of variability (as compared to "robust" scheduling approaches) and instead, only a subset of scheduling decisions is actually imposed at runtime. By not relying on a full schedule, the online adjustment to disturbances does not require another run of a scheduling engine, but can be conducted with little computational effort within the given constraints of the baseline schedule (c.f. Policella et al. 2007).

Pre-selective scheduling strategies, first introduced by Igelmund and Radermacher (1983), can avoid the Graham Anomalies through partial scheduling. The approach requires to designate for each "minimal forbidden set"[42], one operation as the "waiting operation". It is, thus, decided that this waiting operation may not start operation until *at least one* other operation of the critical set has been completed (c.f. Herroelen and Leus 2005; Möhring and Stork 2000).[43] Identifying the waiting operation is highly computational expensive (Herroelen and Leus 2005) and thus simplified approaches have been suggested, especially the *linear pre-selective policies* suggested by Möhring and Stork (2000), which use priority lists to identify the waiting operation, have received attention.

A related concept is discussed as *precedence constraint posting* (PCP) e.g. by Lombardi and Milano (2012). In this concept, the scheduling graph is altered by adding strict precedence constraints between operations (not requiring and AND/OR graph) to avoid resource utilization conflicts. This concept was extended by Policella et al. (2007) to the idea of *partial order schedules* that introduce precedence constraints in such a way that any possible temporal solution within the partial order schedule is guaranteed to be feasible (i.e. not overload the given capacity).

[42] A minimal forbidden set is the smallest possible set of operations that cannot be processed simultaneously because of capacity constraints (Lombardi and Milano 2012; Möhring and Stork 2000).

[43] This logic cannot be captured through standard scheduling graphs. Pre-selective policies instead have to be expressed through AND/OR graphs (Möhring and Stork 2000).

Another form of introducing partial schedules are *ordered group assignments*, introduced by Billaut and Roubellat (1996). In an ordered group assignment, operations are assigned to groups in such a way that within a group every temporal permutation of operations is permissible, but the temporal sequence of groups is fixed. Artigues et al. (2005) find that such an approach offered flexibility potentials "for free", in the sense that the worst-case performance did not deteriorate when more scheduling decisions were left for runtime decision-making. Ordered group assignments are also used as a pre-processing step by Wu et al. (1999), who term their approach Preprocess First, Schedule Later (PFSL). They present a branch-and-bound approach to determine the optimal group assignment. They compare fully static, fully dynamic[44], and PFSL-based scheduling in a stochastic job-shop environment, comparing the total weighted tardiness. They find the PFSL approach to outperform static scheduling and being at par or better with fully dynamic scheduling.

The idea of partial scheduling has been popular in the domain of agent-based production control from early on: Already one of the earliest attempts for a control architecture, the "Yet Another Manufacturing System" (YAMS) approach by Van Dyke Parunak (1987) used the aforementioned Contract Net Protocol to let agents determine fine scheduling, while a coarse-grained schedule is provided as input (Heragu et al. 2002). In Blunck and Bendul (2015), the authors convey the notion of a graph-based partial schedule to flexible job-shops where the a-priori assignment of tasks to machines is relaxed and an additional assignment problem has to be solved (Dauzère-Pérès et al. 1998). However, they cannot confirm the hypothesis that partial schedules outperform full and/or no a-priori scheduling. While analytical evidence is missing, it is frequently assumed that the sequencing problem can (with little loss in performance) be solved separately from (and notably: after) the allocation problem (c.f. e.g. Berger et al. 2010; Brandimarte 1993), creating a "natural" problem decomposition to apply partial scheduling.

Both the problem of myopic (even paradox) behavior and the solution approach through partial schedules, re-occur in a distinctly different setting and model class: Building upon the investigations on Braess' Paradox, selfish routing games, and the price of anarchy (already discussed in Section 3.3.1), researchers have investigated how to reduce the cost of anarchy by predetermining the path choice (routing) for a subset of the flow. While the research described here usually assumes non-atomic flows, the idea in principle translates to pre-determining the allocation decisions on a subset of the orders. Initial experiments on this so-called *Stackelberg Routing*, where a share $\alpha \in [0, 1]$ of the flow is pre-assigned a routing and the remaining flow responds selfishly to this situation (mimicking a Stackelberg game, hence its name), were published by Korilis et al. (1997a). Roughgarden (2004) showed that it is $\mathcal{NP}$-hard to find a strategy for the pre-assignment that minimizes the social cost (and hence completely removes the cost of anarchy), but there are (at least for very simple networks) simple strategies (such as assigning traffic to least attractive route option) that can be shown analytically to reduce the cost of anarchy. Karakostas and Kolliopoulos (2009) later extended the statement to hold for more complex networks, but it was shown in Bonifaci et al. (2010) that network families still exist where the price of anarchy grows with network size.

[44] using a dynamic dispatching heuristic, c.f. Section 2.2.2.

Given that scheduling problems are frequently expressed through graphs (starting with Roy and Sussmann 1964; c.f. also Błażewicz et al. 2000), schedule flexibility lends itself to quantification by using graph-theoretic approaches. Kolisch (2001, Appendix A.1) review the early work on how to assess the *network restrictiveness* of a schedule. Artigues et al. (2005, p. 318) count the number of still realizable production schedules as a measure of flexibility. The approach and implementation is limited to the ordered group assignment approach taken in the publication and cannot be compared across scheduling problems, as it is in absolute values. Aloulou and Portmann (2005) likewise measure flexibility by counting the number of still unresolved precedence relationships, normalized by the total number of possible precedence relationships in a graph (the number of undirected edges in a complete subgraph formed between all operations in the problem). Blunck and Bendul (2015) use the same nominator, but normalize by the total number of possibly unresolved scheduling decisions to attain a measure $\in [0, 1]$.

3.5 Myopia-Control during Production Control

During production control, agents can make decisions within a solution space that has been constrained by both system design and offline scheduling decisions. Unlike decisions made by the designer or planner in previous steps, decisions made during production execution are subject to strict time constraints and occur in a naturally distributed environment, i.e. multiple decision-making entities take decisions simultaneously.

This section explores three dimensions of agent decision-making through which the impact of myopic behavior may be reduced: (1) by extending the temporal information horizon of agents to include either past or expected events further into the future, (2) by taking measures to suppress "nervousness" in the system, and finally (3) by changing the decision making function of agents from strictly selfish/competitive to more altruistic behavior.

3.5.1 Adjustment of the Temporal Information Horizon

Many simple implementations of distributed production control have product agents taking decisions purely based on the current system state. A good example is an allocation decision based on current queue length levels across machine alternatives — an example of the wider class of dispatching rule-based control approaches, discussed in Section 2.2.2. This approach (found to be widely applied in agent-based manufacturing control by Tay and Ho 2008) is based simply on the system state at the current time, with no regard for past or expected future developments. From this baseline scenario, two approaches to further the information horizon for decision-making are possible: looking into the past (in the assumption that future behavior will be similar) or anticipating future system evolution to test and compare alternative courses of action. Here analytical and experimental evidence is compiled that the two approaches can improve the performance of distributed PPC approaches.

Looking into the Past

Maintaining knowledge of past system states can improve agents' decision-making. By having information about past decisions and outcomes, agents can exploit existing knowledge and more complex cooperative social interactions can arise.

This arguably most prominent example of utilizing past information for informed decision-making is the use of *stigmergy* (Marzo Serugendo et al. 2004; Berger et al. 2010; Bonabeau et al. 1999, Ch. 1.2.3), also mentioned by Zambrano Rey et al. (2014) as an example for myopia reduction. In this bio-inspired coordination approach, agents leave artificial *pheromones* to inform subsequent agents of the relative "success" of decision alternatives. With every agent leaving it's own pheromones, the perceived environment is subject to constant evolution (Bonabeau et al. 1999, Ch. 1.2.3; van der Vecht et al. 2007; Zambrano Rey et al. 2013). The computerized adaptation of stigmergy is today the primary form of agent interaction in swarm intelligence systems (Kassabalidis et al. 2001). By accumulating the past decisions, pheromones can provide agents with almost global information about the (longer-term average) system state (c.f. Armbruster et al. 2006; Peeters et al. 2001; Van Dyke Parunak 1997). Van Dyke Parunak (1997) and Van Dyke Parunak and Brueckner (2001) argue that the micro-level dissipation effects at the micro-level of pheromones create enough entropy for the macro-level system to coordinate (reduce entropy) without violating the second law of thermodynamics.

Another popular bio-inspired approach for the coordination of agents in manufacturing and logistics contexts is the *bee-foraging* approach. Here agents communicate not through pheromones, but through message passing. Bartholdi et al. (1993) were the first to analytically investigate the foraging behavior, showing that such coordination approach ensures an (otherwise inexistent) upper-bound on the loss of performance of distributed resource allocation as compared to an optimal (centrally determined) allocation. Quijano and Passino (2010) investigate the emergent behavior of *multiple* hives competing against each other as a simplified model for resource allocation problems between selfish agents, showing that the resulting allocations are NEs and thus foraging could provide global information. Scholz-Reiter et al. (2008) demonstrate the applicability of the bee-foraging method, as an alternative to pheromone-based approaches in manufacturing settings and demonstrate it's ability to respond to unexpected events (machine failures).

Where agents tend to follow past decisions, the process becomes *self-catalytic*, i.e. popular decision alternatives are reinforced (Kassabalidis et al. 2001; Marzo Serugendo et al. 2004). It is, therefore, important to consider the trade-off between the inclusion of past information and responsiveness to recent changes in the environment. To this end, the system may be partly filled with purely exploratory agents that are not bound by the accumulated social information (c.f. Dechaume-Moncharmont et al. 2005). When using a stigmergy approach, the pheromone evaporation rate presents another handle to ensure that the system state perceived by the agents does not overly incorporate outdated information (c.f. Armbruster et al. 2006). Through these measures a more continuous design space (as indicated in the classification model) can be achieved.

Creating Plans for the Future

Where past system behavior is assumed to carry insufficient descriptive power to prescribe future behavior, it may be advisable to generate a plan for the future (almost) at runtime. Such plans can then either bind (c.f. Wong et al. 2006)[45], direct/advice (c.f. Ouelhadj and Petrovic 2009) agents in their decision-making over the planning horizon, or merely enhance the agents' perceptions of their surroundings (Zambrano Rey 2014, Ch. 2). In the frameworks for myopia control of van der Vecht et al. (2007) and Zambrano Rey et al. (2013) this constitutes the difference between influence by belief alteration (through communication) and goal/task determination.

The results reviewed here generally require enhanced (as compared to the distributed "baseline" scenario) processing power to perform the planning tasks. It is not surprising that additional computational prowess in itself can reduce myopic behavior: Argoneto et al. (2008, p. 44) notes that "Sophisticated individual reasoning can increase MAS coherence because each individual agent can reason about non-local effects of local actions".[46] On the other hand, an unexpected relationship seems to hold with respect to dynamic complexity: Moreira et al. (2004) found evidence that while sophisticated local heuristics often show little robustness toward random noise, simpler strategies may not only be more robust, but gain efficiency with increasing random noise.

Beyond the degree of intervention with agent autonomy discussed above, a differentiation along two dimensions seems sensible to provide more clarity and structure to the approaches reviewed here. First, planning between agents can result either in a *joint plan* where agents collectively agree on one plan that encompasses all agents, or form plans individually, thereby collectively forming what is known as a *multi-plan* (Ossowski and Omicini 2002). The second dimension considers the population of agents and whether it is homogeneous (all agents possess the same set of capabilities), or heterogeneous (in particular: an additional decision-making agent or service providing entity exists) (c.f. Section 2.1.2).[47]

Joint Plan Generation:

Not surprisingly, approaches that set out to create a joint plan generally require the presence of a centralized entity to assemble and provide a global picture of the system plant and/or to calculate/assess global schedules.

Most obtrusive from the viewpoint of the individual agent's autonomy is the (partial) resolution of scheduling conflicts through a central decision making entity (Zambrano Rey 2014, Ch. 2.2). For approaches in this category, similar rationales apply as discussed under Section 3.4. Examples where optimization approaches are used *online* to resolve

[45] Similar to the impact of offline scheduling, discussed in Section 3.4.

[46] Note that "non-local effects of local actions" match precisely the definition of (social) myopia by Zambrano Rey et al. (2013).

[47] This second line of discrimination is similarly also drawn up by Zambrano Rey (2014), who discriminates between fully heterarchical and semi-heterarchical control systems. However, this thesis will also consider approaches where a central entity provides global information or other decision-supporting (but no decision-making) capabilities as instances of heterogeneous MAS.

upcoming scheduling conflicts include the works of Berger et al. (2010) and Leitão and Restivo (2006). Online simulations can likewise be applied — not to form a new joint plan, but to assess the performance of possible scenarios, the impact of variability on plan performance, etc. (Monostori et al. 2010). Examples include the simulation capabilities assigned to the staff holon of a HMS variation by Cardin and Castagna (2009) and the evaluation of decision-making alternatives through a petri-net based representation of the manufacturing system, as proposed by Leitão et al. (2010).

Multi-Plan Generation:

In the absence of a global entity with access to a global system state, local agents may improve their decision-making by increasing the "interval of time at which performance is measured" (Herbon et al. 2004, p. 689), generally called the *planning horizon*. By extending the planning horizon, agents can include more information in their decision-making; however, they are also more likely to make decisions based on erroneous or outdated information (ibid.). Changing the planning horizon is a frequently used structural parameter to alter the behavior of distributed PPC systems (Brennan and Norrie 2003). Comparisons of the performance of distributed control architectures as a function of the planning horizon have been undertaken by Brennan (2000), Brennan and Norrie (1999), Mönch and Driessel (2005), and Rogers and Brennan (1997). An analytical investigation and approach to define the "effective information horizon" is provided by Herbon et al. (2004). Estrada and Vargas-Estrada (2013) investigate the impact of increased information horizon in continuous coordination problems and found that accounting for higher-degree neighbors improves the adjustment speed across network architectures. In Section 5.3.4, upper bounds for the price of anarchy are reviewed in model classes that allow analytical investigation. Upper bounds for situations with local information only significantly exceed the upper bounds applying when global information is available, giving further evidence that the impact of myopic behavior (at least in the worst-case scenario) can be reduced through global information (Gairing et al. 2008; Papadimitriou and Valiant 2010).

One way through which agents may attain and keep up-to-date the necessary information about the remote parts of the system plant is through the use of *forward agents*. Forward agents were first explored in the literature on routing in communication networks: Both Di Caro and Dorigo (1997) and Heusse et al. (1998) describe routing algorithms, where, in addition to actual payload packages, routing agents flow through the network. They are emitted by the nodes in the communication network and aimed at analyzing the congestion of upcoming (communication network) links to compare path alternatives. Network nodes can then update their forwarding tables based on the information collected by the routing agents. Heusse et al. (1998) observe that this process creates an "autocatalytic effect" comparable to stigmergy.

The concept has attracted the attention of PPC designers, in particular in the domain of HMS: Zambrano Rey et al. (2012) introduce *adjunct product holons* that are dispatched by a product holon (the agent representing a product) to gather information about the current and expected future state of resources and inform the decision-making of the product holon. The *exploring ants* emitted by Karuna et al. (2006) serve a purpose similar to the routing agents discussed above. However, Karuna et al. (2006) extend the scope of functions that such emitted agents can serve by also introducing *intention ants* that, based on the information collected by their exploring cousins, express the intention of a product holon to reserve time at a certain machine, thereby creating a schedule. These ideas lead to the concept of *Delegate MAS* (D-MAS) being proposed by Holvoet et al. (2009) and Verstraete et al. (2008a), as an extension to the PROSA reference architecture. D-MAS provides a framework for product agents to explore the system state and reduce conflicts among their individual plans by expressing and sharing their intentions through emitted delegates. Weyns et al. (2007) use a variation of delegate MAS to improve the coordination between autonomously guided vehicles (AGVs).

Disciplines outside PPC have discussed incentives as a way of influencing multi-plan formation, with the aim of moving it closer to social optima. In the discussion of road- and communication-network users modeled as non-atomic flows on congestion networks (c.f. also Chapter 5), *tolls/pricing* are discussed as a mean of driving individual agent planning (agents are assumed to decide on their complete routing at the start of the journey) toward less myopic behavior. Cocchi et al. (1993) study the interplay between service disciplines and pricing policies, finding that pricing policies are inevitable for maximizing system performance. In their discussion of data packet routing in TCP networks, Gibbens and Kelly (1999) propose to calculate each resource *shadow price*, the marginal increase in cost (here calculated as packages lost) for a marginal increase in load, and to charge network users (senders of data packets) with this shadow price, to create awareness about the resource implications of their decisions. This finding is in line with the results derived from the analytical calculation of socially optimal flow distributions in congestion networks (c.f. Section 5.7), thus implying that the globally optimal flow-distribution may be reached, when agents are charged with the *marginal cost* (on all agents) when choosing a given route (c.f. e.g. Roughgarden 2005, Remark 2.4.7). The possibility of using incentives (positive or negative) to enhance coordination has also been discussed in the context of organization theory, where Siggelkow and Rivkin (2005) find in minimal model experiments that firm-wide incentives can substitute organizational hierarchy as a coordination measure. While toll/incentive systems are less technically elaborate as compared to forward agents and related concepts, setting and enforcing them does require a central entity with sufficient information and authority to calculate, set, and enforce them.

While incentives influence the way in which agents perceive their environment and decision alternatives, central entities may also influence the coordination process between agents: Lim et al. (2009) present an example where a global optimization problem adjusts the bidding power of agents when negotiating production schedules to improve overall performance.

Without engaging in any centralized decision-making, central service providers may also support multi-plan generation and the quality of the attained plans by providing simulation capacities throughout the system. This gives agents the capacity to evaluate their individual decision alternatives by means of a (global) simulation model. Such approach was among the first myopia control approaches investigated in literature: Already in Duffie and Prabhu (1994), a “look ahead cooperative scheduling” approach was proposed in which a system part can request the start of a simulation of the entire system for a given time period to analyze decision alternatives. Similar ideas were implemented, e.g. by Cardin and Castagna (2009) in the context of HMS and Rolón and Martínez (2012).

Finally, a combination of joint and multi-plan generation is imaginable: Leitão and Restivo (2008) present an extension of the ADACOR framework to incorporate online plan-generation capabilities in which multiple operational holons first create new local schedules in a distributed manner. Supervisor holons optimize and spread these local plans across the systems.

Stylized Model Experiment

The discussion about the impact of the agent's planning horizon on their performance in a multi-plan generation environment can be enhanced through the stylized model introduced in Section 3.3.1.

Instead of varying the wiring probability between consecutive levels, now the weight given to observations from different levels is varied, generally giving more weight to the immediate next level(s) and decreasing weights for the observerd WIP on levels further away. As in previous experiments, agents consider all possible paths across the next four levels in their decision-making. However, where previously the WIP level in machines was simply added up across all levels, these WIP levels are now weighted with q^{i-1}, $i \in \{1, \ldots, 4\}$, where i is the difference between the level of the considered server and the current level.

By varying $q \in [0, 1]$, the parameter changes the relative weight given to observations further into the future. For $q = 1$, the original situation is maintained, where the WIP levels are simply summed up. For $q = 0$, all information but the WIP of the immediately following level is disregarded, effectively reducing the information horizon to a single level.

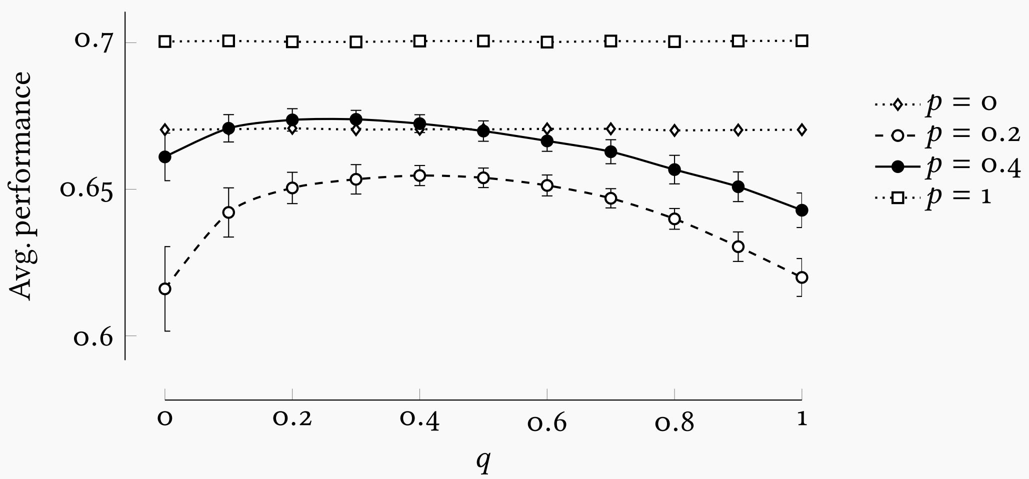

Figure 3.6: Measured performance as function of network WIP devaluation factor for different connection probabilities p. Statistics calculated over 100 simulation runs.

Figure 3.6 shows the result of the experiment setup. For medium levels of p (i.e. scenarios associated with high complexity, c.f. Fig. 3.4), the results clearly indicate a curvilinear relationship of performance as a function of the WIP devaluation applied, thus confirming the hypothesis e.g. of Fig. 1.4. The flat performance curve for $p = 0$ comes as no surprise: With no flexibility available to the agents, more informed decision making cannot have any impact on the performance. For $p = 1$, the benefit of increased flexibility becomes visible as compared to $p = 0$. The flat performance curve is also readily explained: Since after every level, every server at the next level is accessible, the WIP levels on following levels (or their weighting) have no implication on the choice of server at the immediately next level.
The results, especially the peak in performance for medium values of q, underline the ambivalent nature of larger information horizons: While they do give agents access to larger parts of the system, information about distal parts of the system does not seem to improve decision-making. A decision policy that values the near spatial neighborhood higher then the "larger picture", literally meeting the definition of higher levels of myopia (c.f. Section 3.1.1), outperforms policies where agents give more emphasis to remote information. This underlines that a reduction in myopic behavior is not a goal in itself and yet can be used as a guideline to find good system designs that exhibit good performance (c.f. Section 3.1.3).

3.5.2 Nervousness Alleviation

The discussion so far has covered large parts of the design decisions commonly associated with PPC (c.f. Section 2.1.3). In this and the following subsection, the thesis introspects the decision-making *given* the constraints and information provided through previous system design and scheduling decisions. Two important properties of agents' decision-making functions are discussed: their willingness to change a previously prepared plan (a previously made commitment) and (in Section 3.5.3) the perception of utility used in determining the "best" plan or commitment.

Where agents are free to change their decisions, system nervousness may arise in the agent society, since the re-consideration of one agent may trigger re-assessment and likewise plan changes by it's neighbors. This dynamics can lead to deadlock situations, where agents get locked in a cycle of constant response to each other (c.f. Fabiunke 1999, and Section 4.2.4) or at least hinder convergence. Barbosa et al. (2015, p. 106) state that "One of the main problems that could appear in self-organized distributed autonomous systems is chaotic system behavior where holons are continually changing their behavior or enter a continuous cycle of constant evolution/adaptation". Nervousness alleviation measures are designed to suppress such dynamics. They may be likened then to shock absorbers in cars that are meant to dampen the impact of and response to abrupt changes in the environment by the controlled system (ibid.).

Hogg and Huberman (1991) present a meta-heuristic to combat system nervousness in distributed systems. Using an example of resource allocation, they show that a combination of agent strategy selection (increasing the share of agents with good performing strategies) and changes to strategy payoffs can reduce system chaos and yield a performance of the distributed system "which in some cases approaches the optimal one that would be obtained by an omniscient central controller" (ibid., p. 1331).

Several methods of nervousness alleviation have also been tried within PPC settings: Van Dyke Parunak et al. (2003) claim that MAS performance can be improved if individual agent activity is reduced. To achieve that, they suggest agents maintain an internal pheromone level (similar to stigmergy approaches discussed in Section 3.5.1), which increases when the agent observes its own behavior as "hyperactive" and which reduces over time, providing a form of local memory for the agent to fine-tune its behavior. Karuna et al. (2005, 2006) define so-called "socially acceptable behavior" for agents that can require them to either (1) wait before changing their decision, even when they observe a new, preferable decision alternative, (2) deny a given proportion of commitment changes independent of expected improvement, or (3) limit the number of changes per time (Karuna et al. 2006). This dampening approach is shown to avoid the periodic switching of agents between, e.g. parallel resources and similar alternating agent behavior (ibid.). Barbosa et al. (2015) propose for the ADACOR² reference model a "system stabilizer", based on a PID controller to provide a parameterizable way of defining agent's response to environmental changes. The experiments show an improvement with ADACOR² over strictly hierarchical, heterarchical, and ADACOR (first evolution) controlled systems, though the effect of nervousness alleviation is not investigated separately.

Nervousness-reducing measures will also be discussed and applied in the context of the minimal model, investigated in Chapter 4, where similar ideas to suppress agent hyperactivity are applied (c.f. e.g. the works of Fitzpatrick and Meertens (2001) and Hadzhiev et al. (2009) and discussions in Sections 4.2.4 and 4.3.2).

3.5.3 Competition Approach

Already, Section 2.1.2 has discussed autonomy (implying selfishness) of the agents on the one hand and the requirement for system level coordination on the other hand as a central design trade-off within MAS. Similarly, Ahrens (1996) define two different basic principles for coordination among decision-making entities: cooperation and competition. Applying the measures introduced in this subsection, the designer of distributed PPC systems can shift agent behavior toward more inter-agent cooperation.

Several authors (Alshabi et al. 2007; Barber et al. 2001; Decker 1987; Falcone and Castelfranchi 2001; Moulin and Chaib-Draa 1996) have referred to the idea of a continuum of agent behaviors between completely selfish and completely complying with the benevolence assumption. Kalenka and Jennings (1999) point to a spectrum of possible decision-making functions that implement such continuum ordered by increased "social disposition". Alshabi et al. (2007) and Decker (1987) conceptualize a continuum of agent behaviors between autonomous and cooperative, where completely cooperative agents could even change their goals, if needed, to meet the need of other agents in the MAS.

The primary concern with selfishly acting agents is the prioritization of private utility over the social good (Falcone and Castelfranchi 2001; Vánca 2014). Jennings and Campos (1997) argue that "socially responsible agents", i.e. selfish agents, willing to provide some consideration to the greater good of the overall system are a promising approach to attain coordination and hence coherence. Similar conclusions and concepts have been presented by other authors (e.g. Alshabi et al. 2007; Ossowski 1999; Ossowski and García-Serrano 1999). Kalenka and Jennings (1999) hypothesize that through (limited) social disposition, the social performance will increase with little negative effect to individual performance. This observation seems to be in sync with the findings of Decker (1987), who find most real systems, when placed on a scale between selfish and benevolent, are benevolent only to some small degree.

Two paradigms for such "social" coordination among agents are distinguished by Ossowski and García-Serrano (1999): (1) adjusting the behavior directly by altering the agent's concept of rationality or (2) by expressing the agents' interdependencies structurally or through inclusion in the utility function. An example of the first approach is the concept of "social rationality" introduced by (Jennings and Campos 1997), where agent can choose certain plan if it is good for itself or society.

As within the MAS domain, changing the decision-making priorities of agents (i.e. modulating the degree of selfishness) has been a frequently discussed idea within the PPC community: Already Hatvany (1985) suggested "cooperating heterarchies" to avoid the "primitive anarchy" the he assumed would follow from complete autonomy. Likewise, already the initial definition of HMS sought to attribute agents with both autonomy *and* cooperativeness (distinguishing them from "free agents") (Christensen 1994). Boccalatte et al. (2004) devise an extension to the ContractNet protocol in which agents representing jobs become more likely to refrain from requesting machine time (hence giving capacity to other jobs) when their slack[48] is high. Zambrano Rey et al. (2013) showed in a simulation study that altering the decision function of PPC agents toward more benevolence (using the "Competition Approach Paradigm", proposed by Fedoruk and Denzinger (2006)) can increase schedule reliability in a simulated HMS.

3.6 The Resulting Classification Model

Figure 3.7 shows the resulting classification model for myopia reduction during manufacturing system design, scheduling, and control. Within each decision step introduced in Section 3.2.3, the dimensions (rows) indicate decisions through which the degree of myopia exhibited by the resulting system can be affected. Within every dimension, general decision alternatives are paraphrased. The gray arrows associate with every decision alternative a degree of myopia reduction. Decision alternatives located at the wide base of an arrow can be associated with little reduction to the impact of myopic behavior. On the other hand, a design decision that corresponds to an arrow tip can be expected to

[48] Difference between the amount of time to the planned finishing date and the amount of workcontent (in hours) still to be completed.

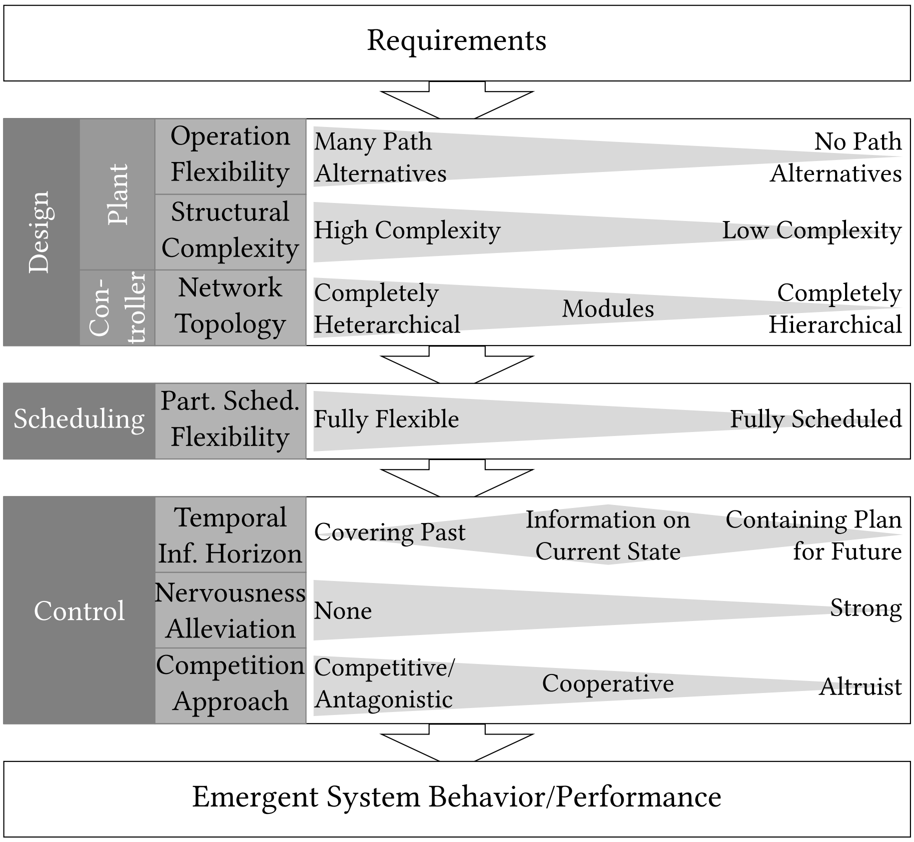

Figure 3.7: Proposed classification for the control of myopic decision-making in PPC. The impact of myopia is reduced where design decisions tend toward the pointed end of the gray arrows. The more the design decision is located towards the base of the arrow, the higher the expected impact of myopic decision making.

significantly reduce the impact of myopic behavior, given the literature reviewed in this chapter.

3.7 Discussion & Limitations

This chapter has addressed research question Q_1, finding that there exists a substantial body of literature across domains which implies that a set of design decisions can influence the impact of myopic decision-making on the performance of distributed PPC systems by either addressing the triggers of myopic decision-making or by limiting the decision space and thereby the potential negative impact of myopic decision-making.

While not part of this dissertation, the classification model presented here mainly as the aggregated result of a literature review has already passed initial validation as an aid to PPC system designers. In Wang et al. (2017), the authors map four popular distributed PPC approaches, two each from the stream of PULL production and agent-based manufacturing control, to the classification model, positioning each approach along all design decisions of the classification model. By observing both intra-group similarities and inter-group differences, the authors substantiate the idea that the presented classification model can be used to classify and compare distributed PPC approaches. They find particular "fingerprints" in terms of where and how the impact of myopic behavior is constrained. The results also indicate that of the distributed PPC approaches that have found recognition in science and practice, in fact all already apply measures to reduce the impact of myopic behavior in one form or the other. This observation gives further evidence of the wish of manufacturing system designers, to combine the advantages of distributed and hierarchical PPC or at least constrain the impact of myopic behavior.

The chapter has highlighted a transdisciplinary effort to understand and design complex systems from agent-level interactions. It has phrased efforts across multiple disciplines in terms of a common goal (reducing myopic behavior and/or the impact it has on system performance) and classified them according to a relatively small set of design decisions. Through this, it can help not only researchers in the domain of distributed and agent-based PPC, but in related scientific fields as well to understand their work in a broader context. While the design decisions presented in Fig. 3.1 are tailored to the domain of manufacturing system design and control, the basic hierarchical sequence of decisions arguably applies to a wide range of engineered and social systems. This makes the underlying approach to understand designing distributed systems as an attempt to limit and shape the impact of myopic decision-making, applicable to other design domains as well (the design of organizations and MAS in other domains are examples).

For researchers who work on hybrid PPC, the mapping of existing distributed and hybrid PPC approaches to the classification model (as done in ibid.) can help to identify "empty spots", so far unexplored regions, in the design space. Thus, the classification model may point to promising fields of research that can extend our understanding of the design space in a meaningful fashion.

For practitioners in the domain of hybrid PPC, the presented classification model in its short visual depiction (Fig. 3.7), as well as the discussion and motivation elaborated in Sections 3.3 to 3.5, provides a wealth of interdisciplinary design inspirations to explore the performance potential of their manufacturing system (plant and controller) under hybrid PPC. By pointing at design dimensions on different time-scales of decision-making, the classification provided here can provide input during the design of new manufacturing systems as well as during day-to-day operations in which scheduling and control decisions have to be taken.

Limitations

There are three main limitations to the model presented here:

First, the model neglects interrelationships between the design dimensions. The design dimensions for distributed systems are not independent, but may require, support, or impair each other (already noted by Gasser 1991). This was most prominently noted in the discussion of "hierarchy" (Section 3.3.2). This implies that the internal consistency of the model cannot be ensured since not all relationships between concepts have explicitly been considered (Wacker 1998).

Second, the manufacturing system requirements, while mentioned as an input, are insufficiently represented in the model. The design dimensions are discussed w.r.t. their impact on myopic behavior, but the degree of myopic behavior tolerable, even desirable, in a given PPC setting, has to depend on the requirements posed on that system. To the model's defense, this is not an easy feat! The aim to analytically understand interdependencies among the traits of the problem description and the (likely) best control approach has kept production researchers busy for decades. In 1997, Rogers and Brennan put this question to the spotlight: "Of primary importance is the fundamental question concerning the choice of control architecture: *i.e.* is it possible to determine whether a specific control architecture is appropriate for solving a given manufacturing system control problem?" (Rogers and Brennan 1997, p. 881). However, to the best of the author's knowledge, no contributions beyond "rules of thumb" have been made so far.

Finally, the model dimensions are not operationalized or quantified — thus hindering more rigorous model validation beyond the attempts reported by Wang et al. (2017). Most recommendations remain qualitative in nature. The chapter has pointed at forms of quantification where possible (e.g. in the stylized model used to assess system complexity in Section 3.3.1 or for the assessment of partial schedules in Section 3.4). However, for most design dimensions and myopic behavior, there is no academic consensus as to how to measure them.

3.8 Directions for Future Research

As a heuristic framework, the benefit to researchers and practitioners can only be assessed through repeated application. The steps taken in Wang et al. (2017) can provide an early indicator of the capacity of the presented model to describe existing distributed PPC approaches in terms of their measures to reduce myopic decision-making and the impact thereof. They cannot replace practical tests, where the classification model is applied to either design a manufacturing system with distributed/hybrid PPC "from the ground up" or where design dimensions proposed in the model are used to alter the degree of myopic behavior in existing PPC approaches. Observed improvements in performance would validate the hypothesis underlying this chapter (and in fact this thesis), that the purposeful positioning of a PPC system between the poles of fully distributed and hierarchical can improve performance (c.f. Section 2.3) and that the developed classification model, in particular can support the design of PPC approaches that come close(r) to the region of maximum performance.

To understand the mechanisms behind particular design dimensions in a better way, specialized analytical models will be necessary. The stylized example discussed in Sections 3.3.1 and 3.5.1 and the model discussed in Chapter 4 are examples of such analytical investigation.

Chapter Four

Balancing Hierarchical and Heterarchical Control Architectures — A Minimal Model

"We will be successful in engineering agent-based systems just to the degree that we understand the interplay between disorder and order."

Van Dyke Parunak and Brueckner (2001, p. 124)

Submission under preparation

The model and results presented in this chapter are currently prepared for publication.

Distributed PPC systems — and distributed control systems in general — have been described in antithesis to a traditional centralized controller as a configuration (a network) of agents collectively forming the system behavior through their individual actions. In this chapter, the impact of the structure of this agent network on the collective performance of agents will be investigated. In particular, the impact of hierarchy in the control network will be the subject of experimentation and analysis. From the various design decisions Chapter 3 has pointed to, through which the space between hierarchical and strictly distributed PPC may be described and bridged, the benefit of hierarchy was found to be widely assumed and discussed Section 3.3.2. Research question Q_2 hence especially asks about the role of hierarchy in the design of distributed PPC systems that attain the "optimal" performance. This research question is addressed here.

4.1 A Model to Analyze the Impact of Control Network Hierarchy

4.1.1 Model Requirements and Analysis Steps

This chapter then intends to analyze the impact of *hierarchy in the control network*, masking — as has Section 3.3.2 — the effect of network structure from other properties (like increased computational power or decision-making authority) often attributed to agents of a higher hierarchical level. To this end, a simple minimal model is necessary that can study the network structure among otherwise equal agents (homogeneous MAS). To *demonstrate* the impact of hierarchy, a measure of performance is necessary. To *understand* the role of hierarchy, the agents and their interaction should lend themselves to analytical investigation. As the results are to be transferable to PPC settings, the model should encapsulate a problem that resembles the scheduling problem (resource allocation and sequencing), faced by (distributed) PPC systems. The Graph Coloring Dynamics (GCD) model, to be introduced in Section 4.2, meets this criteria and will be used in this chapter.

Section 2.1.3 has already pointed to the common conception of PPC problems as RAS. The control of a RAS then has a two-fold objective (Reveliotis 2005, Ch. 2):

1. To ensure the correctness and inherent consistency of resource allocation and
2. to ensure the efficiency of resource allocation with respect to some performance measure.

The chapter will consider both objectives sequentially. Section 4.2 will introduce GCD as a model language, which is prone to model the distributed control of RASs, seeking a correct and consistent resource allocation. The first analytical investigations in this chapter (Sections 4.4 and 4.5) will focus exclusively on studying the impact of control network hierarchy on the ability of distributed control systems to find a correct and consistent resource allocation. Section 4.6 then presents a model extension of the GCD model to also consider the second objective of resource allocation control, namely performance optimization.

4.1.2 Hypotheses on the Role of Hierarchy — The Organization Theory Perspective

Before discussing the model in more detail, a set of hypotheses on the impact, role, and mechanisms of leadership in distributed control settings is derived in this section. To this end, this chapter turns to organization theory, in particular Complexity Leadership Theory (CLT) (c.f. Section 1.5.3). The hypotheses to be derived here extend the research questions developed in Section 1.4. Unlike them, the hypotheses represent testable statements that motivate the model development and the measures analyzed to identify "leadership" on behalf of agents higher in the hierarchy.

CLT is a suitable source of inspiration and hypothesis development in this context, since researchers in this domain — as will discussed below — are actively concerned with understanding the roles of "leaders" in complex systems and have developed a coherent grasp of theory that can be used to develop analyses for the GCD model.

Why Complex Leadership Theory

This thesis is not the first to suggest that management science in general (and CLT in particular) may provide useful input for the understanding of complex systems beyond the domain of management and organization research: Solow and Szmerekovsky (2006) suggest that the role of leadership, in particular, is a promising field where management science can actually contribute to the understanding of complex systems. They state:

> *"While many complex systems may initially arise from self-organization [...], as many of these systems evolve it is the emergence of some sort of centralized organization that allows for high levels of performance that might not have been achieved otherwise. [...] To the extent, then, that the behavior of certain complex systems is affected greatly by central organization, our understanding of the behavior of these systems should include the study of how central organization and leadership affect system performance."*
>
> — Solow and Szmerekovsky (2006, p. 52)

Above quote also defines the expectation from leaders in complex systems in particular: to aid the coordination process between system entities and to increase performance. In fact, we may take this as the very definition of leaders in such systems: Calvert (1992, p. 7) argues that "leaders are needed because of, and derive their powers and capabilities from, their ability to solve problems of coordination". Entities that can provide such capabilities, even in a homogeneous MAS as considered here, may rightfully be called leaders.

It should be noted that organization theory knows at least four different ways through which central "leaders" can affect organization behavior: *Authoritarian, motivational, cooperational, and passive-responsive* (Solow and Szmerekovsky 2006). This research will focus primarily on *cooperational* leadership, i.e. "the leader's role of achieving cooperation among the team members" (Solow and Leenawong 2003, p. 67), since it corresponds to the model assumptions that leaders attain their prominent role through their positions in the network architecture only.

In this chapter, indications of such "leadership" will be explored particularly in agents of a high central position in the coordination structure associated with a higher level of hierarchy. It should first be verified that the parallelism set out between agents in such network positions and "leaders" in an organization (c.f. Section 4.3.1) is valid. This assumption can indeed be confirmed in the organization theory literature: The combination of a set of links among agents and tasks (structure) and the rules and procedures (dynamic) constitutes the defining characteristics of an organization (Carley and Gasser 1999, Ch. 7.2.2). Carley and Ren (2001) and Carley et al. (2000) then find in a model-based investigation of organization structures during armed forces' operations that "certain agents in important, central network positions were likely to develop certain leader-like characteristics" (Hazy 2007, p. 396). Ibarra (1993) finds central network positions to have stronger correlation with the exercise of power (measured as the degree of involvement in administrative processes) than the formal rank (for innovation processes, the effect of both indicators is indistinguishable). Likewise, central network positions seem to support the work of leaders: High centrality metrics (here: the betweenness centrality, c.f. Costa et al. 2007) of managers of open source projects were found to support the technical success (measured in the number of code submits) by Grewal et al. (2006).[49] Both CAs and

[49] Results for "commercial success", measured in number of downloads, were more nuanced.

ABMs were also suggested as one potential modeling class to understand the emergent patterns and principles of emergence in organizations by Lichtenstein (2007).

Derivation of Hypotheses

As already discussed in Sections 1.5.3 and 2.3, the assumption of a duality between too much and too little hierarchical oversight in organizations has accompanied organization theory since at least the 1930s (Thomas et al. 2005b). A significant contribution was made by March (1991), who discussed the management of national subsidiaries in large, multinational firms as being driven by both *exploration* and *exploitation*, with optimal performance being attained at the mix of the two extremes. More recent research also suggests that an alternation between exploration and exploitation could increase the performance (March 1999; McKelvey 2004). This notion is supported by the minimal model experiments of March (1991) and Siggelkow and Levinthal (2003). In particular, March (1991) finds that a combination of slow learning in rapidly changing environments results in poor performance, while quickly learning organizations may actually perform better under moderate amounts of change in the environment.

Exploration, the self-dependent behavior of employees/subsidiaries to try new decision alternatives, requires “creativity” (Coleman 1999, p. 35 f.). Based on the already cited work of Langton (1990), authors such as Coleman (1999) and Stacey (1993) suggest that companies should be placed on the *edge of chaos* to attain maximum *creativity*. It is thereby assumed that the organizational structure (for example, the presence of hierarchy) plays a central role in positioning an organization within this duality (Bartlett and Ghoshal 1987; Siggelkow and Levinthal 2003), as March (1999, p. 214) states[50]: “Organizational structure can be used to strengthen exploration by undermining the effectiveness of exploitation”.

This discussion leads to the first hypothesis, which picks up in a testable fashion the idea behind research question Q_2:

> H_1*: In complex systems, a balance between centralized and distributed control leads to the highest performance.*

Another reason why CLT appears to be a promising framework to discuss the results of GCD experiments is that it shares the assumption of homogeneity among organization members. Unlike in classical organization theory, Uhl-Bien et al. (2007) argue that leadership in CASs is not a process of authority, but an emergent dynamic amongst agents. Consequently, leaders are not *controlling*, but rather *enabling* the evolution of the system (Marion and Uhl-Bien 2001). Uhl-Bien et al. (2007) phrase this form of leadership *adaptive leadership* and define it as “emergent change behaviors under conditions of interaction, interdependence, asymmetrical information, complex network dynamics, and tension” (ibid., p. 309).

In this context, CLT is particularly interested in understanding how and why leadership in CASs works. As in the discussion of CAs, the notion of *order* and its creation are

[50] referring to Hedberg et al. 1976.

central: McKelvey (2004, p. 6) concludes that complexity theory is really the "science of new order creation". In CAS, organization theorists believe that individual agents act to reduce *tension* (Lichtenstein et al. 2006), which is seen as an important *condition* to enable emergent leadership and innovation (Uhl-Bien and Marion 2009). In the eyes of Uhl-Bien et al. (2007, p. 306), adaptive leadership "originates in struggles among agents and groups over conflicting needs, ideas, or preferences" and can hence arise from information asymmetry. Tensions are challenges to an agent's personal knowledge base (Lichtenstein et al. 2006, 2007). They can lead to a realignment of the agent's cognitive map and through this generate information (Lichtenstein et al. 2007; Uhl-Bien et al. 2007). Lichtenstein et al. (2006, p. 5) conclude that "Therein lay the seeds of adaptive leadership: Agent interactions can generate tension through which novel information can emerge; when those new ideas lead to positive change, adaptive leadership has occurred." The following hypothesis is thus stated has, to the best of the author's knowledge, not been analytically explored before.

> *H_2: Leaders use information asymmetry and facilitate the coordination process; they enable conflict solving among agents by interaction, exchange of information and emergence.*

Computerized model experiments on the impact of leaders on team performance (Solow et al. 2005) focus on the role of leaders as motivators for their team; they do not discuss the power of leaders that arises strictly from their respective positions in the network architecture.

This implies that leadership is a process to be observed over time, rather than being immediately visible from a static system snapshot. Lichtenstein et al. (2006) suggest that leadership occurs in response to events and action cues; they also put forward ideas for approaches to detect emerged leadership in CAS by measuring the dynamics of organization over time. They suggest to model "these data in ways that highlight their longitudinal and relational qualities" and to analyze "these data in terms of their relational qualities and longitudinal dynamics" (ibid., p. 5).

Events are important in the study of self-organized systems: It is a typical property of self-organized criticality that small inputs of new information can lead to massive changes in beliefs across system entities, resembling the "avalanches" observed in the sandpile model of Bak et al. (1987) (Glinton et al. 2010). If leaders can facilitate the spread of this information (tension) in the organization, they may (a) create or (b) facilitate avalanches of agent re-alignments that can successfully resolve conflicts. The importance of (unplanned) events in the context of CAS leadership is also stressed by Vicari et al. (1996, p. 189), who say about the generation of knowledge through experimentation: "Leaps in the knowledge development of a company typically stem from events that the firm has neither planned or hypothesized". In fact, some authors in the domain of organization theory have gone beyond the acknowledgment of the importance of events and actively called for leaders to generate such events that would shake up established patterns: Marion and Uhl-Bien (2001, p. 406) calls for leaders in CAS to "spawn emergent behavior and creative surprises rather than to specify and control organizational activities", and Plowman et al. (2007) in their empirical study of a local congregation come to a similar conclusion: "Our findings show that as enablers, leaders disrupt existing patterns of behavior, encourage novelty, and make sense of emerging events for others" (Plowman et al. 2007, p. 341).

This idea is picked up in Hypothesis H_3.

> H_3*: Leaders disrupt existing patterns of behavior.*

Finally, in Section 2.3, the results of experiments in the domain of organization theory were reviewed, concerning the relationship between environmental conditions and the best-suited organizational structure (c.f. e.g. Fig. 2.3). It was found that more heterarchical and decentralized architectures are better suited for turbulent environments requiring speedy improvement.

This leads to the following hypothesis about the interdependence among the organization structure and the optimization time horizon that was likewise expressed for the domain of PPC (Trentesaux 2009, c.f. also Fig. 2.2):

> H_4*: The optimal balance of centralized and decentralization control architecture is a function of the optimization time horizon considered.*

Before the hypotheses developed here can be tested in Sections 4.4 and 4.5, the next sections introduce the model language (Section 4.2) and the experiment setup (Section 4.3) used in this chapter.

4.2 Model: Graph Coloring Dynamics

This section introduces GCD as a minimal model of resource allocation that lends itself to analytical investigation and to study the impact of network structure on system performance.

4.2.1 Constraint Satisfaction Problems as a Minimal Model for Coordination Problems in Resource Allocation

The first objective for the control of RASs (ensuring correctness and consistency of the resource allocation) can abstractly be described as a Constraint Satisfaction Problem (CSP). CSPs describe the problem of "finding *values* for problem *variables* subject to *constraints* that specify which combinations of values are allowed" (Freuder 1995, pp. 103 f.). Across disciplines, many planning problems can naturally be modeled as such CPSs (Do and Kambhampati 2001). Where control over the decision variables in a CSP is separated across multiple agents, we arrive at a Distributed Constraint Satisfaction Problem (DCSP) (Fabiunke 1999; Solotorevsky et al. 1996; Yokoo et al. 1998; Zhang et al. 2002). Solotorevsky et al. (1996) describe a DCSP as a set of (interconnected) constraint networks where every such network is solved by a separate agent. As such, DCSPs provide a minimal model to which multiple MAS application scenarios, such as the recognition problem, multi-agent truth maintenance, and (critically) allocation and scheduling problems within MAS, can be mapped (Yokoo 2001, Ch. 2.3; c.f. Solotorevsky and Gudes 1997, for an application of DCSP to the nurse scheduling problem).

4.2.2 Graph Coloring as a Minimal Model for Distributed Constraint Satisfaction Problems

In this chapter, a prototypical CSP, Graph Coloring (GC) (Fitzpatrick and Meertens 2002; Freuder 1995; Yokoo 2001, Ch. 1.2) is considered. The GC problem (in its most basic form) is defined on an undirected Graph $G = (V, E)$, where every node $v \in V$ has to be assigned to one of k states called "colors", in such a way that no two connected nodes share the same color (Jensen and Toft 1995, Ch. 1.1). Given a graph G, the smallest number of colors necessary to solve the graph coloring problem is called the graph's *chromatic number* $\chi(G)$.

Determining if a given graph can be colored with a given number of colors is a remarkably challenging combinatorial problem that has defied all attempts for efficient algorithmic treatment. It is part of the famous 21 $\mathcal{NP}$-complete problems listed by Karp (1972) and has attracted much attention in the field of computer science and algorithmics. The underlying CSP nature of the GC problem, where edges represent constraints, has early invoked comparisons with scheduling problems (c.f. Section 2.1.3). GC formulations have hence repeatedly been applied to model general resource allocation problems, including timetabling, multiprocessor scheduling, etc.[51]. De Werra and Hertz (2015, p. 255) note that GC "may provide a natural tool for dealing with a variety of scheduling problems". When applied to scheduling, the graph's vertices represent objects to be scheduled and edges represent scheduling constraints. The change of color allows for a one-dimensional adjustment of each node to satisfy the constraints. The color could either represent a resource or server (as e.g. in Brueckner and Van Dyke Parunak 2006; Windt and Hütt 2010) in a set of multiple identical servers. Then sequencing on each machine is fixed (or externally determined, c.f. Section 4.6) and edges connect operations to be executed in the same time interval (the graph is then an interval graph, c.f. e.g. Halldórsson et al. 2003), or we consider operations on the same resource and every color represents a time window of fixed length (as e.g. in Fitzpatrick and Meertens 2001). In both cases, a higher χ relaxes the problem (for a given conflict graph) (Fitzpatrick and Meertens 2001; Windt and Hütt 2010). While the standard GC problem is too simple for most real-world scheduling applications (Hansen et al. 1997), extensions to the problem syntax are possible to account e.g. for operations of different lengths, include target functions such as makespan minimization, preemptive as well as non-preemptive scheduling, etc. (Halldórsson et al. 2003).

However, this thesis intentionally disregards the opportunities of available model extensions.[52] This thesis will consider the GC model as a minimal example of a flexible job-shop scheduling problem ($J, P|prec|F$; where P indicates the presence of parallel identical machines) with unit-length operations and no "technical constraints", precedence relationships between operations that are pre-determined (e.g. subsequent operations in a job) (c.f. Bongaerts et al. 2000). Solving the flexible job-shop scheduling problem (FJSSP)

[51] c.f. De Werra and Hertz 2015; Fitzpatrick and Meertens 2001, 2002; Halldórsson and Kortsarz 2004; Halldórsson et al. 2003; Hansen et al. 1997; Korst et al. 1994; Leighton 1979; Marx 2004; Myszkowski 2008; Windt and Hütt 2010.

[52] Section 4.6 will extend the model toward realistic performance measures, but not by means of a semantical change to the model.

(c.f. also Dauzère-Pérès et al. 1998) is a challenge also to centralized PPC systems and has been tackled by numerous authors by applying a variety of optimization meta-heuristics and algorithmic approaches.[53]

4.2.3 Graph Coloring Dynamics: Multi-Agent Graph-Coloring as a Cellular-Automaton

While most of the research in the GC problem has focused on OR methods to find or bound the chromatic number for certain families of graphs, the relatively new field of distributed graph coloring (or GCD) investigates the performance of distributed agents in solving the k-coloring problem, i.e. assigning each node v a color $c \in \Sigma_C$ ($|\Sigma_C| = k$), in such a way that no two adjacent nodes have the same color (c.f. Garg et al. 1996; Kuhn and Wattenhofer 2006).[54] The set of available colors is generally equal to or slightly larger than χ. The problem graph in this situation can be understood as an *influence graph* that depicts the interconnections between agents (Berger et al. 2010) through which agents may exchange information. The agents are characterized by their decision heuristic, as described in Section 4.3.2, and a one-dimensional internal state, the node color. In the following, $x_i(t) \in \Sigma_C$ will denote the color (internal state) of node/agent i at time t.

The GCD approach maps the dynamic processes of distributed decision making to a model with discrete time and states and hence firmly within the model language of CAs on graphs. CAs were developed (under a different name) in the 1940s by von Neumann and Ulam (c.f. von Neumann 1963). Their studies have been advanced significantly by the works of Wolfram (c.f. Packard and Wolfram 1985; Wolfram 1983, 1984). In a CA, the state of each cell (agent) at time $t + 1$ is a function of the state of the cell and its (appropriately defined) *neighborhood* at time t. These functions (called *rules* in the context of CAs), obviously have enormous impact on the dynamic evolution of the CA. Depending on the choice of update rule, CAs are capable of deterministic, chaotic, and complex behavior, even universal computation (Langton 1990) and are hence a suitable tool for the investigation of topics related to distributed computation (Crutchfield and Mitchell 1995). The analysis of CAs has contributed, e.g. to the understanding of design properties of metabolic networks (Marr et al. 2007), pattern formation in biological processes (presented in Deutsch and Dormann 2005, Ch. 3), particle mixing phenomena, and signal propagation (Moretti and Muñoz 2013). As Freuder (1995) shows, CAs are a feasible model language to study the GC problem.

Starting with the work Wolfram (1984), CAs have been studied in the context of complexity science: CAs are given as one example by Holland (1998, Ch. 7) for a more general class of CGPs that can exhibit emergent behavior (c.f. Section 1.5.1). CAs can also serve

[53] c.f. e.g. Brandimarte 1993; Dauzère-Pérès and Paulli 1997; Dauzère-Pérès et al. 1998; Fattahi et al. 2007; Hurink et al. 1994; see Ziaee 2014, for a recent review.

[54] It should be mentioned that there are other distributed solution approaches to the graph-coloring problem: Wu et al. (2011) present clustering as a distributed method of solving the graph-coloring problem. However, the algorithm and its parameters need excessive adaptation for every new network. Bio-inspired meta-heuristics are developed and tested in Bessedik et al. (2014).

as models for CASs in (Holland 2002). Discrete difference equations (that allow the explicit calculation of a system state at $t + 1$), such as CAs, have also attracted research in chaos theory because they do not require to solve intractable differential equations (Levy 2000). Though CAs are usually considered to "live" on a (one-dimensional) grid, their definition can be extended to undirected networks of arbitrary shape (c.f. Marr and Hütt 2009). As Marr and Hütt (ibid., p. 546) note, "Cellular automata (CA) on graphs in principle provide the possibility to monitor systematic changes of dynamics under a variation of network topology". The local definition of update rules makes them an appropriate tool to investigate the emergence of patterns and solutions. As Brueckner and Van Dyke Parunak (2006, p. 109) note for the GCD problem: "The agents in the distributed graph-coloring problem interact locally to solve a global problem without being explicitly aware of their joint task. Therefore, we consider the global solution to be emergent."

CAs also have a history of application for the investigation of distributed production-scheduling exercises. In Oliveira and Vidica (2012) and Seredynski and Zomaya (2002), a "dynamic neighborhood" is created to map the irregular structure of a scheduling graph into the regular structure, commonly assumed for CAs. In Swiecicka and Seredynski (2000), the precedence graph is considered as a tree-like structure (with directed edges indicating precedence relationships) from which a neighborhood is created by taking the r 'neighboring' cells in the graph to either side. All papers have in common that they take a two-stage approach in which a suitable ruleset in determined in an a priori teaching phase before the identified rule can be applied for scheduling purposes. Moreover, in all papers, the CAs is "only" responsible for finding a task allocation problem, not a sequencing of tasks. This second elementary problem of production-scheduling is solved using heuristics instead (called "Scheduling Policies" in the CA-scheduling literature (Carneiro and Oliveira 2013)). In allocation processes beyond production scheduling, GCD has been applied to packet switching (Yeo et al. 2002) and register allocation (Smith et al. 2004) tasks.

The GCD approach, therefore, combines CAs as a minimal model of distributed decision-making and self-organization and GC as a minimal model for scheduling and lends itself perfectly for a first stylized discussion of the interactions among agents in heterarchical PPC systems (Windt and Hütt 2010). Instead of extending the underlying problem definition, this thesis takes the GC problem as a minimal interdisciplinary model that allows to investigate the convergence process within communities of decision-making entities as a function of the number of agents, tightness of the scheduling constraints, and complexity of the scheduling problem.

When perceived as a (distributed) CSP, both GC and manufacturing control can be treated as classical coordination problems as abstractly defined by Mesarovic et al. (1970): Given a control problem $\mathcal{D}$, the predicate $\Pi(x, \mathcal{D})$ is true, if and only if x is a solution of $\mathcal{D}$:

$$\Pi(x, \mathcal{D}) \equiv x \text{ is a solution of } \mathcal{D} \tag{4.1}$$

The "overall control problem" (ibid., p. 263) then is to find a (set of) control(s) $m \in M$ that minimize a performance metric $g(m)$ over M.

Critically however, this chapter is not interested in entering the abrasive race of finding the "best" (distributed) control for a GCD problem. This research follows the argument of Hooker (1995, p. 33) that such "emphasis on competition is fundamentally anti-intellectual and does not build the sort of insight that in the long run is conducive to more effective algorithms. It tells us which algorithms are better, but not why". Instead, the research reported here tries to understand *if* and *why* certain algorithm design principles — here in particular the choice between hierarchical and heterarchical agent topology — significantly impact the heuristic performance. In other words: Is there a subset of M for which $g(m)$ is consistently better than the average? For a general discussion of solution strategies for DCSPs, the reader is referred to e.g. Fabiunke and Kock (2000), Vattani (2012), and Yokoo et al. (1998).

4.2.4 Previous Work

GCD has been discussed and investigated as a model for DCSPs before.

Pearl (1988, Ch. 4.1.1) discusses GCD as an example of *constraint propagation*, a conceptualization said to be feasible to study the dynamics of unsupervised parallelism and exploit interdependencies between tasks. In the constraint propagation view "links in the network should be treated as the only mechanisms that direct and propel the flow of data through the process of querying and updating beliefs" (ibid., p. 145). Through these communication lines agents can question the current beliefs of their neighbors and inform them of updates. These updates may in turn cue a change in belief, which "initiates a multidirectional propagation that will continue until equilibrium is reached" (ibid., p. 145).

The first analytical experiments on GCD (to the knowledge of the author) were performed by Fabiunke (1999), who investigates a distributed 2-coloring of a square grid network structure. Agents in each node of the grid apply a "min-conflict" rule, minimizing their locally observed conflicts (number of neighbors with identical color). Fabiunke finds that such purely deterministic min-conflict rule dynamics on the given simple network structure can be trapped in "deadlock" situations, with high conflict numbers, especially when cells update in parallel. To avoid these deadlocks, Fabiunke (ibid.) successfully tests two different approaches to avoid deadlocks, namely

- the inclusion of redundant constraints (i.e. shortcuts in the grid), although finding possible shortcuts can be difficult for real-world problems (ibid.) and
- giving each agent a probability to keep its original color instead of choosing the conflict-minimizing one.

A similar investigation was performed by Zhang et al. (2002), who discuss a GCD model as a minimal model of a Distributed Search Algorithm (DSA), using again mostly grid-like graphs. Just like Fabiunke (1999) did (implicitly), they vary the degree to which nodes make color choices in parallel by setting probabilities for color changes, given a possible improvement, no change, or even a possible deterioration of conflict density. They find that the node update probability together with the constrainedness of the GC problem, can be used to show phase-transition effects in terms of the system's ability to reduce the overall conflict count in the network.

Fitzpatrick and Meertens (2001, 2002) develop a similar approach by applying a min-conflict update rule and a probabilistic activation of nodes. The authors also explicitly discuss their GCD experiment as an abstraction of distributed and parallel resource allocation processes on a random graph of varying edge density. Central to their investigations is the evolution of the number of conflicts in the graph as a function of the activation probability and graph density. The authors also propose to use the relative number of color changes per round as a measure for the communication cost, observing a sharp drop in color change activity early in the solution process (Fitzpatrick and Meertens 2002).

Brueckner and Van Dyke Parunak (2006), building on the mentioned work by Fitzpatrick and Meertens[55] investigate a model of GCD where communication between agents is subject to delays and nodes may "fail" (randomly change their color in every round).[56] Again considering the dynamics on a random graph and a decision heuristic oriented at conflict minimization, Brueckner and Van Dyke Parunak perform parameter sweeps over a variety of parameters governing the network structure and dynamics. They find distinct regions where the DCSP is solved quickly "either because the agent population is sufficiently intelligent to settle on a low-conflict solution, or because the problem is so easy that almost any randomly selected configuration results in a low degree of conflict" (ibid., p. 118). High information delay and high rates of random color choice can lead to "thrashing" behavior, where the system fails to reduce conflicts due to erratic decision-making or decision-making based on outdated information.

An essential motivation for the investigation in this chapter goes back to the work of Kearns et al. (2006), who were the first to investigate the impact of non-random structural anomalies in the graph. In human-subject experiments, they investigate the impact of the graph structure on the time required by human-subject networks to solve a 2-coloring problem. This research stream was further developed in Enemark et al. (2011), Enemark et al. (2014), Judd et al. (2010), McCubbins et al. (2009), and Shore et al. (2015), who investigate more sophisticated graph structures and also confirmed the finding by Kearns et al. (2006) that more links actually made the coloring problem easier (Judd et al. 2010). The experiments also gave rise to renewed computational investigation of the problem by Hadzhiev et al. (2009), characterizing the network dynamics and investigating the impact of shortcuts in a ring graph on solution performance. The identified positive impact of shortcuts on the solution performance (for suitable decision heuristics) mimics the result of Fabiunke (1999) and was confirmed for human-subject networks by Judd et al. (2010) and McCubbins et al. (2009). Another algorithmic investigation was done as the PhD work of Vattani (Vattani 2012, in particular Ch. 2; Coviello et al. 2012), who focused on the algorithmic investigation of the problem and the identification of good local decision heuristics. The problem has since been discussed in the context of

[55] In fact, both research attempts were funded, through the "Autonomous Negotiating Teams" (ANTS) program of the US Defense Advanced Research Projects Agency (DARPA).

[56] Missing and incorrect information transfer had already been touched upon in Fitzpatrick and Meertens (2002).

scheduling (Windt and Hütt 2010), organization theory (Shore et al. 2015), as well as social and political coordination processes (e.g. Enemark et al. 2014).

The idea of using GCD as a model to investigate the role of hierarchy in DCSPs in particular, was nurtured by the finding of Kearns et al. (2006) that the "leader cycle", a ring-like structure with an additional, highly connected leader in the center, showed significantly better performance as compared to a simple ring graph. The idea of adding nodes of elevated (high degree) position in the graph was later picked up by Enemark et al. (2014), who studied the "star network" that featured prominent nodes but abandoned the edges in the original ring graph. They also did not study the effect of any incremental changes in the network (as did not Kearns et al. (2006)), stating that the "networks differ in multiple parameters, so it is difficult to attribute the differential effect to any given parameter with much confidence" (Enemark et al. 2014, footnote on p. 130).

In this chapter, the thesis extends the previous work, investigating in more detail the impact of the network structure — in particular, the presence of hierarchy — on the GCD solution performance. Hierarchy as a feature of network structures (c.f. Section 3.3.2) will be investigated and a cause-and-effect model explanation for the role of hierarchy in DCSPs will be developed. This research will also continue to exploit the natural relationship between the GC model and scheduling tasks by discussing a novel "forward-model" analogy between the GCD model and a manufacturing control environment.

4.3 Model Setup

With the GCD model now introduced and motivated from scheduling and organization theory as a suitable model to investigate the behavior of agents in coordination problems, the model used in this chapter can now be defined. Both the static network layout and the dynamics that evolve on it will be introduced in the following subsections.

4.3.1 Network Layout

To investigate the quality of control system architectures (not find the best distributed GCD solution approach), this thesis will follow the advice of e.g. Hall and Posner (2001), Hernando et al. (2016), Hooker (1995), and McGeoch (1996) and use not only one limited set of test-instances, but apply a "factorial design" (Hooker 1995, p. 38) approach that allows to generate problem instances based on a limited set of parameters believed to be critical to performance and compare the *relative* performance of simple solution algorithms as a function of these system parameters. This section presents the stepwise generation of the problem graph (c.f. Fig. 4.1) and introduces the relevant parameters. The mapping of network parameters to the features of the scheduling problem are summarized in Table 4.1.

First, a set of possible states (colors) for each node is defined. As discussed above, the size of this set can be interpreted as the number of parallel identical servers available for the PPC problem. These colors form the set Σ_C. Since in the following experiments this

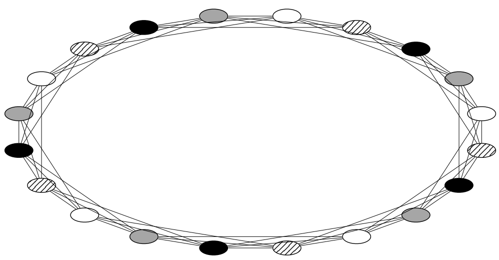

(a) Step 1: Regular ring graph.

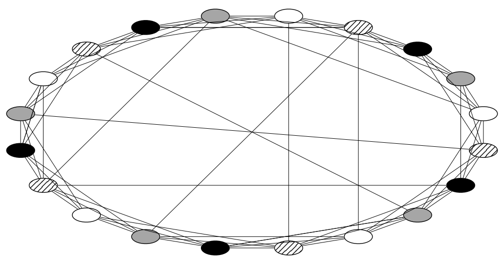

(b) Step 2: Shortcuts are added between ring nodes.

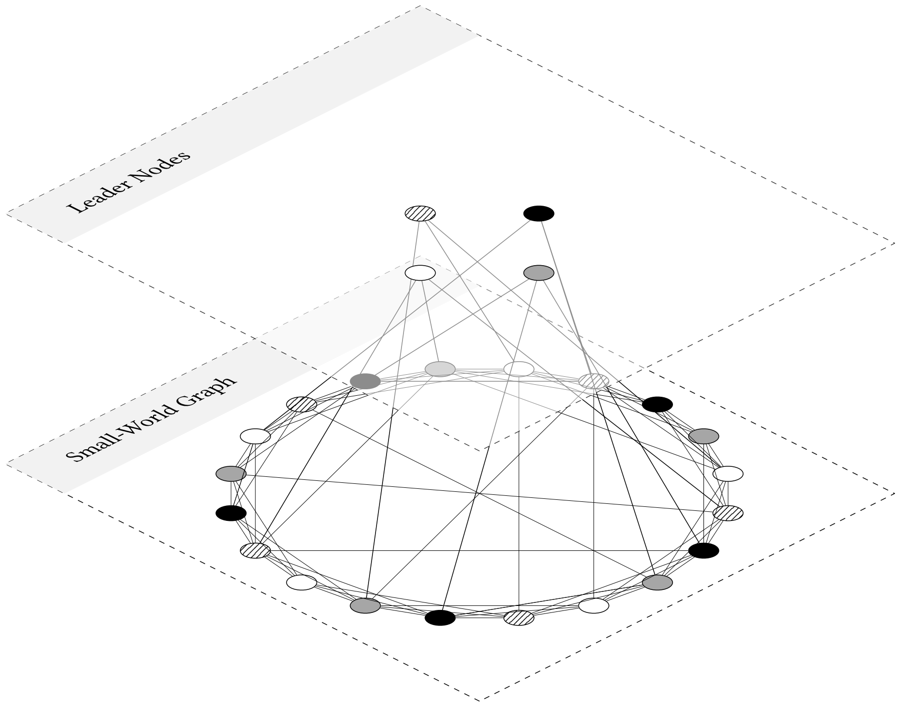

(c) Step 3: The (possible) addition and linking of so called "leader nodes" leads to the final problem setup.

Figure 4.1: Step-wise creation of the Watts-Strogatz inspired network setup for the GCD experiments. The shown coloring is one feasible solution of the GC-problem for this graph (here: $N = 20, l = \chi = 4, c = 3, k_L = 3, h = 0$).

number is equal to the graph's chromatic number, the size of the set can be identified as $\chi = |\Sigma_C|$.

This research employs a variation of the arguably best-known approach to generate networks with both high local clustering and small average path length (collectively known as the "small-world properties"): the famous *Watts-Strogatz Model* (Watts and Strogatz 1998).

As proposed by Watts and Strogatz (ibid.), a regular ring graph (Fig. 4.1(a)) lays the foundation for the network. In the ring graph, every node is connected to its spatial neighbors. Translated to the domain of PPC, scheduling constraints are assumed to occur primarily between operations that are "close" to each other in space and/or time. The "toroidal arrangement" toward a ring can be interpreted as the result of applying the *periodic boundary condition*, an approach developed in the field of partial differential equations in solid-state physics. Applied in CAs, the periodic boundary condition means that the opposing ends of the grid are "tied together" to feed also the peripheral cells with reasonable input and avoid artifacts at the ends of the grid. For the experiments in this research, ring graphs of $N = 60$ nodes will be analyzed where every node is connected to its $c = \chi - 1$ neighbors on either side. This setup guarantees, when N is an integer multiple of χ, χ-colorability of the graph. In particular, any permutation of all available colors maintained along the ring will solve the graph. At the same time, colorings with $< \chi$ colors are ruled out, as the edges in the ring form complete sub-graphs of size χ. This gives us χ-nary CA tasked with finding a χ-coloring of a graph with chromatic number χ.

Unlike the original Watts-Strogatz model, random rewiring is *not* performed. Instead, additional *shortcuts* between ring-nodes are added to the graph (Fig. 4.1(b)). As the shortcuts used, e.g. by Fabiunke (1999), these additional edges do not further reduce the solution space (any color configuration solving the original ring graph will also solve the ring graph with shortcuts), but are meant to connect otherwise distal nodes to each other in random fashion. In the PPC analogy, shortcuts represent "long-range" interdependencies between operations. In the experiments, $s = 30$ shortcuts will commonly be added to the ring. Given that the ring was originally designed to be solvable by any permutation of the χ colors, χ-colorability can easily be maintained when the ring nodes are sequentially mapped to modulo classes with respect to χ (i.e. the i'th node is mapped to class $a(i) = (i \bmod \chi)$), and shortcuts between nodes i and j are permitted *if and only if* $a(i) \neq a(j)$.

Adding shortcuts increases the connectivity of the graph. Nodes can "see" the colors of more nodes directly, and their local information horizon grows. It is, therefore, reasonable to model the transition between fully distributed control and centralized global control through the addition of shortcuts. An interesting alternative implementation of this transition would be a change in update rules (c.f. Marr and Hütt 2009) — in particular, an extension of the neighborhood considered for node updates. However, for the network sizes assumed for production control environments (in the order of hundreds of nodes), the network diameter is so small that the extended neighborhood of one node would essentially cover a large share of the graph. Thus, this research will stick to the network-structure driven approach to emulate the effect of "central coordination".

In particular, l so-called "leader nodes" may be added to the problem graph (Fig. 4.1(c)). The cases $l = 0$ (no leader nodes) and $l = \chi$ will be considered. Each leader node is connected to k_L nodes in the ring graph. Again, χ-colorability is preserved by assigning the leader nodes to modulo classes (one per class) and limiting the set of possible edge-recipients among the ring-nodes to those of different modulo class. Given N and χ, k_L finds its natural upper bound at $\frac{\chi-1}{\chi} \cdot N$. As this upper bound is a function of N and χ, one can express the *relative* prevalence of leader node interaction across different network and color set sizes. In particular, this research will (inspired by Fig. 1.4) refer to the quantity

$$1 - \frac{k_L}{\frac{\chi-1}{\chi} \cdot N} = 1 - \frac{k_L \cdot \chi}{(\chi - 1) \cdot N} \tag{4.2}$$

as the *degree of autonomous control* (this is equivalent to $1 - \lambda$ in the model of Solow and Szmerekovsky (2006), discussed in Section 2.3.2). The leader nodes are comparable to mediator agents or federation architectures discussed in MAS research (c.f. Section 3.3.2). They also bear resemblance to the idea of *liaisons* between interdependent departments, considered in organization theory as an approach to improve coordination in decentralized organizations (suggested by Galbraith 1973, Ch. 5, 1974; investigated empirically by Siggelkow and Rivkin 2005), although this research does not separate the ring nodes into departments or assign leader nodes to particular parts of the ring graph.

Leader nodes are visualized in Fig. 4.1(c) as being at an elevated level. This is for visualization purposes only. The nodes are absolutely identical in the applied update rules (c.f. Section 4.3.2). Any observed difference in emergent system behavior must be caused by their integration in the network, maintaining the common homogeneity assumption in CAs (Marr and Hütt 2009) and the research on DSAs (Zhang et al. 2002).

Finally, the graph is parametrized by a variable $h \in [0, 1]$, which indicates the relative graph density among leader nodes. For a complete undirected graph among χ nodes, $\chi \cdot (\chi - 1)/2$ edges are necessary; for any value of h, a random subset of $h \cdot \chi \cdot (\chi - 1)/2$ of these edges will be formed. For obvious reasons, h can only reasonably be defined for $l > 0$.

In summary, a coloring problem is an undirected graph, defined by the tuple $(\chi, N, c, s, l, k_L, h)$. The mapping of parameter interpretations in the domain of PPC is again summarized in Table 4.1. Parameter sweeps over the possible values of k_L and h seem appropriate to investigate the effect of increased leader node involvement in the graph (movement toward central coordination) and increased coordination among leader nodes respectively.

In the context of the discussion in Section 4.1.2, it is important to note that, as can be seen in Fig. 4.2, leader nodes also attain high levels of betweenness centrality, given that they hold (already for small values of k_L) more edges that constitute shortcuts in the graph. Hence, leader nodes in the problem graphs have comparable network structural attributes, as observed for leaders in complex organizations.

Assessing Graph Complexity

With the parameters defined to model the transition between distributed and centralized control, another dimension of PPC problems can be mapped to the model: complexity.

	Interpretation within GCD	Interpretation in context of PPC
N	Number of ring nodes	Number of operations in the system
l	Number of "leader nodes"	supervisory control instances in the shop
χ	Graph's chromatic number	Number of machines
c	Number of connected neighbors	Level of interaction among operations close in space and time
k_L	Number of edges per leader node	Level of interaction between leader nodes and operations
h	Network density among leader nodes	Level of exchange and interaction between control instances

Table 4.1: Mapping of parameters of the GCD graph generation model to properties of a manufacturing (control) system.

The complexity of the PPC problem has been mentioned as a challenge to traditional HPP systems e.g. in Sections 1.1.1 and 2.2.1 and as part of the classification model developed in Chapter 3 was discussed in Section 3.3.1. It is also assumed to affect the effectiveness of distributed PPC by Philipp et al. (2007) and Scholz-Reiter et al. (2009a), where it forms an additional independent variable, influencing the logistics target achievement in Fig. 1.4.

The question remains as to which network parameter should be used to vary it. Since the leader nodes and the edges connecting them to the ring graph are considered part of the control structure, it makes sense to disregard them for the scope of complexity analysis. The parameters χ, N, and s are left as candidates.

N, the number of nodes in the graph, is better described as the problem size. Size and complexity, however, are conceptually different (c.f. e.g. Dewar and Hage 1978, for a discussion in the context of organization theory). Shortcuts (and hence their number s) is a better candidate at first sight, as they make the graph less regular and meet our conceptualization of "complexity". One should be cautioned though by the previous discussion on the role of shortcuts, being largely perceived and *simplifications* of the problem (Fabiunke 1999; Hadzhiev et al. 2009; Judd et al. 2010) (c.f. also Section 4.3.3). The graph's chromatic number χ is arguably the most difficult to assess in terms of its contribution to graph complexity. In the context of GC, an increase in the chromatic number for a given problem graph eases the problem. The same would apply for the scheduling on parallel identical servers (discussed in the context of PPC in Windt and Hütt 2010). Yet, in the model reported here, changes in χ also change the network: Most significantly, the set of possible permutations from which one has to be selected and implemented across the ring grows more than exponentially with χ (c.f. Section 4.5.1). For this reason, this research will use χ as the gauge to assess the problem complexity in this chapter.

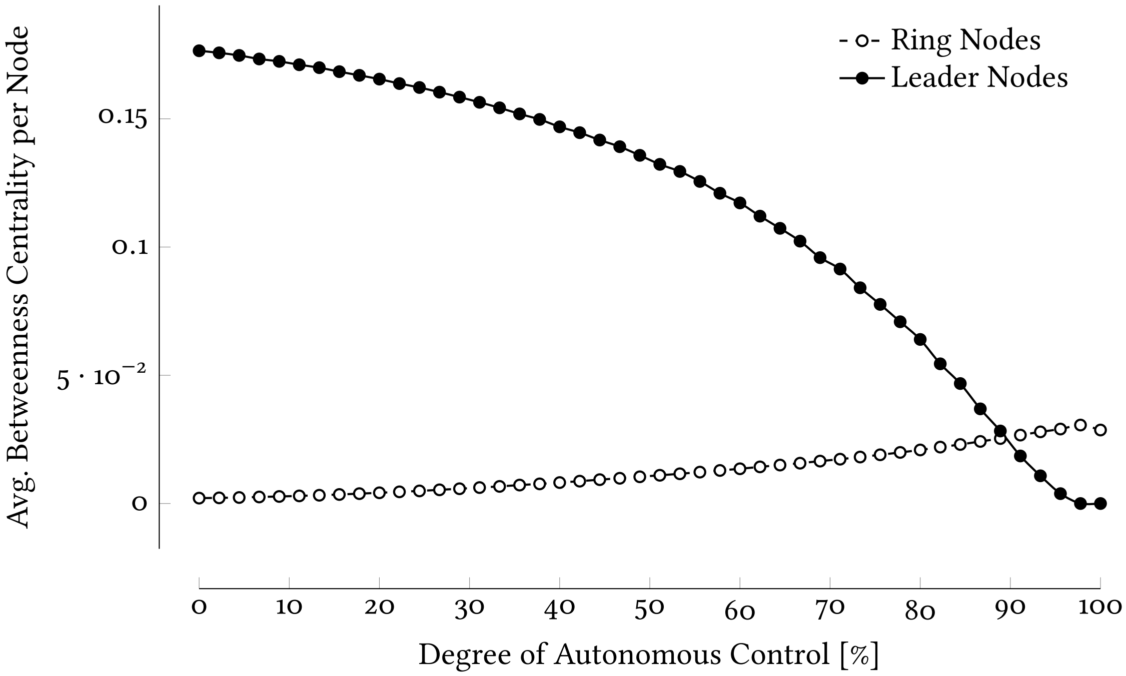

Figure 4.2: Betweenness centrality for leader and ring nodes as a function of the degree of autonomous control (mediated by changing k_L) for unconnected leader nodes, $\chi = 4, N = 60, s = 30, h = 0$, averages taken over 100 samples.

4.3.2 Applied Network Dynamics

The network dynamics determines the dynamic evolution of the node states over time, as it specifies for every node how to assess information sources and take decisions on that basis. The network dynamics then match perfectly the previously (Section 2.1) mentioned definition of *control* as a process of updating information and beliefs Pearl (1988, Ch. 6.4.1).

This research will focus on strategies that mimic the "natural heuristics" generic decision strategies observed by Kearns et al. (2006) in their human-subject studies and later formalized and tested by Hadzhiev et al. (2009). In their computerized investigation of GCD and its performance dependency on the network structure (especially the number of shortcuts), Hadzhiev et al. (ibid.) develop a classification of behaviors that every node could apply. In particular, they decompose the behavior into a *neighborhood assessment strategy* that determines how a node changes its color, given a color distribution among its neighbors, and a *temporal organization strategy* that determines which nodes are allowed to re-assess their color (execute their neighborhood assessment strategy), thereby defining how node activity spreads across the graph.

Such an approach turns the dynamics into a production-rule formalism, i.e. a sequence of "if (premise), then (action)" rules that, according to Pearl (1988, pp. 148) has particular appeal to AI researchers, since "both the activation and the action are meaningful, because they engage semantically related propositions". Collectively, these two strategies establish a *spreading activation*, where activity in a node (leading to a potential color change) is triggered by "changes occuring in logically related propositions" (ibid., p. 147).

Out of the classification of dynamics set up and explored by Hadzhiev et al. (2009), the *AW*-rule will be applied here. This means that the authorization to change colors (called *attention* by Hadzhiev et al. (ibid.)) is forwarded along all edges adjacent to a node that has just changed its color. The strategy also incorporates so-called *attention waves* to avoid continuous node excitation. Under this concept, a node can only be activated (move to state E and invoke a color change) if it has previously been in the quiescent state Q. After any invocation, the node will transfer from E to refractory state R and remain there for a fixed number of time steps. This refractory time is set to 2 time steps through the following experiments. By blocking immediate re-invocation, the nodes transform the undirected spreading activation into directed "waves" that can move through the network. Once activated, nodes apply a *strategic waiting* approach, where nodes generally seek to change to the conflict minimizing color, but will retain their current color with probability p, if the current color is conflict minimizing. With probability $(1 - p)$, will the node choose from all conflict minimizing colors or (if the current color is the only such color) from Σ_C. $p = 0.9$ will be set for all experiments. In addition to the described waves, attention may be given to a node "out of thin air" through a random excitation scheme. The probability for each node to be randomly excited in any given round is set to 0.05 throughout all experiments in this chapter.

The choice for the *AW* strategy and the parameter defined above, is largely motivated by the promising performance in the face of an increasing number of edges shown by this strategy in the experiments of Hadzhiev et al. (ibid.). However, based on the discussion in this chapter and this thesis so far, further good arguments can be made:

Attention waves alter the node update policy from parallel to sequential (c.f. also Deutsch and Dormann 2005, Ch. 4.3.4). This prevents indefinite dynamic loops observed for update mechanisms that are both parallel and synchronous as previously discussed and observed e.g. in Brueckner and Van Dyke Parunak (2006), Fabiunke (1999), Zhang et al. (2002), and Pearl (1988, Ch. 4.1.1). The notion of "strategic waiting" in the neighborhood assessment policy bears distinct similarity with e.g. the anti-system nervousness strategy presented by Karuna et al. (2006). More generally, "calming" agents is a frequently applied myopia reduction strategy, as reviewed in in Section 3.5.3 (c.f. also Hogg and Huberman 1991; Van Dyke Parunak et al. 2003). The strategy also features the basic idea of the *min-conflict rule* to solve DCSP already applied/suggested in Fabiunke (1999), Fitzpatrick and Meertens (2001, 2002), and Van Dyke Parunak (1996) and usually considered a good *repair heuristic* to work based on not conflict-free (in particular: random) initial variable assignments (Minton et al. 1992). In the context of CAs, the min-conflict rule can be understood as a particular form of an *outer-totalistic update rule* (Marr and Hütt 2009; Wolfram 1983) where the action taken by a cell in the CA does not depend on the exact allocation of states in the neighborhood, but aggregate measures, thus allowing a concise definition of rules for CAs of arbitrary size and shape (Marr and Hütt 2009). Finally, the random changes of decisions by agents is a proven method to leave local optima (Pearl 1988, Ch. 4.1.1) and also provides a natural initiation of the solution process after the experiment start.

When discussing the results in the context of CLT, caution has to be raised concerning the verisimilitude of the behavior with respect to humans interacting in organizations.

Mason and Watts (2012) rightly argue that agent-based simulations are at risk of conceptually misrepresenting the behavior of human decision-makers, rendering results of little practical applicability. Lazer and Friedman (2007, p. 672) demand that "Research based on formal representations of human behavior needs to convince readers that the necessarily constrained assumptions in a model somehow capture the essence of some empirically relevant set of circumstances". Notably, in this case, choosing the conflict minimizing color was reported by the participants in the experiments of Kearns et al. (2006). The computerized investigation also showed an increase in solution speed with increasing connectivity of the graph (Hadzhiev et al. 2009), a result reported also in human subject experiments (c.f. e.g. Kearns et al. 2006; McCubbins et al. 2009). Thus, one may assume some degree of conformity between the experimental setup and humans in coordination settings (with low network visibility). Obviously, only a computerized experiment setup allows a high number of repetitions across a wide range of network structures as required here and allows us to set experiment parameters at will (c.f. Dal Forno and Merlone 2007).

Human-subject experiments (such as Enemark et al. 2014; Kearns et al. 2006) have also focused on the effect of increased information horizons on the decision-making performance. This research will refrain from such extensions here, as an extension of the agent's information horizon beyond its immediate neighbors would

1. quickly imply that agents can "see" the entire network, given the small diameter of small-world graphs (already discussed in Section 4.3.1) and
2. significantly complicate the formulation of agent decision-making logic.

4.3.3 Performance Measurement

The GCD problem then is to find, from a random initial color distribution, a color assignment for all $N + l$ agents that satisfies all color inequality constraints. In the initial experiments, the number of color changes (across all nodes) required to reach this solution (referred to as $r_C(G)$) will be used as the sole performance metric, assessing the control networks ability to find a coherent and consistent solution quickly (c.f. Section 4.1.1). A different performance measure, closer to those applied to PPC systems, will be presented and applied in Section 4.6.

Result Normalization

It was observed early in the investigation of GCD with human subjects that additional edges aided the solution process (Kearns et al. 2006; McCubbins et al. 2009). An explanation was provided in later work (Enemark et al. 2011; Judd et al. 2010), who found that redundant edges (edges that allow information exchange but do not further constraint the solution) in particular, help in the solution process. Both the shortcuts and edges to leader nodes are entered, maintaining χ-colorability, and hence fall into this category. Edges that represent additionally constrain feasible color distributions, on the other hand, impede the solution process (examples are the additional shortcuts between clusters used in JUDD et al. (2010)). This conclusion is in line with the results of other experiments with both human and computerized decision-makers (ENEMARK et al. 2011; FABIUNKE 1999; HADZHIEV et al. 2009; JUDD et al. 2010); it confirms theoretical considerations by GALBRAITH (1974) that lateral relationships between otherwise independent decision-making entities should increase the capacity to process information. YOKOO (2001, Ch. 1.5.1) investigates the landscape for CSP algorithms and find that once a threshold is reached, the problem difficulty decreases with number of shortcuts. This threshold is reached once the problem is sufficiently constrained to reduce the number of possible solutions to a minimum. With the set of possible solutions already being fully constrained by the edges among neighboring ring nodes, it can safely be assumed that the experiments reported here have reached/passed this threshold.

It is to be expected then that the addition of shortcuts to and from leader nodes will also have a positive effect on the overall solution speed. Since this thesis is interested in the effect of network *structure* on the solution speed, one should control for the overarching effect of additional edges independent of their location in the graph. One can do so by normalizing the number of color changes required to solve the network *with* leader nodes by a network without leader nodes, but with the same number of nodes and shortcuts. To state this explicitly: the *relative performance* of a GCD network $G = (\chi, N, c, s, l, k_L, h)$ with leader nodes ($l = \chi$) is computed by comparing the number of color changes required to solve the network ($r_C(G)$) with the number of color changes necessary to solve a *reference network* G_{ref}. Note that G (for $h \neq 0$) has a total of $N + \chi$ nodes and a total number of shortcuts given by

$$\underbrace{s}_{\text{shortcuts in ring}} + \underbrace{\chi \cdot k_L}_{\text{edges from leader nodes}} + \underbrace{h \cdot \frac{\chi \cdot (\chi - 1)}{2}}_{\text{edges between leader nodes}} . \tag{4.3}$$

Table 4.2 shows how the parameters of compared test and reference networks compare.

Parameter	χ	N	c	s	l	k_L	h
Test-network G	χ	N	$\chi - 1$	s	χ	k_L	h
Ref.-network G_{ref}	χ	$N + \chi$	$\chi - 1$	$s + \chi \cdot k_L + h \cdot \frac{\chi \cdot (\chi - 1)}{2}$	0	0	0

TABLE 4.2: Mapping of test- and reference networks for result normalization in the GCD experiments.

The relative performance is then expressed similar to Eq. (3.1).[57] In the model at hand, a *lower* number of color changes required is associated with a higher performance. Hence, the relative performance is expressed as

$$\delta_C(G, G_{ref}) = \frac{r_C(G_{ref})}{r_C(G)} . \tag{4.4}$$

[57]Notably, the reference model here is not assumed to exhibit a lower degree of myopic behavior but provides a benchmark scenario to account for the effect of additional edges

Measure of Uncertainty

Owing to the stochastic nature of large parts of the network generation and dynamics, a significant variability in performance metrics between iterations is to be expected. In the following figures, standard errors (standard deviation of the mean value estimate) are reported instead of standard deviations in the figures in the remainder of this chapter. The standard error is attained by dividing the standard deviation by the square root of the number of samples.

For the ratio of reference and test network performance, the standard error can be computed based on commonly known results from uncertainty propagation laws. In particular, given the standard deviations and number of samples for both test and reference networks, the standard error associated with the relative performance, the ratio $y = \frac{x_1}{x_2}$, can be calculated as

$$\Delta y = \sqrt{\left(\frac{\Delta x_1}{x_2}\right)^2 + \left(-\frac{x_1 \Delta x_2}{x_2^2}\right)^2} \tag{4.5}$$

where the values for Δx_1 and Δx_2 are the standard errors of the reference and test networks respectively.

4.4 Numerical Results

All preliminary work is in place now to use the above-described GCD model to assess the ability of different network configurations to solve the distributed GC problem. We first focus on variations in the leader node involvement parameter k_L. In the presentation of the results, we will pick up the visual account of Fig. 1.4 by showing the relative performance acc. to Eq. (4.4) as a function of the degree of autonomous control (Eq. (4.2)) and the complexity measured by the graph's chromatic number (c.f. Section 4.3.1). To start, we investigate the case with unconnected leader nodes ($h = 0$).

The result is shown in Fig. 4.3. The result shows important characteristics hypothesized before. Most importantly, it shows a clear peak in relative performance for medium values of autonomous control — with decreasing performance values to either extreme of the spectrum. The figure also shows a general downward trend as complexity increases, as had been hypothesized in Fig. 1.4. For very high degrees of autonomous control (low values of k_l), a slight increase in relative performance can be seen. However, this appears to be an artifact of the setup.

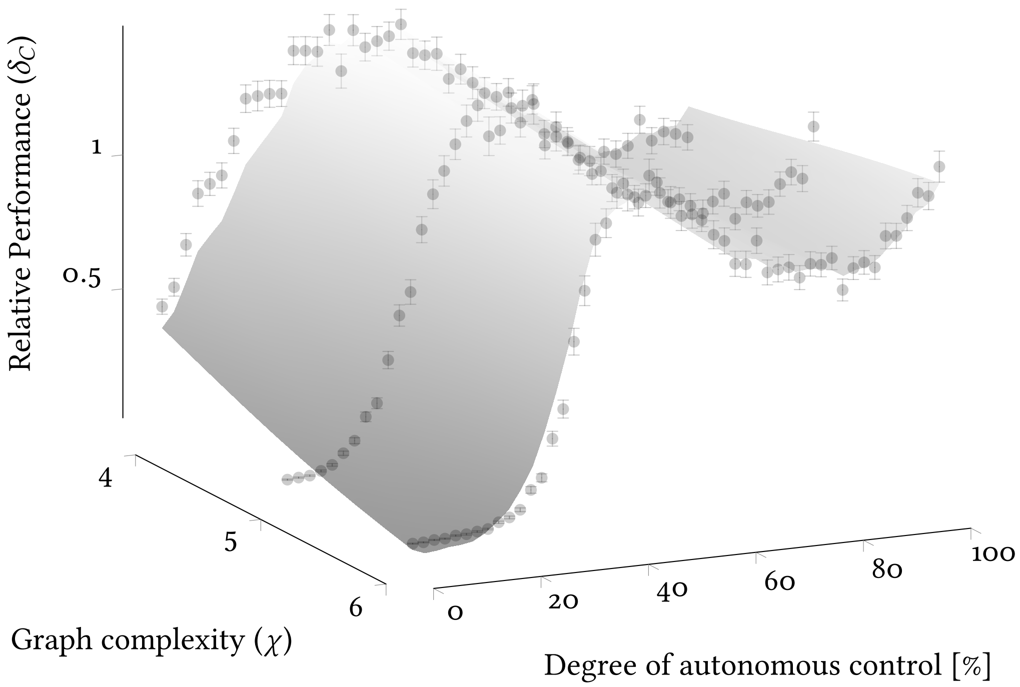

Figure 4.3: Relative performance as a function of the degree of autonomous control and graph complexity for unconnected leader nodes, $N = 60, s = 30$, averages taken over 400 samples.

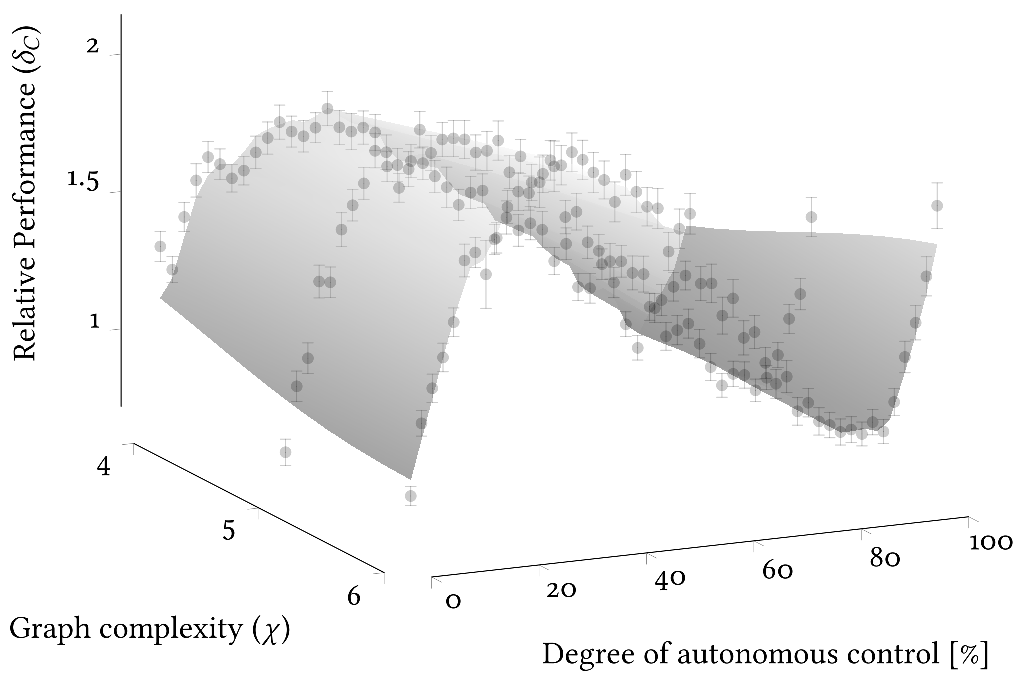

Figure 4.4: Relative performance as a function of the degree of autonomous control and graph complexity for fully connected leader nodes, $N = 60, s = 30$, averages taken over 400 samples.

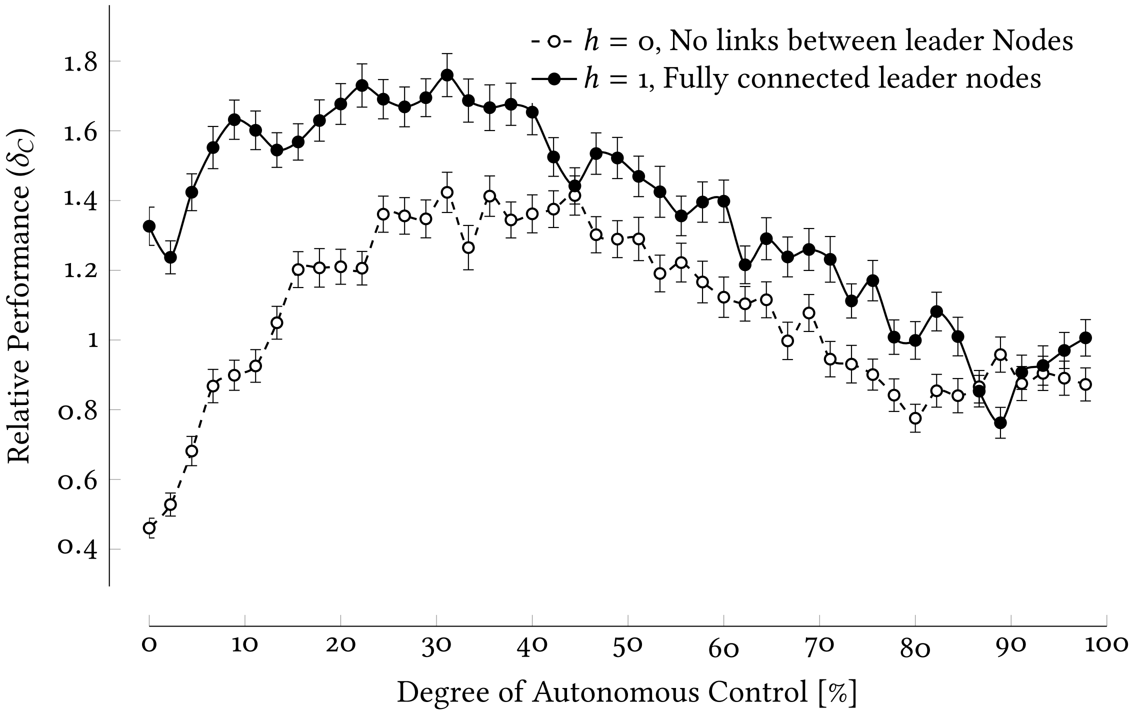

(a) Comparison of relative performance at $\chi = 4$ for both unconnected ($h = 0$) and fully connected leader nodes ($h = 1$), $\chi = 4, N = 60, s = 30$, averages taken over 400 samples.

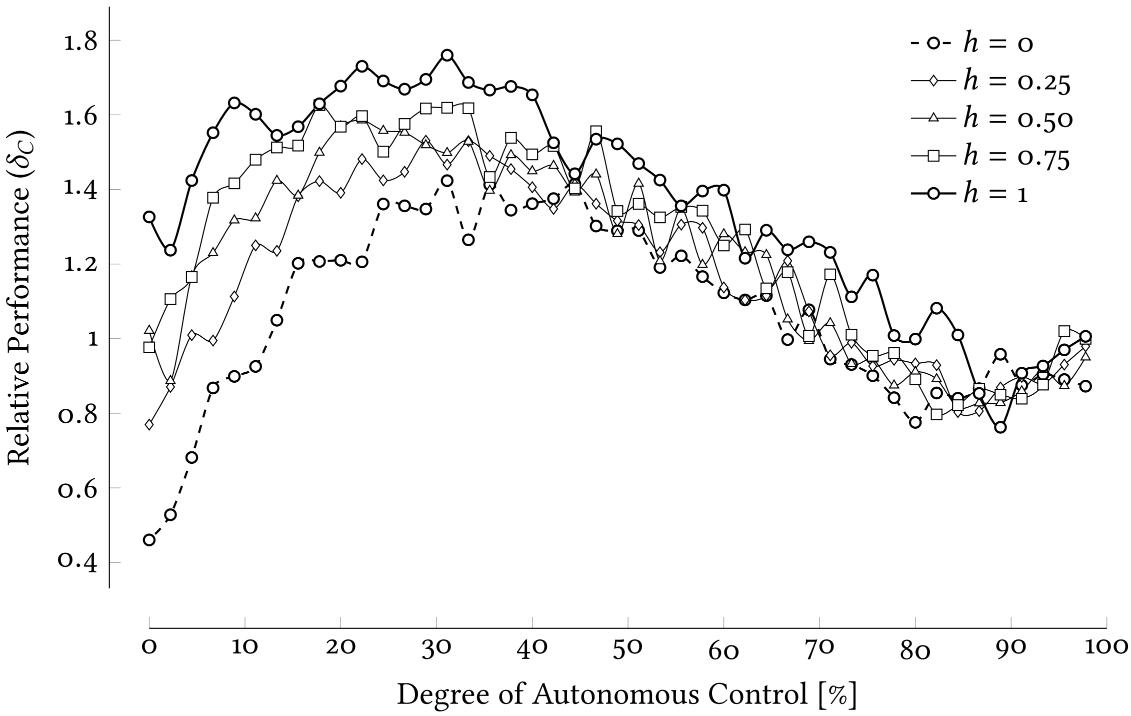

(b) Evolution of relative performance, as the degree of leader node connectivity (h) is increased incrementally, $\chi = 4, N = 60, s = 30$, averages taken over 400 samples, error bars omitted for visual clarity.

Figure 4.5: Impact of connections between leader nodes on the relative performance as a function of the degree of autonomous control.

4.4.1 The Role of Intra-Leader Node Communication

Now, the impact of changes in the parameter h is investigated, starting with the polar opposite of Fig. 4.3, a complete network formed between the leader nodes ($h = 1$). The result is shown in Fig. 4.4

The comparison of Figs. 4.3 and 4.4 shows a distinctly different shape of the curve for connected leader nodes. For a better comparison, Fig. 4.5(a) leaves aside the complexity axis and compares unconnected ($h = 0$) and connected leader nodes ($h = 1$) for one level of complexity (χ).

The results show a significant difference, especially for low levels of autonomous control (high levels of k_L). The drop in performance for highly hierarchical control networks, hypothesized in Section 2.3, seems much less distinct when leader nodes are interconnected and hence able to share information directly. Figure 4.5(b) shows further that this result is not an artifact at the polar ends of the range of possible values of h. In fact, as h is increased, we see (a) a continuous upward shift in relative performance and (b) a shift in the maximum attained performance toward network configurations with more hierarchical features.

4.5 Understanding the Role of Leader Nodes

With leader node connection with the ring graph and among each other showing a clear impact on the solution performance, this section sets out to understand the mechanisms that lead to the observed phenomena.

Leader Nodes do not Settle First

An obvious first hypothesis for the (at least partially) positive impact of leader nodes could be that these nodes “agree” on a solution for the entire graph before the remaining graph nodes settle. The hypothesis basically assumes that leader nodes solve a simplified version of the problem “internally” before propagating the attained result throughout the entire graph. The observed positive effect of connected leader nodes fits nicely within the hypothesis, as the links between leader nodes allow direct communication among leader nodes and could easily be imagined to support such internalization of the solution process.

All questions of precise mechanisms aside, for this hypothesis to show any promise, one would need to observe that leader nodes “settle” first, which means that they show a distinct reduction in color change activity — not only over time, but more pronounced than the ring nodes. The obvious first test for this hypothesis then is to measure the color change activity (the probability of color changes per round) over the solution process and compare the results. This is done in Fig. 4.6.

The result quickly refutes the above-stated hypothesis. While a reduction in activity — across all nodes — can clearly been observed, there is no evidence that leader nodes were

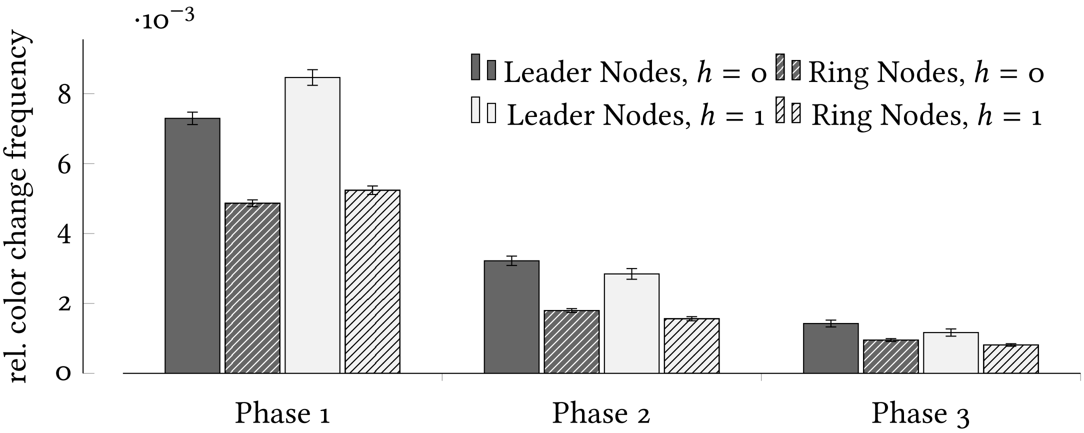

Figure 4.6: Relative color change frequency for both leader nodes (solid bars) and ring nodes (striped bars) and both unconnected (dark bars) and fully connected (lighter bars) leader nodes in each third of the solution process (total number of rounds required until solution of the graph), averages taken over 400 samples.

to settle earlier. Also, no difference between connected and unconnected leader nodes is visible, further stymieing hopes to explain leader node behavior and impact through changes in activity over time.

4.5.1 The Evolution of Order

Since such simple quantitative analysis of leader node activity falls short of capturing the mechanisms by which leader nodes support the solution process, this thesis now starts investigating Hypotheses H_2 and H_3. To this end, this subsection will derive a quantifiable definition of (partial) solutions for this specific GC problem that will be used to visualize and measure the spread of "order" throughout the agent population.

The evolution of a single node i over time is described by the time series of its color selection $x_i(t)$. The color selections of all nodes in the graph at some point t form a *state vector* that describes the color distribution across the entire graph at t. Given the ring-based construction of the graph (Section 4.3.1), we know unequivocally that only a continuous permutation of the χ colors along the ring and across the leader nodes can solve the GC problem. Hence, among the $\chi^{(N+l)}$ *possible* state vectors of the graph, only $(\chi - 1)!$ different state vectors are valid solutions, as they represent *circular permutations* of the elements in Σ_C. Circular permutations are permutations with no defined "first element" (Brualdi 2009, Ch. 2.2). In particular, there are $(\chi - 1)!$ circular permutations of χ elements. The condition for finding a feasible solution to the GC problem can be re-stated as follows: the underlying GC problem is solved, *if and only if* all nodes are part of the same circular permutation.

Critically, one can identify circular permutations also locally in the graph. In particular, looking at any χ subsequent ring nodes, or the χ leader nodes in conjunction, one can

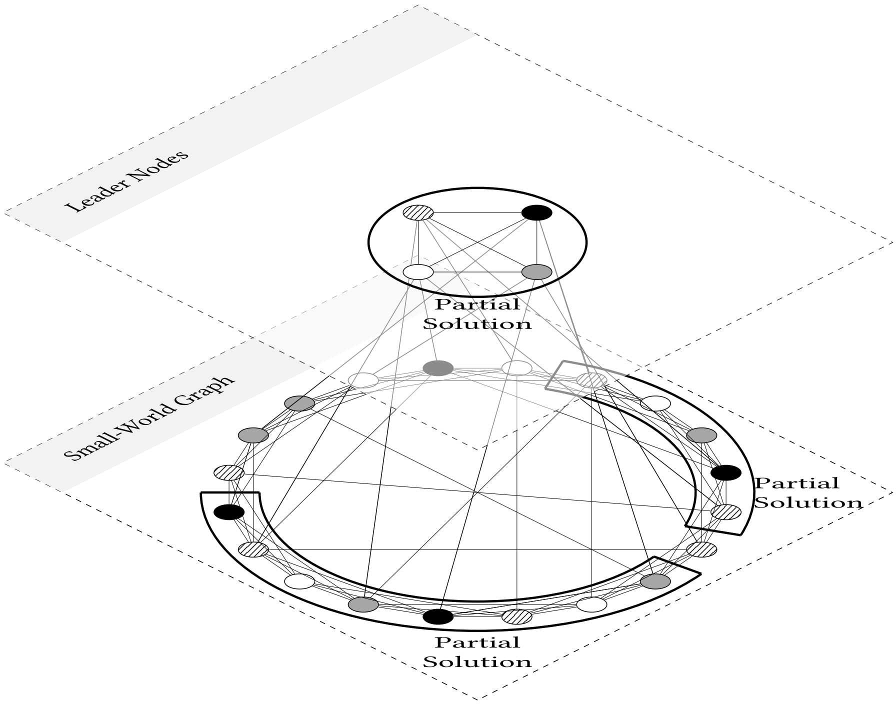

Figure 4.7: Partially solved GCD network with sections of locally coherent node-colorings (solution regimes).

determine, which circular permutation (if any) these nodes form. The nodes will not form any valid circular permutation *if and only if* at least one color occurs more than once.[58] Fig. 4.7 shows a partially solved ring graph where some parts already show color permutations.

Given that established circular permutations on χ nodes represent a locally consistent solution to the GC problem, the $(\chi - 1)!$ circular permutations will be referred to as *solution regimes* (σ_i). With $\Sigma_S^{(0)}$ being the set of all solution regimes.

$$\Sigma_S^{(0)} = \{\sigma_1, \sigma_2, ..., \sigma_{(\chi-1)!}\} \tag{4.6}$$

Node i is said to be assigned to a solution regime at time t, if the sequence of colors $x_i(t), x_{i+1}(t), \ldots, x_{i+(\chi-1)}(t) \in \Sigma_S^{(0)}$, thus establishing for every node a new time-series, $s_i(t)$, that tracks the evolution of solution regimes for one node over time. $s_i(t) = 0$ will be assigned if node i is not part of any solution regime at time t. Looking again at all time-series $s_i(t)$ in conjunction, one can understand and visualize the evolution of order in both space (across the ring and leader nodes) and time. Figure 4.8 visualizes the evolution of the number of conflicts ($c(t)$) across the entire graph and the evolution of solution regimes on ring nodes over time for one instance of the GCD model. White patches indicate nodes that do not belong to a solution regime at this point of time ($s_i(t) = 0$). Note that regions associated with the same solution regime may be incompatible, due to either "unaligned" nodes in between and/or different "offsets" of the same circular permutation.

[58]Since we look at χ nodes and $|\Sigma_C| = \chi$, this is equivalent to saying that at least one color does not occur among the χ nodes.

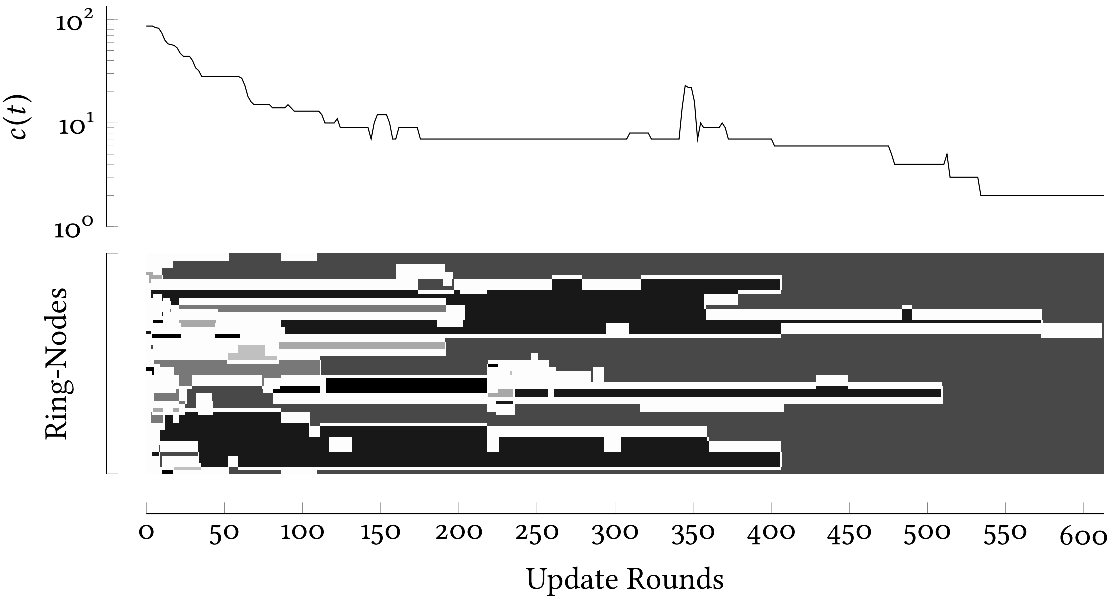

Figure 4.8: Evolution of solution regimes on a ring-graph (ring-nodes only). White indicates no solution regime ($s_i(t) = 0$), $\chi = 4, N = 60, s = 30, k_L = 30, h = 1$.

The two parts of Fig. 4.8 enhance our understanding of the temporal evolution of the solution process, developed with Fig. 4.6. The decrease in node activity (Fig. 4.6) goes along with a sharp drop in conflict-reduction speed. The initially high number of conflicts, which resulted from the random initial coloring, is quickly reduced as nodes show high activity. The solution regime plot unveils that this is achieved by nodes establishing local solution patterns that are consistent for sections of the ring, thus reducing the number of conflicting edges in the ring. The conflict reduction slows down significantly, as local conflicting solution regimes "compete" to spread across the ring. We see solution regimes disappearing and spreading as, for the largest part of the solution process, conflicts are eliminated in a painstaking back-and-forth among agents, most of them situated in a situation with just one conflict-minimizing color (the color corresponding to their local solution regime).

Such compound analysis of the composite behavior of neighboring CA cells is not new. Very early in the discussion of CAs in complexity science, *phases* and *domains* were introduced as attributes assigned to a local neighborhood of cells: "Phases in cellular automata may in general be described by 'order parameters' that specify the spatially periodic patterns of sites corresponding to each phase" finds (Packard and Wolfram 1985, p. 927). In the investigated model, a solution regime can be understood as an order parameter, as it describes the spatial configuration of cells (the sequential update rules implies and Fig. 4.8 confirms that the temporal evolution is, in fact, much reduced as compared to many update rules classically studies in CAs). A *domain* then refers to the region that has adopted a certain order parameter, belonging to a certain phase (Lizier et al. 2008; Packard and Wolfram 1985). They are visualized by the different colored patches visualized in Fig. 4.8, each defined by the locally established phase. Lizier et al. (2008) then investigate through the "application of a measure of local information transfer into each spatiotemporal point" (ibid., p. 026110-2) that in particular *particles*, the boundaries between conflicting domains (the white patches of Fig. 4.8) are essential for the transfer of information in the system. A compound investigation of neighboring CA cells was also suggested by Holland (1998) as a vehicle to attain a deeper level of understanding of the dynamics of CAs. In particular, he suggests to interpret the combined states of all cells forming a composite as a state of the composite (ibid., Ch. 10). This idea is picked up here by associating a solution regime time-series $s_i(t)$ describing the composite state to every node time-series $x_i(t)$, describing the temporal evolution of each cell.

As Hypothesis H_3 explicitly addresses the "interruption" of patterns, binary encoding of $s_i(t)$ can also be defined. In particular, let the time-series $s_i^{bin}(t)$ be defined as

$$s_i^{bin}(t) = \begin{cases} 1 & \iff s_i(t) \in \Sigma_S^{(0)} \\ 0 & \text{otherwise.} \end{cases} \tag{4.7}$$

When $s_i^{bin}(t)$ is 1, node i is "in tune" with its direct neighbors. While the color selection of these nodes may be in conflict with other parts of the graph, the nodes can be assumed to be in a local optimum, since the nodes between neighbors constitutes large parts of the total number of edges and those can only be satisfied collectively when a circular permutation is reached (a solution regime $\neq 0$ is attained). When the binary indicator is 0, node i and its immediate neighborhood have been disrupted (or not allowed yet to settle) on a locally consistent color allocation. Especially after the initial phase of high node activity

4.5.2 Mutual Information as a Measure of Leader Node Impact

To measure the impact of leader nodes on the distribution of solution regimes throughout the graph, the impact of the observation of the leader nodes' states on the connected ring nodes must be quantified. The impact of an observation on another process can generally be measured by calculating conditional probabilities (Pearl 1988, Ch. 4.1). Such information transfer "is widely considered to be a vital component of complex nonlinear behavior in spatiotemporal systems" (Lizier et al. 2008, p. 026110-1). The (time-delayed) mutual information is an established way of measuring the flow of information from sender to receiver (Kirst et al. 2016). It is also among the most commonly applied measures used to evaluate and rank information sources (Pearl 1988, Ch. 6.4.2), hence allowing statements about the relative role and importance of system components (here: nodes in the graph).

Given two random processes X and Y that can take values $x \in X$ and $y \in Y$ respectively, the mutual information between X and Y describes the reduction in uncertainty about Y that results from knowing the state of X (c.f. e.g. Lizier et al. 2008). It is also known as the conditional entropy of Y, given X and is analytically defined as (ibid.):

$$I(X;Y) = \sum_{x,y} p(x,y) \log_2 \left(\frac{p(x,y)}{p(x) \cdot p(y)} \right) \tag{4.8}$$

While the base of the logarithm is not important, a base of 2 is commonly used as it renders a result in *bit*. While the formulation in Eq. (4.8) is commutative w.r.t X and Y (the *direction* of information flow is not defined), a direction is added by introducing a time offset between sender and receiver. In particular, one can measure the "immediate" flow of information from X to Y by calculating the impact (in terms of the mutual information) of the stochastic process $\{X_t : t \in T\}$ of the sender on the time-shifted process $\{Y_{t+1} : t \in T\}$ of the receiver. The observation period of the process is, of course, limited, so the *true* individual and joint probabilities cannot be measured. The observed relative frequency of a node state is, however, operationally equivalent to the probability of this state and can hence used to calculate mutual information (ibid.).

Mutual information will be calculated between a leader and connected ring nodes to further investigate the role of leader nodes in the effect observed in Section 4.4. In particular, and in accordance with Hypothesis H_3, the mutual information flow between $s_i^{bin}(t)$ of a leader node and $s_i^{bin}(t+1)$ of a connected ring node (c.f. Eq. (4.7)) is measured. This established a clear direction of the information flow — from the leader to the ring node. The fundamentally changing nature of the solution process over time (in particular, the different behavior at the start of the experiment) observed in Figs. 4.6 and 4.8 reasons a more focused analysis, so the following analyses will focus on the mid 20% of the solution process. This decision is justified, as the initial rounds of the solution process seem largely driven by the evolution of local regimes and the rapid decrease of conflicts in the graph. This initial process phase is largely driven by nodes adjusting the initial coloring to the conflict-minimizing color, resulting in high activity and constant passing of "attention" (c.f. Section 4.3.2) along the graph. The network structure, which establishes long-range connections between nodes via the leader nodes cannot substantially contribute to this local coordination process. Most of the time, however, is spent on aligning competing solution regimes to find a globally consistent solution. The mid 20% of the solution process are deemed prototypical of this part of the solution process, where the network architecture can show its strength. The result of the mutual information analysis for this subset of the solution process is shown in Fig. 4.9. Note that the measured mutual information is an absolute value. Notably, no result normalization is necessary here, unlike previous analyses of the number of color changes.

The results show a similarity between the flow of mutual information between leader and ring nodes and the relative performance. In particular, for $h = 0$, we observe a curvilinear relationship of the mutual information that closely resembles the hypothesized performance curve as a function of the "degree of autonomous control". For $h = 1$, we observe (as in Section 4.4) a more pronounced increase in information transfer with almost no drop in information transfer for very low values of "autonomous control".

This provides the first important piece of the puzzle of understanding the role and effect of leader nodes: leader nodes drive the solution process by steering the process by which ring nodes enter and leave local solution regimes. By cuing ring nodes to leave their local optima (color configurations that minimize conflicts in the immediate neighborhood), leader nodes support the process of solution regime alignment that, as Fig. 4.8 shows, shapes the largest part of the solution process — in particular, the mid 20% for which the mutual information is analyzed.

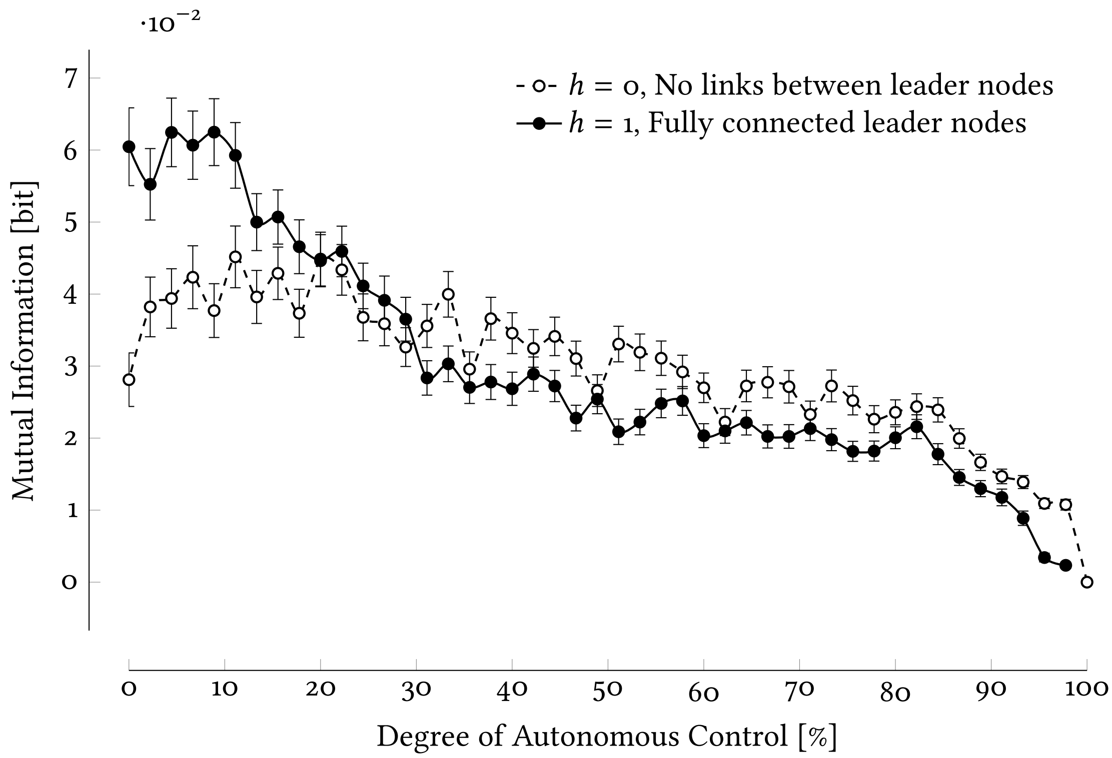

Figure 4.9: Average mutual information transfer across edges from leader to ring nodes in the mid 20% of the solution process for both connected and unconnected leader nodes, $\chi = 4, N = 60, s = 30.$, averages taken over 400 samples.

4.5.3 Investigating Leader Nodes as Noise-Emitters: Simulated Annealing

The hypothesis that leader nodes drive the solution process not by internally agreeing on a solution, but by steering the collective process of "experimentation" and breaking up existing local solutions can further be substantiated, by actively modeling leader nodes as noise-emitters meant to lead ring-nodes out of local optima. This process is similar to a distributed Simulated Annealing (SA) optimization process that has been discussed as a model for conflict-resolution processes in MASs by Arshad and Silaghi (2004) and for organizational adaptation by Carley (1997) and Carley and David M. Svoboda (1996). As a meta-heuristic to solve MILPs, SA has been applied to solve scheduling problems (e.g. by Fattahi et al. 2007) and has also been applied to the GC problem (c.f. Johnson et al. 1991). Should similar behavior be observed in the original model as well as the SA adaptation, this would support the idea that leader nodes are not prescribing solutions, yet much rather drive the solution process over time, and it would provide a powerful metaphor to understand leadership in complex organizations.

SA was introduced as an optimization meta-heuristic by Kirkpatrick et al. (1983). It is inspired by the annealing process found in solids (Eglese 1990; Kirkpatrick et al. 1983). As molecules in a solid cooling down from high temperature incrementally lose their ability to change their positions (as they could in the liquid state), SA describes a local search algorithm where the probability of accepting an inferior solution, necessary to leave local optima, is a function of the difference in the target function value and a control parameter, T, which is known as the *temperature* of the system. As the solution process continues, the temperature decreases and the search algorithm becomes less and less inclined to leave local optima — and explore decision alternatives outside the found (local) optimum — and seeks to exploit/build upon the best solution found so far instead (Eglese 1990).

With this balance between exploration and exploitation, annealing has also been discussed in the context of CLT as an apt model to conceptualize the emergence of order in CASs from local activity by agents. Uhl-Bien et al. (2007) argue:

> *"This capacity to rapidly explore solutions can be illustrated with a problem solving scenario called annealing [...]. In this scenario, multiple agents struggle with localized effects created by a given environmental perturbations (or tension; [...]). As these agents develop localized solutions, work-arounds, or related responses, they affect the behaviors of other interdependently related agents, who subsequently build on the original response to create higher-order responses. [...] In this process interdependent agents and CAS experiment, change, combine strategies, and find loopholes in other strategies [...]."*
>
> — Uhl-Bien et al. (2007, p. 303)

To test whether leader nodes act as noise sources in the solution process, their behavior will be changed to act as observers of a local search exercise, interfering according to the observed "temperature" of the system. Setting the initial system temperature and defining a functional relationship that describes the cooling-down process (the evolution of T over time) are among the design decisions to be made when using SA (Eglese 1990). For the GCD experioment, instead of performing a continuous temperature decrease over time (T is a function of time), the system temperature is coupled to the conflict count time-series $c(t)$. As the number of conflicts in the graph represent a distance measure from a completely solved graph, the number of remaining conflicts provides a gauge to assess the degree to which the search process has finished. To empirically define a functional relationship between $c(t)$ and the probability of leader nodes being activated to re-consider their color-assignment, 100 runs of the normal GCD model are analyzed and the share of leader nodes being excited (c.f. Section 4.3.2) is measured every round. Averaged over multiple simulation runs, one can describe leader node activity as a function of the conflict count $c(t)$. This relationship is shown as a scatter plot in Fig. 4.10.

Visual inspection of the resulting functional relationship across multiple parameter tuples leads to using a third-degree polynomial to smoothen the data and get a continuous functional expression for the leader node excitation probability. This fitting was done separately for every network structure (every tuple G, as defined in Section 4.3.1). Figure 4.10 shows one such fitted polynomial.

For the case of connected leader nodes, the leader nodes are treated as a single decision making entity. Unlike the original GCD experiments, leader nodes are not excited through attention propagation, but the excitation probability is determined by the conflict count and the fitted polynomial instead. Upon excitation, the joint leader nodes now have a different decision space, as they cannot collectively choose one color. Instead, their decision space is the space of feasible solution regimes ($\Sigma_C^{(0)}$), from which the combined

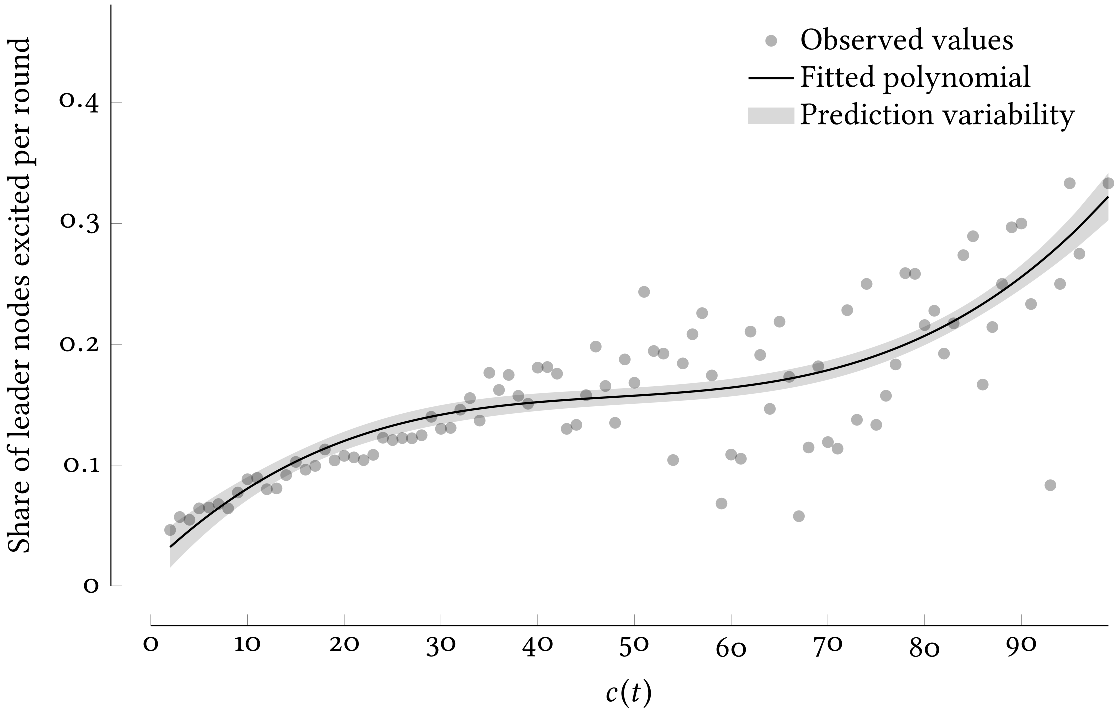

Figure 4.10: Average share of leader nodes excited in a given round as a function of the conflict count in the network and the maximum likelihood fit of a third-degree polynomial to the data, $\chi = 4, N = 60, s = 30, m = 30, h = 1$. Averages takes over 100 experiments.

leader node entity may choose a conflict minimizing option by the same principles, as outlined in Section 4.3.2.

For unconnected leader nodes, such collective decision-making seems an inappropriate assumption as the nodes have no direct link to share information. Instead, each leader node is modeled individually (and with the known solution space of colors Σ_C). Node excitation however is again not arriving through attention waves, but induced through the observed conflict count. In particular, the probabilities measured through the experiments in Fig. 4.10 are converted to account for separate entities, by calculating the corresponding individual excitement probabilities. To calculate the individual excitation probabilities corresponding to the observed overall excitation level in Fig. 4.10, a binomial distribution is used. Let p_{obs} be the probability observed in Fig. 4.10 and p_{ind} the individual excitement probability in question, then it should hold that

$$f(n = \chi, p = p_{ind}, k = \chi) = \binom{n}{k} \cdot p^k (1 - p)^{n-k} = p_{obs}. \tag{4.9}$$

Since $n = \chi$ and hence $\binom{n}{k} = \binom{\chi}{\chi} = 1$ and $(1 - p)^{\chi - \chi} = 1$, this expression is readily transformed to:

$$p_{ind}^{\chi} = p_{obs} \iff p_{ind} = p_{obs}^{\frac{1}{\chi}}. \tag{4.10}$$

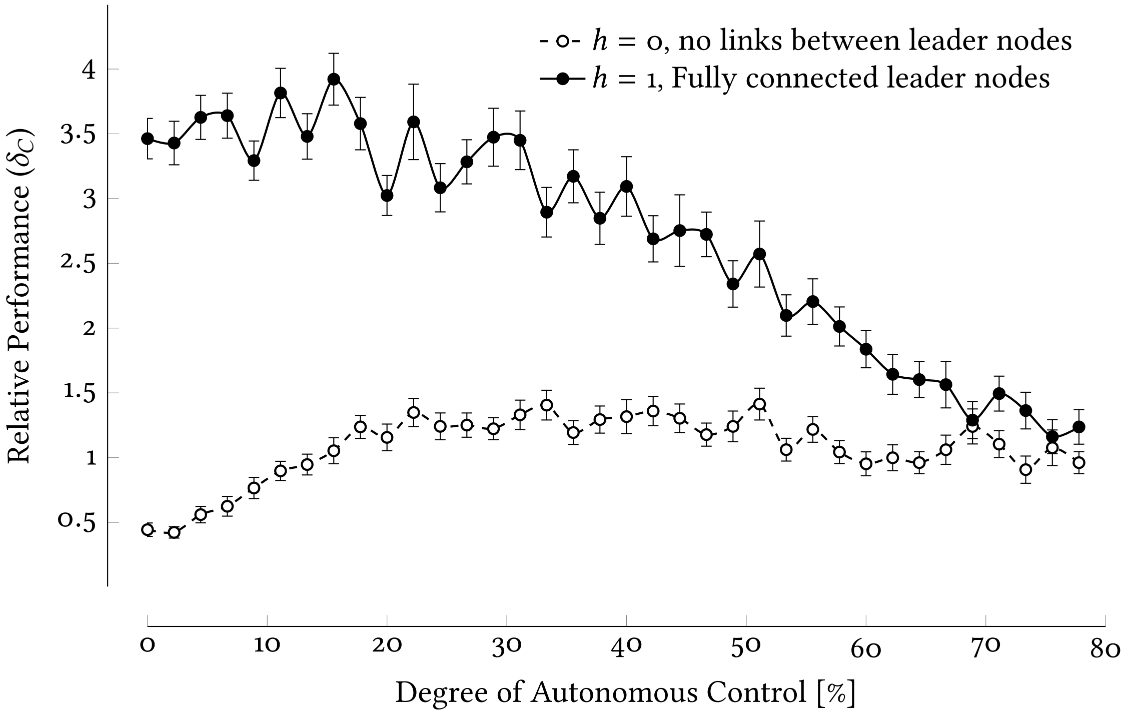

Figure 4.11: Relative performance when leader nodes are controlled through simulated annealing, $N = 60$, $\chi = 4$, $s = 30$, averages taken over 100 samples. Experiments for degrees of autonomous control 80 80% (low values of k_L) were not completed due to excessive time and memory use.

With this work in place, one can now compare the results for the relative performance between the original GCD model and the model where leader nodes are controlled according to the SA fiction. As Fig. 4.11 shows, the results are very comparable. For unconnected leader nodes, Fig. 4.11 shows almost identical relative performance values. In particular, the curvilinear relationship between the degree of autonomous control and the relative performance is maintained. For $h = 1$ (fully connected leader nodes), on the other hand, Fig. 4.11 shows significant absolute differences for fully connected leader nodes. While maintaining a similar overall shape, relative performance is higher in the SA scenario as compared to the standard GCD model. This result however is readily explained: As mentioned above, the leader nodes were modeled in the connected case as a single decision-making entity with a different (in particular: larger) decision-making space. The increased information horizon (the union of the nodes connected to any leader node) and the ability to choose a regime, minimizing conflicts for *all* leader nodes collectively (not simply a color), should give it a decisive advantage.

4.5.4 Resulting Understanding of Leader Nodes in GCD

The results of Sections 4.5.2 and 4.5.3 provide the following understanding of the impact of leader nodes on the DCSP solution process modeled through GCD: The high color change frequency of leader nodes, combined with the mutual information analysis, imply that leader nodes drive the solution process in the ring — in particular, they seem to be involved in disrupting solution regimes in the ring nodes (as implied by the mutual information between the s_i^{bin} time-series).

The simulated annealing models manages to reproduce the increase in performance as leader nodes are given more edges. The role of leader nodes (with relatively a few edges) is readily described as "noise sources" that manage to insert sufficient random information into the ring to prompt the ring graph nodes to leave local, globally incompatible, solution regimes. The receding, yet prominent, color-change activity of leader nodes (as compared to ring nodes) gives them a role reminiscent of noise in simulated annealing. The performance increase however, eventually levels off. This is an example of the law of diminishing returns: the impact of the first few links from leader nodes to ring nodes is dramatic. With higher leader node connectivity, all neighborhoods, on average, are reached and a further increase in the performance requires a much larger number of links to be added.

As leader nodes gain even more edges, the main qualitative difference between connected and unconnected leader nodes is the decrease in relative performance, observed with unconnected leader nodes ($h = 0$), that appears, once leader nodes are connected in excess of roughly $50 - 60\%$ of possible ring nodes. For the case of $\chi = 4$ mainly reported here, this is the case once the average ring node is connected to 2 to 3 leader nodes.

In this regime of pronounced central coordination, the communication between leader nodes becomes inevitable in reducing incompatible information transmitted onto the ring nodes. The more such intra-leader node coordination is missing (the smaller the value of h), the more leader nodes seem overburdened with providing meaningful information to ring nodes. This captures the concern of Horling and Lesser (2004) about federated agent structures, that the *mediators* (the leader nodes) can become the bottleneck of the coordination process. In this part of the parameter space, we notice a decrease in performance, yielding the often invoked, curvilinear relationship between the degree of "distributed" control and (relative) performance.

This understanding of leadership mechanisms in CAS is in line with previous hypotheses in organization theory — in particular in CLT: The analysis reported here largely confirms Hypotheses H_2 and H_3. The curvilinear performance shape, while supporting Hypothesis H_1, was not directly predicted. The results imply that the mechanisms of leadership expressed in Hypotheses H_2 and H_3, while generally true, are constrained by the information processing and communication capacity of and between agents. The outstanding impact of direct communication channels between leader nodes play in regimes of high involvement of leaders has so far not been discussed in the organization theory literature.

4.6 Relating Toward Application: A Forward Model

The previous section has developed a functional understanding of the differences in relative performance — induced by network structure — for the GCD problem, which was established as a minimal example of a DCSP. It could be shown that the control network architecture has a distinct impact of the ability of agent networks to "agree" on a correct and consistent solution. In this section, the focus is on the second function of control in RASs: ensuring efficiency of the attained solution.

The GCD model in the form explored so far does not provide anything resembling classical PPC performance metrics (c.f. Section 2.1.3). In fact, it does not have any notion of time. This section will present an extension of the GCD problem discussed before, called the *forward model*, that introduces time and allows us to measure not only the quality of the coordination process (as done so far, by measuring the number of color changes required), but the quality of the attained solution. In doing so, the gap between GCD as a DCSP as well as the tasks and performance measures usually applied to PPC can be narrowed. In a forward model setting, the data observed from the original model is acted upon to attain a translated dynamic behavior that can be interpreted in terms of a more realistic application domain. As an example from neuro-sciences, Stephan et al. (2008) use a forward model to "translate" the results of their minimal model of neuronal excitations into the domain of functional magnetic resonance imaging (fMRI), attaining similar statistical properties as observed empirically.

Transferring results of a GCD model to classical job-shop scheduling problems (i.e. no machine flexibility) was first considered by Fabiunke. In Fabiunke (1999) and Fabiunke and Kock (2000), the authors constrain the DCSP by a maximum makespan τ that has to be reached. In their model, agents represent operations and have to agree on starting times (node states) such that the resulting schedule is valid and has a makespan $\leq \tau$. This interpretation of color states as timeslots (c.f. also Section 4.2.2) allows them to introduce a constraint on minimum performance (maximum makespan), without changing the nature of the model from a DCSP to a constrained optimization problem.

4.6.1 Forward Model Setup

To map the GCD model toward the domain of PPC, the following modifications to the original model are applied while maintaining the original interpretation of ring nodes representing operations and colors representing machines: A time component is introduced, which is *not* equal to a round in the GCD model. Instead $N + \chi$ (the total number of nodes present in a model with χ leader nodes) rounds are (arbitrarily) set to equal one unit of time, determining a relative dependence between the speed of agent decision-making (occurring in rounds) and time passing in a distributed scheduling setting.

Operations will be allowed to decide for one machine (color) and enqueue for processing there. Across all machines and operations, a processing time of 1 time unit will be assumed. Note that leader nodes are *not* considered as operations, they are solely added to support (or stifle) coordination. Given that agents can only access local information, the best proxy for the optimality of the local decision (machine assignment) by an agent has to be the number of conflicts with its immediate neighbors. Likewise, the time since the last change in color, the result of a change in the agent's "cognitive map" (Lichtenstein et al. 2006, p. 5), has to be the agent's best available proxy to determine whether a sufficiently stable global solution has emerged to commit to the current local assignment, i.e. enqueue at the currently preferred machine (the machine corresponding to the current color).

For obvious reasons, the optimal makespan is $\frac{N}{\chi}$ time units. Deteriorations may occur for one of two reasons: First, when the stream of operations settling for processing at one machine breaks down, the machine may become idle. As the total work content is constant, this "lost" time leads to a higher makespan. On the other hand, if operations prematurely settle for one machine, an uneven load across the χ identical parallel machines means that some machines would be idle, while others still have to finish the operations that have decided to enqueue there. So, both overhasty and overcautious agent behavior can cause deteriorations in shop floor performance. To model this trade-off, a *freezing time* parameter f is introduced that indicates how many time units (again: not rounds) an agent's "cognitive map" (i.e. color choice), the idea of the best local decision in the context of the observed partial problem, has to have remained unchanged in order for the agent to place sufficient "trust" in its perception and decision to "finalize" it and enqueue at the corresponding machine.

In summary, the forward model creates a scheduling (an assignment of operations to machines and a sequencing of operations on each machine) in accordance with the following rules:

1. A ring node (operation) that has not changed its color during the previous f units of time will be assigned to the machine corresponding to its current color for processing.
2. A node assigned to a machine in this way will be "removed" from the graph, i.e., will not be updated or influence the color changing decisions of its neighbors any more.
3. Nodes assigned to one machine will be processed according to a First In – First Out (FIFO) priority rule.
4. Machines need 1 time step to complete one operation. Additional jobs being queued at that machine during this production time will be delayed accordingly.

4.6.2 Forward Model Results

Given the above setup, the performance (here: the makespan) can be measured as a function of the degree of autonomous control, the connection between leader nodes (h), and the freezing parameter f. Since this is an *absolute* value, no result normalization is applied.

The results are shown in Fig. 4.12. All result plots in this section have an inverted y-scale so that high makespans are visually represented as low performance. A "medium" value of $f = 5$ is investigated first. The results are shown in the second stacked subplot of Fig. 4.12. They resemble the original GCD results (Fig. 4.5(a)). In particular, they exhibit a curvilinear relationship between the degree of autonomous control and performance for unconnected leader nodes and a peak, much further shifted toward centralized control, once leader nodes are connected.

In order to test Hypothesis H_4, the freezing parameter f is now changed. The upper part of Fig. 4.12 shows the result for $f = 1$, a short freezing time. Agents under this setting settle quickly for solutions, making "incorrect" freezes (freezes that load the machines unevenly) more likely. Successful control architectures in this environment must be able to support the quick creation of local solution patches to reduce scheduling conflicts. In the optimum, the performance values for $f = 1$ are hence below those for $f = 5$. More importantly, the peak is clearly shifted to the "right", regions of higher degree of autonomous control. In settings where agents (have to) settle quickly for solutions, it appears that network architectures exhibiting low degrees of hierarchy/high degrees of autonomous control are (relatively) best suited.

In comparison, a situation with *long* freezing times — in particular $f = 10$ — is shown in the lower third of Fig. 4.12. The plot has the same y-axis range as the plots above and thus large parts of the results are invisible. Notably however, the peak performance for both connected and unconnected leader nodes is now reached for very low degrees of autonomous control (leader nodes have many edges). In this situation, where agents are "careful" in settling for a color (machine), a successful control architecture must seek to effectively communicate existing scheduling conflicts in the graph while allowing the nodes to freeze to a color (assign themselves to a machines) by muting unhelpful "noise" transmitted along the communication edges.

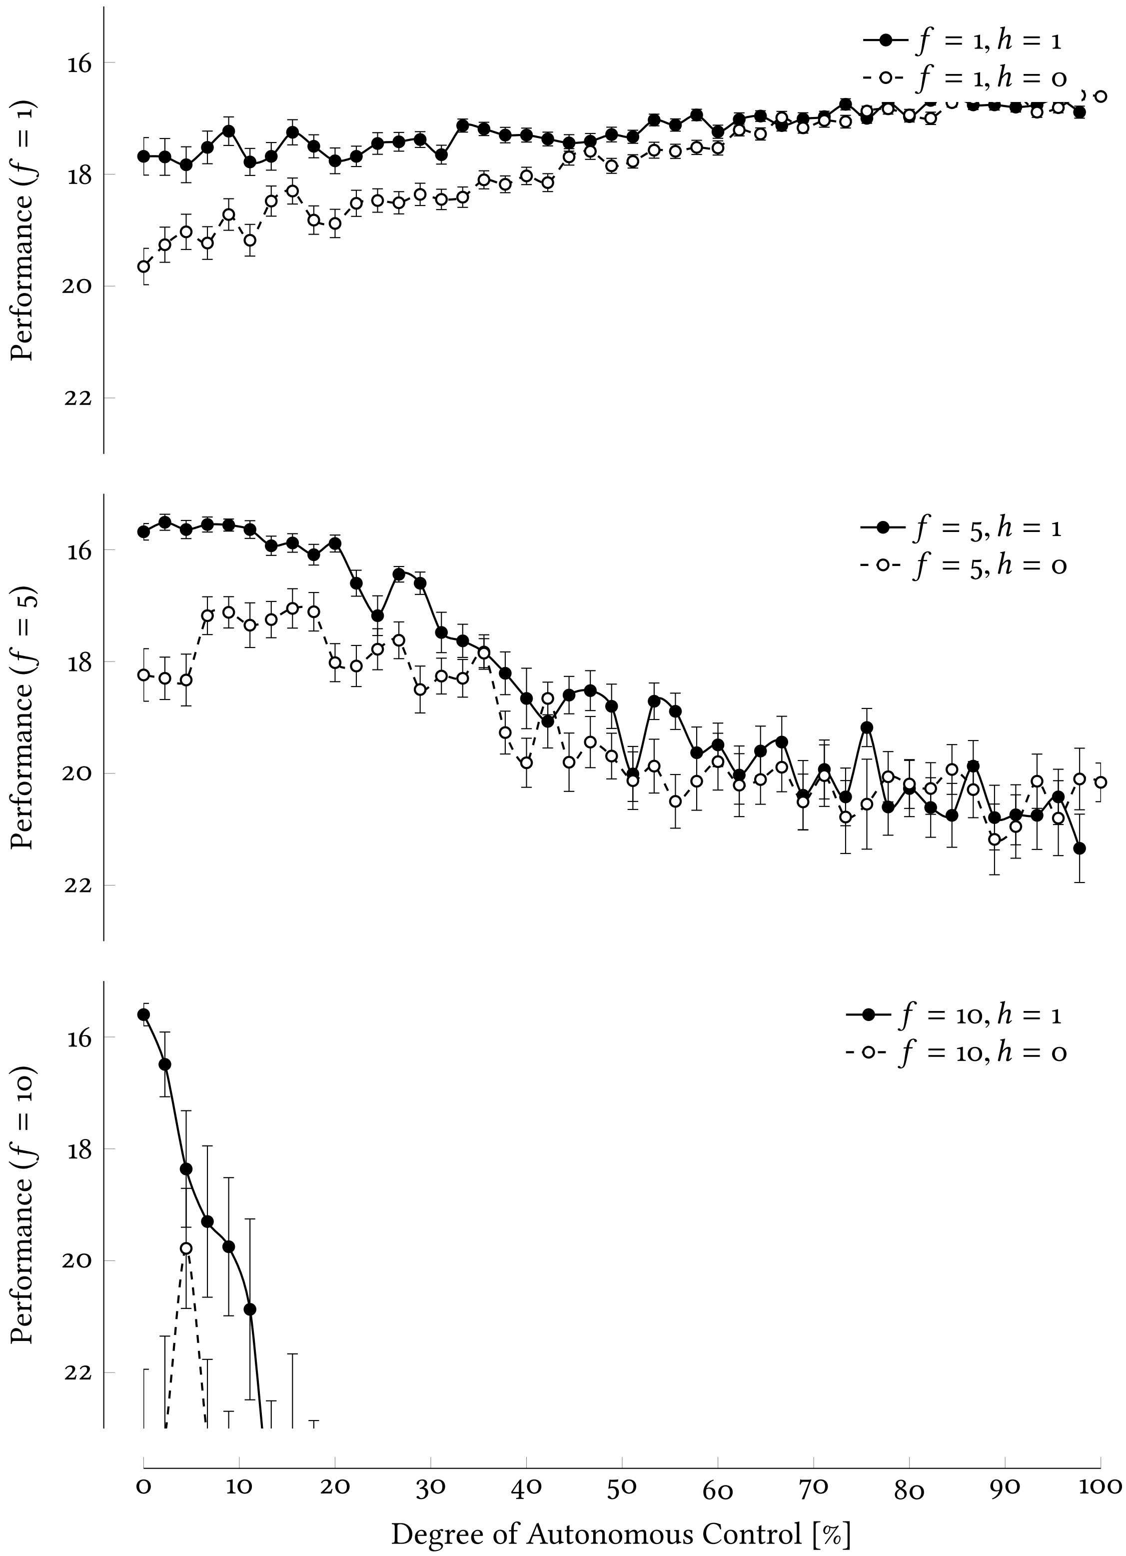

Figure 4.12: Makespan (inverted to show performance) of the Forward Model experiments for freezing parameter values $f = 1, 5, 10$ (top to bottom) as a function of the degree of autonomous control for both unconnected and fully connected leader nodes, $\chi = 4, N = 60, s = 30$, averages taken over 100 samples.

4.7 Discussion

The reported research allows the formulation of a new view on the role of hierarchy as a performance-mediating factor in distributed systems. The results w.r.t. the "optimal" degree of autonomous control as well as its dependence on the optimization horizon give analytical evidence for long-held beliefs not only in the PPC community (c.f. Sections 2.3 and 4.1.2).

Optimal Balance of Centralized and Decentralized Control – the Role of Leadership and Information Time Horizon

The minimal model analysis validates findings from production research and organizational science, indicating that the highest level of performance is achieved for a mixture of centralized and decentralized control. Going beyond the previously established hypotheses, this thesis finds the role of information exchange among leader nodes to be central for both (1) the overall peak performance and (2) the location of the performance maximum on the scale of control architecture.

These findings provide empirical evidence for the propositions of several scholars and validate Hypothesis H_1.

> *H_1: In complex systems, a balance between centralized and distributed control leads to the highest performance.*

Enablers and Mechanisms of Leadership in Complex Systems

The research reported in this chapter shows that (1) the introduction of strategically located nodes (leader nodes) with access to dispersed parts of the organization (the ring graph) can add to the coordination capacity (and overall performance) of a control architecture. (2) Connections between leader nodes that give access to other leader nodes' information can significantly increase this effect. Both findings stress the importance of leader nodes as aggregators and transmitters of global information.

These results validate Hypothesis H_2.

> *H_2: Leaders use information asymmetry and facilitate the coordination process; they enable conflict solving among agents by interaction, exchange of information and emergence.*

By defining solution regimes as circular permutations of colors on the ring, one can analytically conceptualize, visualize, and measure the evolution of local patterns of behavior on the graph. By evaluating the time-series of a leader node, entering or leaving a solution regime, and the time series of the connected ring nodes, one can see a flow of information very similar in shape to the observed relative performance curve.

There is hence strong evidence in favor of Hypothesis H_3.

H_3: Leaders disrupt existing patterns of behavior.

These findings indicate the high importance of hierarchical features in the control architecture even in situations, where nodes in higher levels of hierarchy have no different authority or decision-making capability than their "subordinates". The observed effect is, therefore, an emergent phenomenon mediated (as shown) by the network architecture. The conducted experiments indicate that even in systems relying on emergent organization, some level of hierarchy can be beneficial for system performance, as it provides the ability to aggregate and spread information (c.f. Helbing et al. 2006a).

The investigation of the forward model further allows us to investigate the effect of hierarchy on performance duality in the larger context of changing environmental conditions for companies. The reported results w.r.t. the optimal degree of autonomous control in the forward model are clearly dependent on the freezing time parameter f, representing the time-horizon given for coordination. In particular, more hierarchical control architectures are appropriate in situations where longer time can be given for coordination. As we decrease f, coming close to the "soft-real time" requirements posed on many information systems in the context of production control, there is a clear shift toward more decentralized control architectures. Hypothesis H_4 can, therefore, also be validated:

H_4: The optimal balance of centralized and decentralization control architecture is a function of the optimization time horizon considered.

For industrial practice, these findings strengthen the role of management even in complex perceptions of organizations and indicate that organization structures, communication routines, instruments, etc. should be designed according to the specific structural and dynamic environmental conditions. Leaders can facilitate global optimization in distributed systems by exchanging information and preventing local optima that can emerge due to institutional barriers. The most interesting finding here is that the communication among managers appears to be critical to positively influence performance.

4.8 Directions for Future Research

Four main avenues of future research seem promising:

The first concerns the extension toward different network architectures. This chapter has investigated small-world ring graphs with artificially added levels of hierarchy. A worthwhile extension would be the discussion of networks where "leaders" emerge naturally. The most prominent example being the widely studied class of *scale-free graphs* (Barabási and Albert 1999), where the so-called *hubs*, nodes of outstandingly high degree and centrality, emerge naturally as the result fo a "rich getting richer" network generation scheme. There is an existing body of literature on the role of these hubs in dynamic processes (Goldenberg et al. 2009; Müller-Linow et al. 2008; Santos and Pacheco 2005; Wang and Chen 2002) to which investigation of the GCD model could add. In particular, scale-free graphs do not require a binary classification of nodes into leader and non-leader nodes, as necessary/implied in the investigation of ring graphs. Using the established information transfer as the litmus test of leadership (Section 4.5.2), the GCD model could help researchers to answer the question if and how many hubs in scale-free networks act as leaders. On the side of networks with artificially induced hierarchy, the hierarchical networks presented in (Ravasz and Barabási 2003) provide the opportunity to extend the investigations reported here towards multi-level hierarchies.

Second, the role of information quality seems worthy of investigation. The impact of lost or incorrectly transmitted information in GCD settings has already been investigated in Fitzpatrick and Meertens (2001), but gains additional importance especially in the light of findings on the role of leader node interconnections. Three possible information errors can be distinguished: the erroneous perception of a neighboring node's state, a time-delay in status transmission, and random color changes (errors in true colors). Doing so could enhance our understanding of coordination and leadership as information transmission and computation processes and the impact of network architecture on the effectiveness of these processes.

The third avenue concerns the focus on direct neighbors in the information exchange between agents. Estrada and Vargas-Estrada (2013) show that for continuous consensus problems on graphs and leaders fixed in their state (not adopting to their environment), any node (independent of e.g. centrality) seems evenly capable of leading the coordination process, when agents' decision making also considers their indirect neighbors (larger information horizon).

Finally, the most promising (from a PPC point of view) and most daunting avenue is the evolution toward more realistic application scenarios. The forward model of Section 4.6 has taken an important first step in this direction. However, it falls short of most scheduling scenarios, as it does not consider, e.g. sequencing constraints between multiple operations of one job. A natural extension of the experiments made here would be to apply GCD on directed and disjunct graphs, which would make it possible to apply it also to graph representations of real job-shop (Błażewicz et al. 2000; Roy and Sussmann 1964) and flexible job-shop (Blunck and Bendul 2015; Dauzère-Pérès et al. 1998) scheduling problems.

Chapter Five

Capacity Dimensioning for Distributed Control

> *"It is not always appreciated that in a severely practical subject [...] there is need for theory. However, the history of science suggests that progress in any field of research can best be achieved by a judicious mixture of practical experience, experiment, and theory."*
>
> Wardrop (1952, p. 326)

Previous Publications

The contents of this chapter, in particular the problem description, motivation, and the iterative approach in Section 5.6 have previously been published in:

Henning Blunck, Dieter Armbruster, and Julia Bendul (2016). "Simultaneous Workload Allocation and Capacity Dimensioning for Distributed Production Control." In: *Procedia CIRP* 41. Research and Innovation in Manufacturing: Key Enabling Technologies for the Factories of the Future - Proceedings of the 48th CIRP Conference on Manufacturing Systems, pp. 460 –465. ISSN: 2212-8271. DOI: 10.1016/j.procir.2015.12.117

The analytical approaches (incl. the shown definitions, lemmas, theorems, and corollaries) have been published in:

Henning Blunck, Dieter Armbruster, and Julia Bendul (2017). "Setting production capacities for production agents making selfish routing decisions." In: *International Journal of Computer Integrated Manufacturing*. Special Issue: Cyber-Physical Product Creation for Industry 4.0. DOI: 10.1080/0951192X.2017.1379097

Chapter 3, in particular Section 3.3.1, provides motivation to study the impact of system *design* on the performance of distributed PPC problems. In this context, this chapter discusses the impact of *capacity dimensioning* on the collective behavior of product agents. Unlike in the previous chapter, the goal of the analysis presented here is not to improve a direct measure of performance. Instead, the goal is to attain a (long-term average) predictable flow distribution as the collective result of the decision-making of many selfish acting agents in the sense that it leaves every machine in the system with a-priori determined utilization levels.

Finding that new methods for capacity dimensioning are necessary, this chapter does not only address research question Q_3, but also provides a glimpse into a new approach to manufacturing system design: By truly incorporating a constructionist bottom-up design approach, this chapter highlights the consequences of more strategic decisions (in particular: production system design) when being performed as a function of low-level design decisions on agent behavior.

5.1 Problem Motivation

To understand better, why the capacity dimensioning problem is of interest in the context of distributed PPC, this section will (1) quickly introduce the capacity dimensioning problem (Section 5.1.1) and (2) "traditional" approaches to solve it under a reductionist top-down planning approach (Section 5.1.2). Then, Section 5.1.3 discusses how intelligent products making selfish routing decisions violate the assumptions of the traditional approach.

5.1.1 The Capacity Dimensioning Problem

(Capacity) "Dimensioning is the quantitative determination of capacities/resources such as the number/measurements of all of the equipment [...] and the calculation of costs" (Schenk et al. 2010, p. 81). The process, which includes calculating resource requirements (in units of work content per time period) and subsequent dimensioning of capacities, is an important subproblem in the design of production systems (c.f. Chryssolouris 2006, Ch. 5.1; Schenk et al. 2010, Ch. 3; Koren 2010, Ch. 7.2); it is usually part of a larger sequential factory planning process like the "0+5+X Model" by Schenk et al. (2010).

Capacity decisions have immediate impact on initial investment, operating cost, and the ability of companies to meet future demand, as they provide an upper bound on the amount of goods and services that can be produced and sold (Stevenson 2009, Ch. 5; Hopp and Spearman 2008, Ch. 18). The dimensioning decision is widely acknowledged to impact other performance criteria, such as WIP and throughput time, thereby establishing a trade-off in which the designer of production systems has to navigate (c.f. e.g. Bitran and Morabito 1999; Bitran and Tirupati 1989; Negri da Silva and Morabito 2009; Stevenson 2009, Ch. 5).

5.1.2 The Traditional Approach

As this section will show, models for manufacturing system design commonly use a "fluid modeling approach" where the analysis abstracts from individual items (products) and focuses on aggregated *flows* of products. Such approximation is the natural consequence (on the side of modeling approaches) of the hierarchical decision-making nature in the management of manufacturing systems (c.f. Section 2.1.3) (Chen and Mandelbaum 1994; c.f. also Cassandras and Lafortune 2008, Ch. 11.9), given the strategic nature and importance of the capacity-dimensioning decision.

This section will argue, based on a review of existing literature, that two assumptions are central to the successful application of traditional capacity-dimensioning approaches: (1) flow distribution of products over process path (or machine) alternatives are fixed as an input to the capacity dimensioning process and (2) capacities are calculated based on these flows, adding extra capacity beyond the calculated demand. Figure 5.1 shows the traditional approach formally.

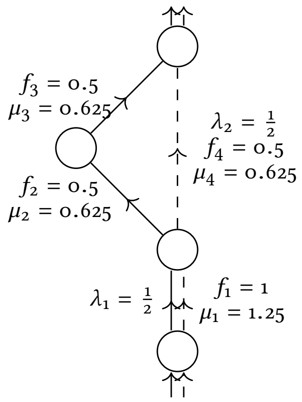

Traditional Approach (Single Commodity Case)

1. Nodes are buffers, edges are machines.
2. Define possible production paths.
   Here: solid and dashed path over 4 different machines combined.
3. Distribute flow demand among (subset of) paths.
   Here: Total flow $\lambda = 1$, path flows $\lambda_1 = \lambda_2 = 0.5$.
4. Calculate capacity demands per machine/edge (f_1-f_4).
5. Set resource capacities (μ_1-μ_4), given u^*, i.e. $\mu_i = \frac{\lambda_i}{u^*}$.
   Here: $u^* = 0.8$ for all edges.

Figure 5.1: Traditional approach to capacity dimensioning. Notably, capacities are set, *after* flow distribution is known, and thus this traditional approach to capacity dimensioning does *not* work when flow distribution is decided at run time! Figure from Blunck et al. (2017).

The Fixed Flow-Distribution Assumption

The traditional way to estimate the amount of required capacity has been to divide the expected work content per time-period to be processed at a given machine by the amount of time available per time period in machine is available for processing. Kusiak (1987) and Miller and Davis (1977) show that academic and practitioner-oriented literature has featured this approach since at least the 1950s.[59]

With the rise of Flexible Manufacturing Systems (FMS) in the 1980s, machine alternatives complicated the problem (hence called the "equipment requirements problem" by Kusiak (1987)), as products could feasibly been produced by different machines, even using different processes. The resulting larger decision space has led to the development of optimization models to find the (cost-) optimal combination of machines for a given production program or to maximize output for a given investment (c.f. Askin 2013; Bard

[59] Miller and Davis (1977) already include an "efficiency of use" factor, which corresponds to the target utilization assumption discussed in the next subsection.

and Feo 1991; Lee et al. 1991; Tetzlaff 1995; Koren 2010, Ch. 7.2; and Lee et al. 2006, for a review).

While these approaches seek to minimize investment cost, the static modeling approach fails to consider dynamic system properties, especially the existence/length of queues in front of machines. These waiting times can be estimated and considered in manufacturing design methods based on queuing networks (reviewed e.g. by Solot and van Vliet 1994; Viswanadham et al. 1992, again focusing on FMS). While for FMS closed queuing networks are commonly used (where tokens represent a limited set of pallets carrying products), open queuing networks for general purpose manufacturing systems have been investigated, e.g. by Chandy et al. (1977), Jackson (1963), Kim and Kameda (1992), and Tantawi and Towsley (1985). Queuing models can be connected with optimization approaches to determine the optimal capacity allocation (in terms of mean latency), given budget constraints (Bitran and Tirupati 1989; Dittel et al. 2011).

Distributed approaches have been taken to the capacity-dimensioning problem as well: A bio-inspired approach to network formation and capacity dimensioning using the slime-mold *physarum polycephalum* is presented by Tero et al. (2010). Argoneto et al. (2008, Ch. 6) present an approach where agents bargain over the distribution of a fixed amount of capacity between machines. The dynamic (and distributed) adjustment of capacities to variations in demand as well as machine failure has been investigated by Duffie and collaborators, e.g. in Duffie and Shi (2010), Duffie et al. (2012), and Jeken et al. (2012), where initial capacities were set “traditionally”.

Importantly, all the presented approaches either take a *flow distribution* (a distribution of product demand over alternative process paths) as input or — in the case of optimization models — establish the flow distribution as an output alongside the capacity dimensioning. The flow distribution is decided during the manufacturing system design stage, and does not account for decision-making at the production execution stage, e.g. through intelligent products (succession instead of hierarchical planning, c.f. Section 2.2.1).

The Target Utilization Assumption

While the above-discussed approaches can calculate the “optimal” distribution of flows across process path alternatives and (by extension) the capacity *demand* by machines, it is widely acknowledged that capacity *supply* should exceed, and not just match, the demand.

This deliberate gap between required and provided capacity is called “tactical underutilization” by Schönsleben (2012, Ch. 13.1.1). The utilization rate more broadly refers to the ratio of used capacity and available capacity, i.e the share of time (a relative number, usually expressed in %) when the machine is busy (c.f. e.g. Papadopolous et al. 1993, Ch. 2.4.1; Hopp and Spearman 2008, Ch. 8.5.1). The intentional decision for underutilization (any utilization below the theoretical optimum of 100%) can be motivated through a variety of considerations:

The first, and most mundane, idea is that the fluid approach to demand calculation will inevitably disregard factors that either add to capacity demand or lower the available capacity. Stevenson (2009, p. 179) differentiates between *effective* and *design capacity*, stating that "Effective capacity is usually less than design capacity owing to realities of changing product mix, the need for periodic maintenance of equipment, lunch breaks, coffee breaks, problems in scheduling and balancing operations, and similar circumstances". Schenk et al. (2010, Ch. 3.2.3) mention, among others, setup times, machine unavailability, time spent producing faulty products, and sick hours. Within the literature on Total Productive Maintenance (TPM), an Overall Equipment Effectiveness (OEE) value measuring the ratio between time spent producing quality output (equal to the calculated capacity demand) and the theoretically available capacity (maximum speed production around the clock) of 85% is already considered "world class" (Dal et al. 2000).

Second, the above-mentioned trade-offs, which are associated with capacity dimensioning, imply that an intentional decision to acquire excess capacity allows the production system designer to position the system along the trade-offs between capacity investment and high utilization on the one hand, and other performance indicators, such as low throughput times and low WIP, on the other hand.

Finally, knowing the utilization rate and either capacity demand or supply allows one to make predictions about the dynamic properties of the respective machine during operation. In particular, the field of *queuing theory* has derived analytical expressions or approximations for many situations to predict dynamical properties — in particular, the *sojourn time* (the total time spent by a product at a machine, both waiting and in process). In general, the expressions will show an increase in sojourn time with increasing capacity demand (utilization) (c.f. e.g. the Overview of results in Papadopolous et al. 1993, Appendix B). For many simple queuing systems, the sojourn time will go to infinity as utilization approaches 100%. Setting a target utilization < 100% then allows to bound the absolute value of the expected sojourn time. Since the first derivative of the clearing function generally increases with increasing utilization as well, limiting the utilization rate also limits the impact of small deviations from the expected utilization on the experienced throughput time.

In this chapter, this intentional under-usage of capacity is conceptualized by assuming a "target utilization" level $u^* \in (0, 1)$ for every machine in such a manner that given capacity demand f, the installed capacity will be $a = \frac{f}{u^*}$, resulting in a capacity utilization of $u = \frac{f}{a} = \frac{f}{f/u^*} = u^*$. Note that the commonly applied idea of intentionally determining bottleneck machines (based on their high investment cost or role for product release, c.f. Goldratt (1999)) is absolutely compatible with this view when such machines are assigned a higher target utilization level. Notably, even for bottlenecks, a 100% target utilization does not lead to optimal target achievement across performance indicators (Chakravorty and Atwater 2006).

5.1.3 Conceptual Problems with the Traditional Approach Under Autonomous Control

One of the major ideas behind the development of distributed PPC systems is to exploit system-inherent flexibility potentials at runtime (Brennan and Norrie 2003; Windt et al. 2010a). Previous chapters have also discussed the notion of selfish decision-making in the context of distributed PPC. Tay and Ho (2008) note that the routing decision (i.e. the decision between alternative production paths) is usually implemented by comparing the queue lengths of the machine candidates which, in turn, depend on the installed capacity. It can, therefore, be assumed that capacity decisions influence the behavior of intelligent products. In fact, the interplay between short-term machine capacity changes and the resulting changes in the behavior of intelligent products has already been studied through both discrete and continuous model classes (c.f. Jeken et al. 2012; Karimi et al. 2010).

The flow distribution in the context of distributed control hence should be considered as a function of the capacity distribution as well. Adding this to the sequential relationship between flow and capacity distribution discussed earlier and assumed in the existing sequential factory planning processes (c.f. Section 5.1.1), we attain a feedback loop between flow and capacity distribution, shown in Fig. 5.2.

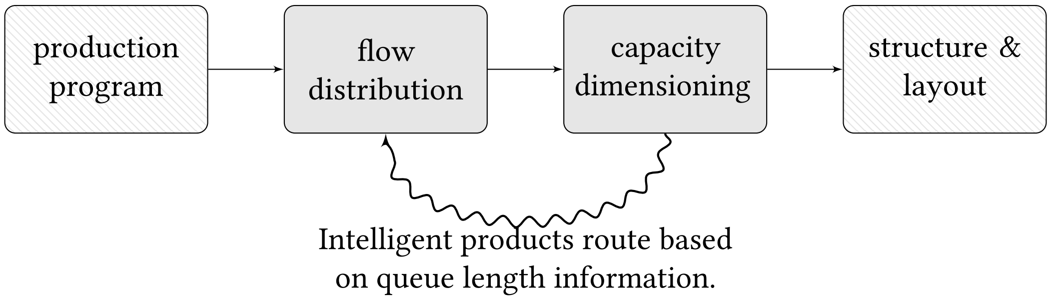

Figure 5.2: Conceptual problem with the traditional capacity dimensioning approach under systems with distributed control. Now, the flow distribution depends on the capacity distribution, forming a circular reference. Planning process adapted from Schenk et al. (2010, Ch. 3.2). Figure from Blunck et al. (2017).

This impact of distributed control on the assumptions of capacity-dimensioning is not a purely academic discussion. As already discussed in Section 1.3, a lack of performance guarantees and predictability is currently hindering the adoption of agent-based distributed PPC systems (Mařík and McFarlane 2005; Trentesaux 2009), with performance predictability only possible through extensive simulation studies (Mařík and McFarlane 2005).

5.2 Goal: Finding Utilization-Attaining Flow-Distributions

Section 5.1.2 has found capacity dimensioning using set target-utilization levels to be widely applied in traditional top-down manufacturing system design methodologies. As capacity dimensioning is part of the wider PPC problem (c.f. Section 2.1.3), setting production capacities, as described in Section 5.1.2, can be considered to represent the "traditional", hierarchical PPC approach. Section 5.1.3 has contrasted this with a new reality where intelligent products make selfish decisions. In this chapter, the idea of intentionally setting target utilization levels as a desirable trait of the traditional hierarchical PPC approach is transferred to the era of distributed PPC.

As Section 5.1.3 has shown, predicting the flow of products and hence the utilization of capacity under a distributed PPC assumption requires to break the circular relationship between capacity and flow distribution constituted by the interplay between classical capacity dimensioning ideas on the one hand and the selfish routing of intelligent products on the other hand. The goal is to define a capacity distribution in such a way that the *resulting* flow distribution of intelligent products, *assuming* selfish decision-making, results in each machine being utilized at the target utilization level. This chapter will call such flow distributions *utilization-attaining*.[60]

This chapter will discuss multiple angles regarding the construction of such utilization-attaining flow distributions. It will try to answer the following questions that are concretions of research question Q_3 in the context of capacity distribution:

1. Do utilization-attaining flow distributions exist?
2. If so, how can they be found/characterized?
3. Can we identify an "optimal" such distribution?
4. Does capacity dimensioning based on utilization-attaining flow distributions negatively influence other desirable properties of manufacturing systems — i.e. are there trade-offs?

5.3 Model Language: Congestion Games on Networks

To understand the interplay between capacity distribution and product flow better, this chapter turns to a model class that combines non-cooperative game theory, which has been shown in Section 1.5.4 to be a good tool to investigate the collective dynamics of agents, and queuing theory, which has already been discussed as a tool to investigate the impact of capacity distribution: This model class is called *network congestion games*. This section will introduce important concepts from both game theory and queuing theory that are necessary to analyze manufacturing systems as congestion games on graphs and provide more reasoning as to why this constitutes a feasible model class for the analysis of distributed PPC systems.

5.3.1 Game Theory: Terms & Definitions

In non-cooperative games, selfish acting agents (called: *players*) seek to adjust their actions to each other (i.e. coordinate) to maximize their individual payoffs. A stable solution to such a coordination game "constitutes an *equilibrium*, a compromise that assures somehow 'maximal' attainment of the different interests of all involved individuals" (Ossowski 1999,

[60]This will be defined more formally in Section 5.4 (Definition 5.1).

p. 23). In particular, an equilibrium in a non-cooperative game among a finite number of players that assigns every player (a mix of) *strategies* (decision alternatives) in such a way that no player has an incentive to change its own strategies for lack of possibility to unilaterally improve its own position is called a *Nash Equilibrium (NE)* (Başar and Olsder 1998, Ch. 3.2). The Nash equilibrium describes for each player the best strategy (mix), i.e. the mix that maximizes profit or minimizes cost, *given* the decisions of all other players. When applied to manufacturing systems under distributed control, the NE helps to approximate in the the collective behavior of product agents long term average taking routing decisions.

In particular, this chapter looks at *congestion games* (introduced by Rosenthal 1973), in which players distribute *demand* over (a subset of) *facilities*. Where multiple players choose the same facilities, they face cost (*congestion*). Players seek to minimize the congestion they observe. Rosenthal already suggested to use their model to analyze the behavior of road users and the interaction between supply and demand. In *network congestion games* in particular, player strategies are equated with paths in a network (Holzman and Law-Yone 1997; Meyers and Schulz 2012). Players with the same set of alternative strategies (paths) represent one *commodity*. Congestion games describe the response of agents, as they make routing decisions and face congestion/cost in return. Hence, this particular game is a good representation of the collective behavior that one can expect from intelligent products which have to make routing decisions in response to (and with the intent to avoid) congestion in terms of waiting times at machines.

Congestion games are part of a larger class of games called *potential games* (Monderer and Shapley 1996), which requires the existence of a real-valued function that for each user identifies the change in utility resulting from a deviation from its current strategy (Altman et al. 2006). Thus, it allows to identify equilibria as the result of an optimization problem (c.f. Sandholm 2001).

Notably, the players in the (network) congestion games of Rosenthal (1973) and in the potential games of Monderer and Shapley (1996) are individual users placing their individual flow on one decision alternative. As this research again tries to abstract from individual users (or agents), the number of players increases, while each players contribution to the congestion decreases. In its limit, it reaches a situation where the flow is *non-atomic* (Altman et al. 2006). The equilibrium in non-atomic congestion games is described by "Wardrop's first principle" for the application to users of road networks.

> *"The journey times in all routes actually used are equal and less than those which would be experienced by a single vehicle on any unused route."*
>
> — Wardrop (1952, p. 345)

This situation (also known as Wardrop equilibrium, c.f. Haurie and Marcotte 1985) can be interpreted as the limit of a game NEs as the number of players increases (Altman et al. 2006; Sandholm 2001). Alternatively, we can shift attention from individual users to commodities and the *flow* constituted by all individual users belonging to this commodity: Instead of a single particle (user), we identify a player as a commodity, distributing a given amount of flow across a set of path alternatives. The strategy space of such player then contains all feasible flow distributions along path alternatives. In particular, the sum of all commodity flows has to add up to the commodity *demand* and negative flows are not permissible. As this perspective presents again a situation with a finite number of players, a connection between the concepts of Wardrop and Nash equilibrium is established, as proposed by Haurie and Marcotte (1985) (c.f. also Altman et al. 2006; Marcotte and Patriksson 2007). This NE interpretation is applied throughout this chapter.

In the domain of traffic systems research, game-theoretic models have been used to investigate the Network Design Problem (NDP) (c.f. Yang and Bell 1998, for a review). Most closely resembling the model investigated in this chapter is the Continuous Network Design Problem (CNDP), which seeks to find Pareto-efficient design alternatives w.r.t. investment cost (in edge capacity) and the latency of the induced Nash flow (Gairing et al. 2014). It was first investigated in the late 1960s (Dafermos 1968; Dafermos and Sparrow 1969) and is still known to be "one of the most difficult and challenging problems in transport" (Yang and Bell 1998, p. 257). Gairing et al. (2014) showed it to be $\mathcal{NP}$-hard. The NDP is a game-theoretic extension of the Continuous Capacity and Flow Assignment (CFA) problem that considers the simultaneous allocation of capacities and flows in traffic and communication networks to find Pareto-efficient combinations of cost and mean delays without assuming selfish behavior on the side of the road users/messages (flows are pre-determined, not in an equilibrium) (Queiroz and Humes 2003). For a review of the CFA problem, the reader is referred to Queiroz and Humes (2003) and Shen et al. (2005).

5.3.2 Queuing Theory: Terms & Definitions

Queuing theory is a field of applied probability theory concerned with the analysis of systems where users are awaiting service, potentially forming queues of waiting users (c.f. e.g. Allen 2014, Ch. 5.0; Kleinrock 1975, Ch. 1). It can provide the cost functions assumed in (network) congestion games to express the experienced cost (congestion) as a function of flow. Analytically known cost functions provide the basis to prescribe the behavior of machines under different levels of workload and hence the basis to prescribe the effect of agent-routing decisions on the overall system.

Among the many properties investigated by researchers in the field is *total time in the system* (also known as the sojourn time), which is the sum of waiting and processing time experienced by the user in the system (Papadopolous et al. 1993, Ch. 2.4.2). In the domain of manufacturing system design and analysis, this time is commonly known as the throughput time (Nyhuis and Wiendahl 2009, Ch. 2.1.2) (where post-process waiting times and transport times are disregarded).

Knowing — and reducing — the throughput time is not only a goal in itself within the commonly assumed goals of production logistics (Section 2.1.3); it is also directly (and linearly) connected with the WIP in the system.[61] For its ability to predict throughput times as a function of design decisions, queuing theory is a well-established tool for the design and analysis of manufacturing systems (Papadopolous et al. 1993; Papadopoulos and Heavey 1996; Subba Rao et al. 1998).

[61]Through one of queuing theory's most fundamental results, *Little's Law*, named after Little (1961)

5.3.3 Modeling Production Systems Under Autonomous Control as Congestion Games

Congestion games are a frequently used tool to investigate networked systems such as communication networks (Altman et al. 2006) and road-traffic systems (Beckmann et al. 1955; Sandholm 2001). Notably in both cases, the agents representing network users (be it road users or data packages) are modeled as competing and selfish. Since both traffic and production systems can be viewed as coupled dynamic queuing systems (Helbing et al. 2006b), it seems reasonable to also consider congestion games as a model class for modeling production systems, especially when under autonomous control, products are likewise considered to act in their own interest (c.f. Sections 2.1.2 and 3.5.3). One can point to a number of similarities to support this argument:

- Intelligent products are frequently believed to make routing decisions to minimize throughput time (Tay and Ho 2008). This is in line with the assumptions of congestion games (minimizing congestion) as well as assumptions in road traffic modeling (Correa et al. 2004).
- Consequently, NEs have been used to predict user behavior in traffic environments (ibid.).
- It is generally assumed that the latency associated with a given flow on a given edge is increasing in flow. This also holds for all analytical expressions for sojourn time known to the author (c.f. Papadopolous et al. 1993, Appendix B).
- Total throughput time for a product through a manufacturing system is *additive* in the sense of Raith et al. (2014, Def. 1) in the same way that the total travel time can be expressed as the sum of travel times per road section.

There also exists a stepwise process for traffic network design (McNally 2000) that includes the task of traffic assignment (Raith et al. 2014), comparable to the above-discussed manufacturing system design processes (Schenk et al. 2010, Ch. 3).

5.3.4 The Inefficiency of Selfish Behavior in Game Theory

Another major benefit of congestion games in the context of this thesis is that the effect of myopic player behavior has been studied intensively and can usually be calculated analytically. Section 3.1.3 has already pointed at multiple avenues to assess the impact of myopic behavior, concluding that for frequently used modeling languages (like DES) only result comparisons can be attained. In this chapter, a modeling language is applied for which the socially optimal (non-myopic) solution can be analytically derived.

Dubey (1986) showed that NEs do not have to maximize social profit (or minimize social cost). That is, there can be a flow distribution — not as the result of the interaction of selfish agents, but central allocation — with a better social cost. In the context of network congestion games, a popular measure of social cost is the combined latency experienced across all paths and is consequently calculated by multiplying the latency of a path with the flow routed over this path (Roughgarden 2005, Ch. 2.1). A famous and obvious case in the domain of congestion games on graphs is the so-called *Braess' Paradox* (Braess 1968; c.f. Kolata 1990, for a report on a "real-life" experiment; and Cohen and Kelly 1990, for an example using queueing networks) that has caused ample discussion and research in the domain of traffic system design. It describes a situation, where *adding* capacity (in the form of an additional road) to a traffic system deteriorates the NE (in terms of latency).[62]

For the congestion games covered in this chapter, both the NE and the socially optimal flow distribution can be calculated as the result of an optimization problem (c.f. Section 5.7). The "price of anarchy" (Section 3.1.3) is, therefore, easy calculable. This has led to significant academic attention to characterize and bound the price of anarchy:

Upper bounds on the price of anarchy were first proven by Roughgarden and Tardos (2002), with the idea being quickly extended and applied to other cost function classes (Aland et al. 2006; Czumaj et al. 2010), games with maxima on edge flows (Correa et al. 2004), finite games (Christodoulou and Koutsoupias 2005), and other generalizations (Aland et al. 2006). Lin et al. (2011) discuss upper bounds for Braess's Paradox when the maximum latency (not the total latency) is considered as social cost measure. Moreover, the discussed papers on selfish routing under incomplete information provide upper bounds on the price of anarchy (Gairing et al. 2008; Papadimitriou and Valiant 2010). With so much attention given to the theoretical maxima of the price of anarchy, it is important to note that the gaps between social optimum and NE observed in *reality* are often much lower and only arise significantly for a small range of flow values (Papadimitriou and Valiant 2010). Youn et al. (2008) measure the price of anarchy in real road-traffic networks, finding that the observed price of anarchy is significant only for a rather small range of flow intensities.

Research has also been published on strategies that seek to minimize or avoid the price of anarchy during system design: Korilis et al. (1995, 1997b, 1999) study the addition of capacity to existing networks in such a way that the network performance is improved (Braess' paradox is avoided, c.f. Section 3.3.1). Dittel et al. (2011) provide an optimization algorithm to choose (from a discrete set of alternatives) the subset of edges in the process network to enable (set up any capacity) and which capacity (from a discrete set of options) to buy. Starting from the opposite angle, Roughgarden (2001) ponder the question as to which edges (roads) should be *removed* from a network to find the best possible NE (in terms of combined latency), showing this design problem to be $\mathcal{NP}$-hard even for linear cost functions. Roughgarden (2005, Sec. 3.5) shows that the cost of anarchy vanishes when the network is designed for low utilization levels (the flow is small as compared to the capacities), which is of little practical applicability in manufacturing system design owing to its inherent inefficiency.[63]

[62]Another famous example in the more general domain of game theory is the prisoner's dilemma game (c.f. e.g. Başar and Olsder 1998, Ch. 3.2), where cooperation among prisoners would optimize the social (overall) profit, yet this strategy combination is not a NE, as both players have an incentive to deviate.
[63]A comparable result is also formulated by Helbing et al. (2006b).

5.3.5 Model Assumptions & Limitations

Despite the encouraging parallels and analytical opportunities, the congestion game modeling language also imposes some relevant assumptions and restrictions when applied to distributed PPC:

The first set of assumptions arises from the translation of the usual assumptions about NEs: NEs (in their basic form) assume global information on the side of agents. This translates into the necessity in the production case for product agents to have the knowledge of the latencies expected at all machines. Section 3.1.2 has discussed the conceptual importance of local information in the context of distributed PPC systems (as do e.g. Leitão 2009; Trentesaux 2009; Schneeweiss 2003a, Ch. 1.1). However, Chapter 3 — in particular, Section 3.5.1 — has also presented various possibilities to provide agents with complete information[64] and agent-communication and -interaction can likewise provide global information: Armbruster et al. (2006), Papadimitriou and Valiant (2010), Peeters et al. (2001), and Van Dyke Parunak (1997) discuss the ability of stigmergy based approaches to provide agents with global information. That product agents have sufficient information about the expected latencies across all path alternatives hence seems a reasonable assumption for many practical implementations of distributed PPC systems.

Investigations into selfish routing under *incomplete* information have been made, e.g. by Papadimitriou and Valiant (2010), who present an iterative approach to determine the Nash flow distribution for single commodity instances, and Gairing et al. (2008), who discuss the existence of equilibria under incomplete information on systems of parallel edges. However, this research will focus on the above-described classical case, given the better analytical tractability and richer set of established results to build upon.

It is further assumed that all agents experience the same latency to cross an edge (i.e. a machine) independent of their respective commodities. As waiting times in queues are concerned, this is the logical consequence of a FIFO prioritization rule. However, it also assumes that processing times at the machines are equal for all products. While this is generally not true, this non-discrimination between product types is a commonly applied simplification in the analysis of open queuing networks (Papadopoulos and Heavey 1996). It is also worthwhile to remember that waiting times account for the majority of throughput time experienced in job-shop environments: Stommel and Kunz (1973) perform a study in individual and small-series manufacturing and report waiting times up to roughly 85% of the total throughput time. Theoretical and practical considerations for job-shop environments lead Wiendahl et al. (2009, Ch. 4.2) to assume similar values. Waiting time shares of 85 – 90% are said to be "frequently quoted" by Spencer (1991, p. 23).

In the following, NEs will provide the analytical means to predict long-term agent behavior (the equilibrium) as a function of machine capacities. The usefulness of such long-term average is of course based on the assumption that the system requirements (product mix and volumes) are stable. Such assumption is at odds with the analysis of the manufacturing

[64]Note that product agents only need to access information about the current state information machines on at least one of their path alternatives, no coordination or scheduling is assumed.

environment in Section 1.1, that has pointed to increasingly volatile and short-lived market conditions. Levy (2000, p. 79) paraphrases the problem: "Even the most complex game theoretic models are only considered useful if they predict an equilibrium outcome. By contrast, chaotic systems do not reach a stable equilibrium". However, this chapter has focused here on a subproblem of PPC that, by the long-term nature of its decisions has to conceptualize a "long-term average" and take decisions based on that. In that sense, applying game theory seems appropriate for the *design* of manufacturing systems, even under distributed PPC.

5.4 Model Notation

The remainder of this chapter will discuss the problem of finding capacities in a manufacturing system using the notation developed in the context of algorithmic game theory and traffic research (c.f. in particular Raith et al. 2014; Roughgarden 2005). Stating the problem in mathematical terms will allow re-formulation of the goal of finding utilization-attaining flow distributions (Section 5.2), more precisely at the end of this section (Definition 5.1). The notation was already used in Fig. 5.1. Under this modeling approach, machines are equated with edges in a directed graph. This semantical mapping is clearly inspired by roads as the sources of latency in traffic networks. For production environments, a machine-centered graph might be advantageous for some material flow networks. However, the developed mathematical relationships and results are independent of the network representation.

Bold printed symbols represent vectors (lower-case letters) and matrices (upper-case letters) respectively.

- Let the tuple $(G, r, \boldsymbol{c})$ be a capacitated production network, represented as a directed graph $G = (V, E)$. $\boldsymbol{c}$ is a vector of cost (latency) functions given for each edge (machine) $e \in E$ and parametrized by the machine capacities (see below). r is a set of *commodities* (product types) that can be produced in G.
- Let $\mathcal{P}$ be the set of all feasible paths between a source and a sink in G. In particular, let $\mathcal{P}_i$ be the set of paths that can be used by the ith commodity.
- Let $\rho = |\mathcal{P}|$ be the number of paths.
- Let $m = |E|$ be the number of possible edges (machines) for which a capacity has to be dimensioned (the capacity may be 0).
- Then $\boldsymbol{\Phi} \in \mathbb{N}^{m \times \rho}$ (where $\mathbb{N}$ is assumed to contain also 0) is the *flow-path incidence matrix*, where elements $\phi_{i,j}$ indicate the number of times path j visits machine i.
- Let $\boldsymbol{\lambda} \in \mathbb{R}^{\rho}$ be the flow distribution among paths. Then λ_P is a scalar, indicating the flow on path alternative $P \in \mathcal{P}$; $\lambda_P \geq 0$.
- The required flow (demand) for commodity i is given by r_i. i.e. $\sum_{P \in \mathcal{P}_i} \lambda_P = r_i \; \forall i$.

- Let $\boldsymbol{f} \in \mathbb{R}^m$ be the edge flow distribution that results from path flow distribution $\boldsymbol{\lambda}$. $\boldsymbol{f}$ can be calculated as
$$\boldsymbol{f} = \boldsymbol{\Phi} \cdot \boldsymbol{\lambda}. \tag{5.1}$$
Then vector elements f_e are scalars, indicating flow on edge e. Given the constraints on $\boldsymbol{\lambda}$ and $\boldsymbol{\Phi}$, $f_e \geq 0 \, \forall \, e$.

The cost functions $c_e(f_e) \in c$ are assumed to be *semi-convex*[65] in the flow on that edge (machine) and assumed to be parametrized by some non-negative value, a_e, which represents the *capacity* of that machine. The vector $\boldsymbol{a}$ of capacity values of all machines will be referred to as the *capacity distribution*. The observed *utilization* of an edge with capacity a_e and flow f_e is calculated as $u_e = \frac{f_e}{a_e}$. The desired target utilization for one edge is denoted with u_e^*.

Given a vector $\boldsymbol{c}$ of edge latencies, we can calculate the latency of a path as the sum of the constituting edges (Raith et al. 2014), multiplied by the number of times, this edge (machine) appears in the path. The resulting path-latency vector $\boldsymbol{c}_P$ is then calculated as
$$\boldsymbol{c}_P = \boldsymbol{\Phi}^T \cdot \boldsymbol{c}. \tag{5.2}$$

One can now re-phrase the goal of this chapter (Section 5.2) in mathematical terms, formally defining utilization-attaining flow distributions as follows:

Definition 5.1 (Utilization-attaining flow distribution). *A flow distribution $\boldsymbol{f}$ is called utilization-attaining if the utilization of all machines with positive flow is equal to some constant u_e^* ($u_e^* \in (0,1) \; \forall e | f_e > 0$).*

Note that this definition explicitly accounts for machines that are not assigned any flow ($f_e = 0$) and whose utilization level cannot reasonably be defined. As this research considers the capacity dimensioning process — i.e. a part of the strategic design process — it will be assumed that such machines are not acquired.

5.5 Considered Latency Function Classes

As for the nature of the cost functions c_e, no further assumption beyond semi-convexity has been made so far. This chapter will discuss two cost function classes that are introduced with the respective expressions for the sojourn time below.

Affine Cost Function

This simple cost function class has received significant attention in the scientific debate on congestion games (c.f. Roughgarden 2005, Sec. 3.7 for references) due to its simplicity and mathematical tractability. The latency of a machine in the affine cost function class is given as
$$c_e(f_e) = f_e / a_e + b_e \tag{5.3}$$

[65] This is equivalent to demanding that $c_e(f_e) \cdot f_e$ is convex (Roughgarden 2005, Definition 2.4.2).

M/M/1 Cost Function

The $M/M/1$ cost function expresses analytically the behavior of a single server under exponentially distributed inter-arrival[66] and processing times. The $M/M/1$ queuing model is a common base model assumed in queueing theory (Arnold and Furmans 2007, Ch. 4.2). Kleinrock (1975, p. 94) calls it the "simplest nontrivial interesting system". Unlike the affine case, latencies in an $M/M/1$ system "explode" as the utilization approaches 1. In particular, the sojourn time is analytically expressed as

$$c_e(f_e) = \frac{1}{a_e - f_e} \tag{5.4}$$

and hence has a singularity for $f_e = a_e$. Generally, the sojourn time expression (as well as statements about semi-convexity) assume that $f_e < a_e \leftrightarrow u_e < 1$. This is commonly known as the stability condition (ibid., Ch. 3.2).

Beyond its popularity and simplicity, the $M/M/1$ cost function class provides a number of other benefits that make it worthwhile to consider:

- The $M/M/1$ cost function class is a special case of the more general cost function classes (incl. $M/M/n$, $M/G/1$ and $G/G/1$), for which sojourn time functions can be expressed as generalizations of the expression shown above (c.f. Kingman 1961; Allen 2014, Ch. 3.4).
- Poisson processes are frequently assumed in lieu of more precise information to describe "realistic" (not optimal) situations in shop floors. Hopp and Spearman (2008, Ch. 7.3.3) assume exponential processing times in their "practical worst-case" scenarios. Poisson arrival processes are also commonly assumed when modeling the arrival of independent customers/orders (Kleinrock 1975, Ch. 2.5; Wolff 1982).
- Since the output process and the split or merge of output/input streams can easily be expressed analytically (Papadopolous et al. 1993, Ch. 2.2.2), networks of $M/M/1$ queuing systems can be analyzed, forming, e.g. job shops or road networks (c.f. Jackson 1963).

5.6 An Iterative Approach

A first line of attack toward finding utilization-attaining flow distributions is to iteratively find Nash flows, given a capacity distribution and then adjust the capacity distribution in such a way that the calculated flow is utilization-attaining. Then re-calculate the Nash flow for this new capacity distribution, and so on and so forth, until an (almost) utilization-attaining flow distribution is reached. Such iterative alternation between flow assignments with a fixed capacity and capacity assignment assuming fixed flows was also among the first solution approaches explored for the CFA (Fratta et al. 1973).

[66] Time between two subsequent jobs arriving. Arrival processes with exponentially distributed inter-arrival times are also known as a *Poisson arrival processes*

In Blunck et al. (2016), such iterative algorithm to identify approximately utilization-attaining flow distributions is proposed. These are flow-distributions that attain at every machine with positive flow a utilization that is *close* (e.g. $\pm\epsilon$) of u_e^*. The proposed approach is shown in Algorithm 5.1.

Given a capacity distribution, the Nash flow is approximated using the Method of Successive Averages (MSA). The algorithm (Powell and Sheffi 1982; Sheffi and Powell 1981) is widely used for the calculation of equilibria in traffic applications (Sbayti et al. 2007). The algorithm is in itself iterative (the inner loop in Algorithm 5.1 from lines 5 to 12). The idea is to calculate the path cost, given an initial flow distribution over the capacity and to create an alternative flow distribution $\boldsymbol{h}$ using an "All-or-Nothing assignment" in step 9):

Definition 5.2 (All-or-Nothing Assignment)**.** *Given a capacity distribution $\boldsymbol{a}$ and a flow distribution across path alternatives $\boldsymbol{\lambda}$, the All-or-Nothing Assignment policy creates a new path flow distribution $\boldsymbol{h}$ that for all commodities assigns all flow to the path alternatives with the lowest latency, given $\boldsymbol{a}$ and $\boldsymbol{\lambda}$.*

The new flow is a linear combination of the originally assumed flow and $\boldsymbol{h}$ (step 10). Since the weights used to calculate the next iteration's flow distribution change over time, the MSA converges to a flow distribution that is a Nash flow (Powell and Sheffi 1982).

Algorithm 5.1 Iterative Capacity and Flow Adjustment (Blunck et al. 2016)

1: $\boldsymbol{\lambda}^0 \leftarrow \boldsymbol{\lambda}_{initial}$
2: **repeat** ▷ Main iteration loop between capacity and flow adjustment
3:     $\boldsymbol{a} \leftarrow \frac{\Phi \cdot \boldsymbol{\lambda}^0}{u^*}$ ▷ Adjust capacity to flow
4:     $i \leftarrow 0$
5:     **repeat** ▷ Apply MSA
6:         $\boldsymbol{f} \leftarrow \Phi \cdot \boldsymbol{\lambda}^i$ ▷ Calculate edge flows from path flows, Eq. (5.1)
7:         $\boldsymbol{c}_e \leftarrow \boldsymbol{c}(\boldsymbol{f})$ ▷ Calculate edge cost. Note that $\boldsymbol{c}$ depends on $\boldsymbol{a}$
8:         $\boldsymbol{c}_P \leftarrow \Phi^T \cdot \boldsymbol{c}_e$ ▷ Calculate path cost from edge cost, Eq. (5.2)
9:         $\boldsymbol{h} \leftarrow \textsc{AllOrNothingAssignment}(\boldsymbol{c}_P)$ ▷ c.f. Definition 5.2
10:         $\boldsymbol{\lambda}^{i+1} \leftarrow (1-\theta) \cdot \boldsymbol{\lambda}^i + \theta \cdot \boldsymbol{h}$ ▷ With $\theta = \frac{1}{i+1}$
11:         $i \leftarrow i + 1$
12:     **until** Convergence criterion is met ▷ convergence criterion: $\lambda^{i+1} \approx \lambda^i$
13:     $\boldsymbol{f} \leftarrow \Phi \cdot \boldsymbol{\lambda}^i$ ▷ Calculate edge flows from path flows
14:     $\boldsymbol{u} \leftarrow \frac{\boldsymbol{f}}{\boldsymbol{a}}$ ▷ Calculate actual machine utilization
15:     $\boldsymbol{\lambda}^0 \leftarrow \boldsymbol{\lambda}^i$
16: **until** $u_e \in [e_e^* \pm \epsilon]\ \forall e \in E$ ▷ Convergence criterion for capacity assignment

5.6.1 Application to "real" Production Networks

In Blunck et al. (2016), Algorithm 5.1 is applied to five real production networks described by Becker et al. (2014).

Path alternatives are generated by reading production feedback data that gives the products' routings through the system plant. Commodities are then formed by assuming that all products which share the same first and last machine (the same start- and end-node) belong to one commodity. The set of different paths taken by these products forms the set of path alternatives for the commodity. While this is almost certainly not the correct representation of the real system, this straight-forward approach can be applied in situations where only feedback data is available and is likely to *over*estimate the number of path alternatives (system flexibility).

To speed up the solution process, "Pareto-dominated" path alternatives are deleted before the run. A path is considered dominated *if and only if* there exists another path of the same commodity that is shorter (visits fewer worksystems) and the set of machines used in the alternative path is a subset of the machines used in the dominated path. The total flow demand r_i for the commodities is set to the number of products observed between the start- and end-node.

As an initial flow distribution (Step 1 in Algorithm 5.1), the "Equal Share Assignment" rule (Wang et al. 2010; see also Raith et al. 2014, Sec. 4.2.1) is used; this rule distributes the commodity demand evenly among all alternative paths for that commodity.

$M/M/1$ cost functions are assumed for all machines. However, since in the MSA algorithm, a violation of the stability condition (flows in excess of capacity) cannot be ruled out, the cost function is slightly modified to remove the singularity (when flow is equal to capacity), while maintaining its strictly increasing character. The following latency function $c_e(f_e)$ follows the original $M/M/1$ latency function (referred to as c_{org} in Eq. (5.5)) for $u \leq 0.98 \leftrightarrow f_e \leq 0.98a_e$. Above that mark, c_{org} is approximated by the first-order Taylor approximation around $f_e = 0.98a_e$, which is well defined even for $f_e \geq a_e$. The adjusted latency function is then piecewise defined as:

$$c_e(f_e) = \begin{cases} 10^{20} & \text{if } a_e = 0 \\ c_{org}(f_e) & \text{if } a_e \neq 0 \,\&\, f_e \leq 0.98a_e \\ c_{org}(0.98a_e) + (f_e - 0.98a_e) \cdot c'_{org}(0.98a_e) & \text{else} \end{cases} \tag{5.5}$$

where c'_{org} refers to the first derivative of the original $M/M/1$ latency function w.r.t. flow. Equation (5.5) also handles situations where the capacity of a machine is set to 0 during one algorithm iteration. By assigning these machines a forbiddingly high latency, flows over 0-capacity machines are prevented. In the reported experiments, a target utilization level of $u^* = 0.8$ is assumed across all machines and a flow distribution is considered sufficiently close to utilization-attaining, if the realized utilization levels across all machines fall within $u^* \pm 2\%$.

Fig. 5.3 shows the attained utilization levels over the iterations steps (outer loop of Algorithm 5.1). In the last iteration, the stop condition is reached, as all utilization levels are sufficiently close to the 0.8 target.

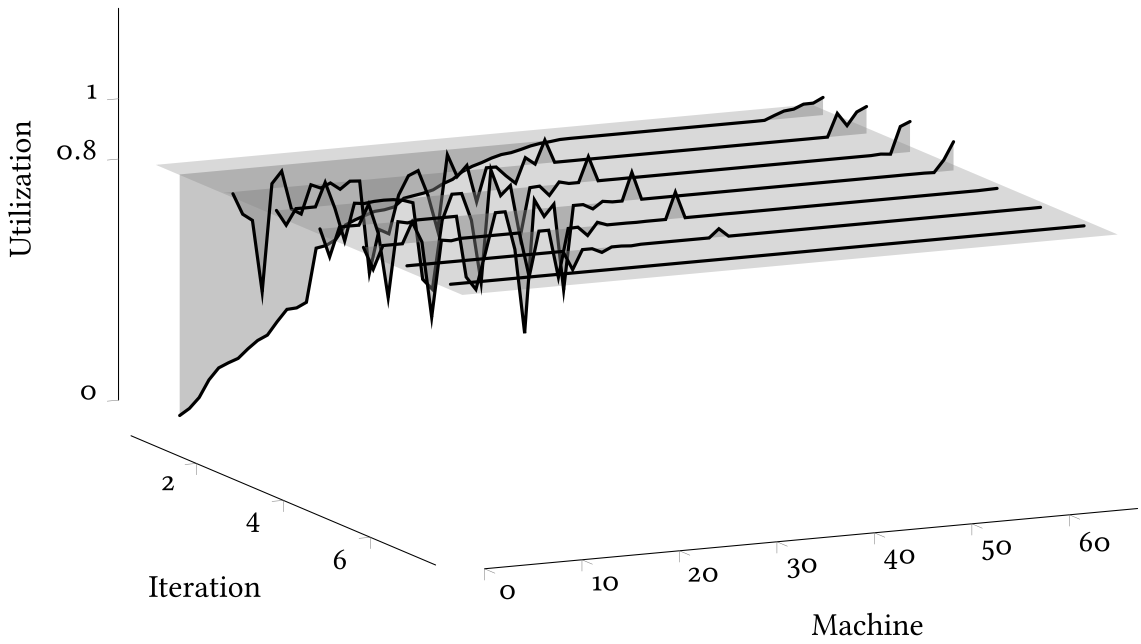

Figure 5.3: Attained utilization levels during an iterative capacity allocation process (shown here for Company F). After iteration 7, all utilization levels are within the ±2% of the target utilization level.

5.6.2 Problems with the Iterative Approach

While the above-presented iterative approach was demonstrated to converge to a capacity distribution in reasonable time (for none of the five tested company networks, more then 20 capacity-assignment iterations were needed), the iterative heuristic exhibits several disadvantages that motivate further investigation of the problem in the following sections.

- The solution depends on the initial flow distribution. This is a common problem for iterative approaches and testing multiple start solutions is hence advised to come to reasonable solutions (Queiroz and Humes 2003).
- In particular, the prohibitive latency associated in Eq. (5.5) with machines with capacity 0 implies that any machine once removed from the set of actually installed machines (those with positive capacity), will not return to the set in further iterations, adding to the path-dependent nature of the system evolution over iterations.
- No guarantee can be given for the existence of a (mostly) utilization-attaining flow distribution.
- It is unclear if *exact* solutions ($\epsilon = 0$) exist.
- No knowledge can be extracted about the properties of utilization-attaining flow distributions.
- The price of anarchy associated with the found solution is no clear. I.e. are there better or worse utilization-attaining flow distribution in terms of social cost? Can a socially preferable (but not utilization-attaining) flow distributions be constructed on the capacity distribution found through the iterative approach?

5.7 Finding Flow Distributions given Fixed Capacities

Before further investigating the capacity-dimensioning problem under the assumptions of selfish products and target-utilizations, an analytical way of calculating both the socially optimal and Nash flow distribution, given fixed capacities, is introduced here.

Analytical expressions to attain the two types of distributions, assuming semi-convex latency functions, have been known for over 50 years (as developed in Beckmann et al. 1955, Ch. 3; Dafermos and Sparrow 1969; see also Roughgarden 2005, Proposition 2.6.1). The combination of convex target functions and linear constraints places both optimization problems introduced below into the set of *convex optimization problems*, which are computationally easy to solve (Boyd and Vandenberghe 2004, Ch. 1.3). In particular, Beckmann et al. (1955) showed that for any transportation problem (choice of network, commodities, flows, semi-convex latency functions), both optimization problems have a clearly defined optimal solution (the solution space is not empty). Both the NE and the social optimum are hence always defined and obtainable as the result of the optimization problem.

Socially Optimal Flow Distribution

It is known that the socially optimal flow allocation, given a set of capacities, can be found by solving the following optimization problem (Roughgarden 2005, Sec. 2.4):

$$\min_{\boldsymbol{\lambda}} \quad \sum_{e \in E} c(f_e) \cdot f_e \tag{5.6a}$$

subject to

$$\sum_{P \in \mathcal{P}_i} \lambda_P = r_i \qquad \forall i \tag{5.6b}$$

$$f_e = \boldsymbol{\Phi}_{(e)} \cdot \boldsymbol{\lambda} \qquad \forall e \in E \tag{5.6c}$$

$$\lambda_P \geq 0 \qquad \forall P \in \mathcal{P} \tag{5.6d}$$

where $\boldsymbol{\Phi}_{(e)}$ denotes the e'th row vector of the path incidence matrix.

Nash Flow Distribution

The Nash flow distribution can be calculated with the identically constrained optimization problem, but with a different target function. As already mentioned, congestion games are a special case of potential games for which the NE can be found using a "Potential function". For the case of congestion games, this potential function is (Beckmann et al. 1955; c.f. also Roughgarden 2007, Sec. 18.3.1):

$$h_e(f_e) = \int_0^{f_e} c_e(t)\, dt \tag{5.7}$$

leading to the following optimization problem:

$$\min_{\boldsymbol{\lambda}} \quad \sum_{e \in E} h_e(f_e) \tag{5.8a}$$

subject to

$$\sum_{P \in \mathcal{P}_i} \lambda_P = r_i \qquad \forall i \tag{5.8b}$$

$$f_e = \boldsymbol{\Phi}_{(e)} \cdot \boldsymbol{\lambda} \qquad \forall e \in E \tag{5.8c}$$

$$\lambda_P \geq 0 \qquad \forall P \in \mathcal{P} \tag{5.8d}$$

5.8 Simultaneous Capacity and Flow Assignment

In this first analytical investigation of the problem, flows and capacities will be set simultaneously, in such a way that the target-utilization is always (and exactly) observed. In the context of the design of manufacturing systems under distributed PPC, this implies that the designer insists on maintaining a set target-utilization, as is custom in the design of manufacturing systems under hierarchical control.

5.8.1 Implementing the Target Utilization Assumption

The concept of utilization-attaining flow distributions (Definition 5.1) can be expressed by implementing a constraint that connects flow and capacity per machine:

$$a_e = \frac{1}{u_e^*} \cdot f_e \qquad \forall \left\{ e \in E \,\middle|\, f_e > 0 \right\} \tag{5.9}$$

Implementing these constraints in the latency functions of Section 5.5 and the respective social cost functions ($c_e(f_e) \cdot f_e$), provides new expressions for latency and social cost. In particular, implementing Eq. (5.9) allows us to express the capacity a_e as a function of the flow. The number of free variables is, therefore, reduced. The expressed assumption is that — during the system design stage — since flows are re-allocated to alternative paths and the capacity associated with this delta in flow “automatically shifts” with the workload, thereby maintaining the target utilization on the old path (where capacity is reduced on machines that are not part of the new path) and the new one (where capacity is appropriately increased where necessary). This “capacity follows flow” assumption

		Cost Function Class	
Cost Type	Expression	Affine	$M/M/1$
Latency	$c_e(f_e)$	$=u_e^* + b_e$ $=$constant	$=\frac{u_e^*}{1-u_e^*} \cdot f_e^{-1} \cdot \operatorname{sgn}(f_e)$ $\propto f_e^{-1}$
Social Cost	$c_e(f_e) \cdot f_e$	$= (u_e^* + b_e) \cdot f_e$ $\propto f_e$	$=\frac{u_e^*}{1-u_e^*} \cdot \operatorname{sgn}(f_e)$ $=$constant

Table 5.1: Overview of the functional form of the latency and social cost functions with the implemented utilization constraint. $\operatorname{sgn}(x)$ indicates the sign function, returning 1 for strictly positive values of x and 0 for $x = 0$ (note that $f_e \geq 0$ by definition).

then leads to simplified expressions for both the latency of a machine and the social cost accrued at one machine. They are expressed formally in Table 5.1 and visualized in Fig. 5.4.

With the utilization-attaining property "baked into" the problem formulation, one can now investigate the re-phrased problem for the existence of NEs and/or social optima. Of course, such swift change of capacity does not/cannot occur during the operation of the manufacturing system, but it is an approach to conceptualize the design stage. However, Nash- and socially optimal flow distributions are still valuable results, as we learn from Lemmata 5.1 and 5.2:

Lemma 5.1 (Equilibrium Persistency)**.** *A utilization attaining Nash-equilibrium obtained through the simultaneous assignment of flows and capacities will maintain its path and social cost, utilization levels, and NE property, under a relaxation of the utilization constraint, allowing variation of flows and keeping the capacities at the equilibrium levels.*

Proof. This is the immediate consequence of the experimental setup. We calculate flows and the corresponding capacities in such a way that, given the capacity values, the resulting Nash flow will be a utilization-attaining NE. □

Lemma 5.2 (Price of Predictability)**.** *A flow and capacity distribution that generates minimum social cost under the utilization constraint may not remain the social cost minimum under a variation of flows keeping the capacities at the equilibrium levels.*

Lemma 5.2 is a consequence of the additional constraints imposed through Eq. (5.9). As these additional constraints are dropped, the optimization problem may attain a better target function value.

In summary, any NE found during the following investigation of optimization problems that contain the fixed utilization constraint will remain a NE, if we set the capacities as calculated using the utilization constraint, and then let the commodities route their flows freely (without adjusting capacities accordingly) over these capacities. However, flow distributions that are *not* NE may exist in such capacitated networks that lead to lower

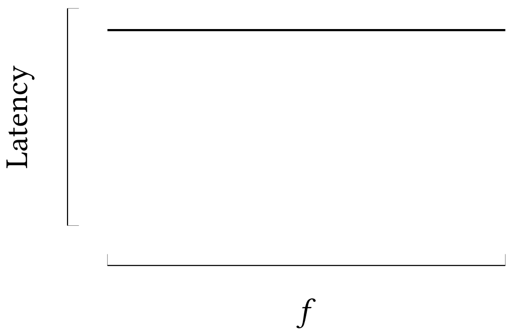

(a) Latency as function of flow for affine cost functions, assuming fixed utilization.

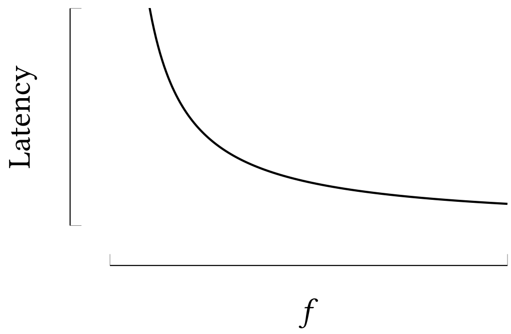

(b) Latency as function of flow for $M/M/1$ cost functions, assuming fixed utilization.

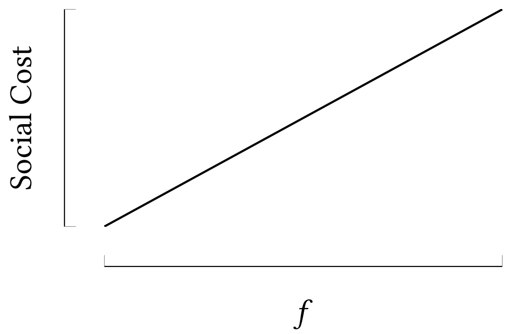

(c) Social cost as function of flow for affine cost functions, assuming fixed utilization.

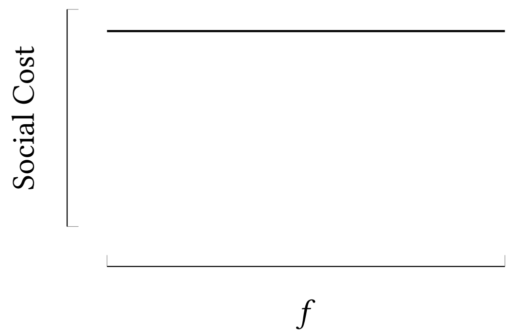

(d) Social Cost as function of flow for $M/M/1$ cost functions, assuming fixed utilization.

Figure 5.4: Visual approximation of the latency (first row) and social cost (second row) functions under the fixed utilization assumption on arbitrary scales. Results for affine cost functions are shown in the left column. The right column shows results for $M/M/1$ cost functions.

social cost, as calculated for the socially optimal solution of the optimization problem *with* utilization constraints.

5.8.2 Results for Affine Cost Functions

Existence & Identification of Nash-Equilibria

As shown in Table 5.1, the latency for affine cost functions is constant under the assumption of simultaneous flow and capacity assignment. This implies the following theorem:

Theorem 5.1 (Shortest paths yield a utilization-attaining Nash equilibrium)**.** *Given a set of source-sink pairs with associated flow demands, known target utilizations u_e^* and affine cost function parameters $u_e^* + b_e$ for all edges $e \in E$,* ***any*** *flow distribution among the shortest weighted paths for each commodity is a NE.*

Proof. To be a NE, the flow distribution has to exhibit two properties (Wardrop 1952):

1. the latency of all paths with flow has to be equal,
2. any path without flow has to have a higher latency than that of used paths.

The two properties obviously hold for the set of shortest weighted paths that must be of equal length (otherwise, one would be shorter) and shorter than any alternative path for that commodity (otherwise, it would be a shortest path). Since the latency term is constant, the path lengths do not change as we apply more flow. □

Note that this, in combination with Lemma 5.1, implies

Corollary 5.1 (Existence and equivalence of utilization-attaining Nash-equilibria for affine cost functions)**.** *There always exists at least one utilization-attaining Nash-equilibrium under affine cost functions. Should multiple path alternatives of equal length exist, any flow distribution among them induces the same latency and social cost.*

Socially-Optimal Flow-Distributions

The simple analytical expression for the latency also implies a simply tractable form for the social cost of affine cost functions under the simultaneous flow and capacity distribution assumption. In particular, it allows to state

Theorem 5.2. *The social cost for affine cost functions for any flow distribution over shortest weighted paths are the same and are the minimum value.*

Proof. The resulting social cost term is linear in the edge flow f_e. This means that the additional social cost, induced by adding flow on a certain path (a certain set of edges) is independent of the flow already associated with these edges and a linear function of solely the additional flow Δf_e. Hence, regardless of previous allocation decisions, any additional flow (e.g. flow of an additional source-sink pair) is best allocated to the path alternatives with the smallest sum of slopes $\sum_{P \in \mathcal{P}: e \in P} (u_e^* + b_e)$ which is just the a-priori known "weight" measure of the length of the path and hence equal across all shortest weighted paths; it is optimal in the sense that there is no alternative path for that commodity which would yield less additional social cost when given the same amount of additional flow. □

When considered together, Theorems 5.1 and 5.2 state that by using shortest paths, we can always construct a Nash-equilibrium as well as a socially optimal flow distribution. In fact, should multiple path alternatives of identical "length" exist, *any* flow distribution among these path alternatives is equivalent w.r.t. latency and social cost. Since both the Nash and the socially optimal flow distribution have the same social cost associated with it, the identified NE is the socially optimal utilization-attaining flow distribution. Note that this does *not* imply that the price of anarchy (Sections 3.1.3 and 5.3.4) is 1, since we cannot make statements about the existence of a better flow-distribution (in terms of social cost), *given* the capacity distribution calculated here (c.f. Lemma 5.2).

Notably, this result holds for *all* affine cost functions. A similar result (social cost of NE equals social optimum) is also known in the setting with fixed capacities and without target utilization constraints (Roughgarden 2005, Lemma 3.2.2). However, that result only holds for strictly linear latency functions (i.e. $b_e = 0 \; \forall e \in E$). The broader applicability of the above-mentioned conclusion is the result of the additional assumption on the ratio between flow and capacity.

5.8.3 Results for $M/M/1$ Cost Functions

Socially-Optimal Flow Distributions

For the $M/M/1$ cost function class, the latency expression is less appealing. However, for the social cost, Table 5.1 shows a constant expression in the $M/M/1$ column. This implies that every machine *that has non-zero flow* will contribute constant social cost independent of their actual flow. This immediately leads to

Theorem 5.3 (Minimal set of machines minimizes social cost)**.** *A machine (edge) is called* active *if the flow through the machine is nonzero. Then the lowest social cost for a utilization-attaining flow distribution in a network of $M/M/1$ cost functions is attained by flows that use the minimum weighted number of active machines (weighted by $\frac{u_e^*}{1-u_e^*}$).*

Proof. Follows immediately from the considerations made above. □

The implied optimization problem has inherently non-continuous (binary) decisions. It can hence be expressed as a MILP as follows. Let

- $\boldsymbol{B}$ be the *binary* path incidence matrix. With matrix elements $b_{e,P}$ indicating, whether or not edge e is in path P ($b_{i,j} = 1 \iff \phi_{i,j} > 0$, otherwise: 0). Column vectors are referred to as $\boldsymbol{B}_{(P)}$.
- x_e be the binary decision variable, indicating whether or not to choose machine e, composing together the decision vector $\mathbf{x}$.
- y_j be an auxiliary binary variable, indicating if all machines belonging to path j are part of the selected subset.
- M be a large scalar ("Big M Method").

$$\min_{x} \quad \sum_{e \in E} \frac{u_e^*}{1 - u_e^*} \cdot x_e \tag{5.10a}$$

subject to

$$\sum_{P \in \mathcal{P}_i} y_P \geq 1 \qquad \forall i, \tag{5.10b}$$

$$-(1 - y_P) \cdot M \leq \mathbf{x} \cdot B_{(P)} - \sum_{e \in E} b_{e,P} \qquad \forall P \in \mathcal{P}, \tag{5.10c}$$

$$x_e \in \{0, 1\} \qquad \forall e \in E, \tag{5.10d}$$

$$y_P \in \{0, 1\} \qquad \forall P \in \mathcal{P}. \tag{5.10e}$$

The optimization problem seeks to minimize the weighted set of "selected" machines. The selection is expressed in the binary decision variables $x_e \, \forall e$. The additional binary indicator variables y_P are tied to the selection of x_e's through Eq. (5.10c), which allows any given y_P to be equal to 1, *if and only if* the set of machines necessary for P is a subset of the set of selected machines. Equation (5.10b) ensures that at least one path alternative will be available for every commodity.

On the size of the MILP, the following can be noted:

Corollary 5.2. *Since the y_p's in Eq.* (5.10) *are tied to the values of x_e, the number of truly independent binary decision variables increases linearly with the number of edges (machines) in the system, which is generally much slower than the increase in path alternatives resulting from additional machines (which is often multiplicative).*

Existence & Identification of Nash-Equilibria

To investigate the existence and nature of utilization-attaining NEs under the assumption of simultaneous flow and capacity assignment, the potential function (Eq. (5.7)) for the $M/M/1$ latency case has to be further investigated. Using the expression from Table 5.1, the potential function can be expressed as:

$$h_e(f_e) = \int_0^{f_e} \frac{u^*}{1 - u_e^*} \cdot t^{-1} \, dt \tag{5.11}$$

where the expression in the integral, as the original $M/M/1$ latency function exhibits a singularity albeit this time at $t = 0$.

In order to still be able to approximate the problem, despite the singularity, the integral is re-normalized by introducing a small constant ϵ. Leading to

$$h_e(f_e) = \int_0^{f_e} \frac{u_e^*}{1-u_e^*} \cdot (t+\epsilon)^{-1}\, dt \tag{5.12a}$$

$$= \frac{u_e^*}{1-u_e^*} \cdot \ln\left(\frac{f_e+\epsilon}{\epsilon}\right) \tag{5.12b}$$

$$= \frac{u_e^*}{1-u_e^*} \cdot \ln(f_e+\epsilon) - \underbrace{\ln(\epsilon) \cdot \frac{u_e^*}{1-u_e^*}}_{\text{constant}} . \tag{5.12c}$$

Note that the constant subtrahend in Eq. (5.12c) can be ignored for optimization purposes.

The restated target function is concave, with the same set of linear constraints as in Eq. (5.8). For such concave minimization problems, a powerful optimization technique exists in the *Falk-Hoffmann Algorithm* (Falk and Hoffman 1986). However, even without such tools, one can draw conclusions about the nature of the optimal solution of this particular optimization problem (i.e. the Nash-flow) by observing that the decrease in latency (under the assumption that capacity is shifted with flow) increases, as the original flow level decreases. Assume a shift of a small amount of flow θ from one path alternative to another one of the same commodity. Then the latency of a machine no longer visited by this marginal flow decreases by

$$h_e(f_e+\theta) - h_e(f_e) \propto \ln(f_e+\epsilon+\theta) - \ln(f+\epsilon) \tag{5.13a}$$

$$= \ln\left(1 + \frac{\theta}{f_e+\epsilon}\right) \tag{5.13b}$$

which increases as f_e decreases. Moreover, note that the coefficient in Eq. (5.12c) is identical to the weights used in Eq. (5.10).

This allows us to draw the following important conclusion:

Theorem 5.4 (Nash equilibrium over minimal machine set)**.** *The utilization-attaining Nash flow distribution is given by the socially optimal, utilization-attaining flow distribution, i.e., by a flow allocation on path alternatives that minimize the weighted set of machines with a non-zero flow.*

Proof. Follows, as discussed above, from the nature of the target function in Eq. (5.12) and the social-cost properties of the minimal weighted set approach in Theorem 5.3. □

In particular, the set is identical to the solution of Eq. (5.10). Meaning that, as for the affine case, the social cost of the Nash-equilibrium and the socially-optimal flow distribution are equal. The found Nash-equilibrium is hence the socially optimal utilization-attaining flow distribution.[67]

In contrast to the affine case, however, not all flow distributions over the pre-determined set of machines (and hence paths) are equivalent in terms of latency and social cost: The concave nature of Eq. (5.12) means that, at least in the single-commodity case, the flow over the used machines is maximized (constrained by the overall flow constraint), as any deviation is penalized $\propto \frac{1}{f}$.

5.8.4 Merits of and Problems with the Simultaneous Allocation Approach

The mental model of assigning capacity and flow simultaneously that stood at the outset of Section 5.8, has significantly advanced our understanding of utilization-attaining flow distributions.

In particular, this approach was able to show the existence of utilization-attaining NEs that have the additional property of being the socially optimal flow distribution that attains the target utilization across all machines. This research has also presented construction mechanisms for both Nash- and socially optimal flow distributions under the "capacity follows flow" assumption that build upon well-established optimization techniques ranging from shortest-path problems to concave optimization.

The obtained results for both affine and $M/M/1$ cost functions are similar in the sense that they restrict the set of machines and/or path alternatives to the bare minimum, leading to manufacturing systems with low degrees of flexibility.

The following section will build upon the obtained results and discuss, if and to which extent, "more flexible" utilization-attaining flow distributions (i.e. flow distributions that use a higher number of path alternatives per commodity) are attainable by the system designer, leading to the notion of a trade-off between cost (in terms of investment and the cost of WIP discussed so far) on the one hand and flexibility on the other hand (Goerner et al. 2009; Kitano 2004, 2007; Meyer 2016, Ch. 6.2).

5.9 The "Price of Flexibility"

In this section, it is discussed if, given the analytical expressions of Table 5.1, one could construct alternative utilization-attaining Nash-equilibria from those derived in Section 5.8 with the general goal of maintaining more available path alternatives and hence higher levels of process flexibility in the equilibrium solution.

[67] As discussed in the case of affine cost functions, this does not rule out the existence of a socially preferable flow distribution that violates the utilization-constraint.

5.9.1 ...for Affine Cost Functions

The following corollary follows immediately from Theorem 5.1:

Corollary 5.3 (No utilization-attaining NE on non-shortest paths or paths of different lengths)**.** *There exists no utilization-attaining NE on networks with affine cost functions that gives strictly positive flow to path alternatives of different weighted lengths.*

Proof. Follows immediately from the proof of Theorem 5.1. The first condition of Wardrop's principle (equal latency of all paths with positive flow) cannot be met when paths of different weighted lengths are used. □

This means that other Nash-equilibria can only be attained, when the machine "cost" (the terms $u_e + b_e$) is adjusted in such a way that multiple path alternatives attain the minimal total weighted length and hence form a Nash-equilibrium under Theorem 5.1.

Since the social cost (in terms of WIP) of the affine cost function NEs was found to be only a function of the weighted "length" of the used shortest paths, any artificial adjustment in the edge lengths to incorporate NEs with more path alternatives is destined to change the resulting social cost, as the designer makes adjustments to the machine parameters b_e and u_e. This change constitutes a price of flexibility. Note that the cost for tied-up capital is, albeit the sole cost factor considered here so far, not the only cost component bound to change under such changes in machine target values: In particular, a downward adjustment of u_e increases the amount of extra capacity installed, leading to increased investment cost and a repositioning of the system in the trade-off discussed above. Similarly, investing in machines with lower b_e (lower workload independent throughput time) is also likely to incur higher investment cost in more powerful and faster machines.

5.9.2 ...for $M/M/1$ Cost Functions

The discussion of the price of flexibility for the $M/M/1$ cost function case starts with an encouraging observation:

Corollary 5.4 (Existence of utilization-attaining NEs for $M/M/1$ cost functions)**.** *Utilization-attaining flow distributions on a larger set of machines than that defined in Theorem 5.3 are possible.*

As the simplest motivating example, consider a network of n parallel edges and a single commodity. For all n machines (edges), we seek to attain the same target-utilization level u^*. Note that we can use *any* number of edges (not just 1, which would be the solution according to Theorem 5.3) and attain a utilization-attaining Nash-flow as long as we apply equal flow to all machines. Note that a configuration where m ($1 \leq m \leq n$) machines have a non-zero flow would induce m times the social cost associated with the socially optimal utilization-attaining flow distribution with just one machine of non-zero flow. A similar result was previously obtained by Korilis et al. (1997b).

The argument gets more difficult as we consider multi-commodity instances of less trivial structure.[68] In fact, our research does not provide any guarantee for the existence of or can characterize the properties of utilization-attaining flow distributions with a larger-than-necessary set of $M/M/1$ machines.

As with the affine cost function case, the key to more flexible utilization-attaining flow distributions seems to lie in the adjustment of the parameters. In the $M/M/1$ case, this is only the target-utilization level. In particular, the manufacturing system designer may decide to seek a Nash-equilibrium with increased process flexibility. Here, a problem where the system designer wants to set a particular flow distribution of choice as the Nash-flow is considered. In such case, the designer is willing to accept deviations from the originally envisioned target utilization rates ($u_e \neq u_e^*$) for achieving the desired flow distribution as the NE.

One can analyze this situation analytically: Let $\boldsymbol{\lambda}$ be the flow-distribution across path alternatives that the designer wants to enforce as the Nash-flow. In particular, this selection partitions $\mathcal{P}$ into two mutually exclusive sets $\mathcal{P}_{in}$ and $\mathcal{P}_{out}$ that take path alternatives with a non-zero flow and those without any flow respectively. Also, through $\boldsymbol{\Phi}$, the $\boldsymbol{\lambda}$ establishes an edge flow distribution $\boldsymbol{f}$. The edge latency then is a function of f_e as well as u_e which appears in the term $d(u_e) = \frac{u_e}{1-u_e}$. In particular, $c_e(f_e, u_e) = d(u_e) \cdot f_e^{-1}$ and hence a linear function of $d(u_e)$ (since $\boldsymbol{f}$ is known and fixed through the choice of $\boldsymbol{\lambda}$). We can then seek to minimize the maximum absolute deviation between $d(u_e^*)$ (d analyzed for the desired target utilization rate) and $d(u_e)$ (d analyzed for the attained utilization rate) across all edges. This difference is referred to as Δ_e. The problem to minimize $\max_{e \in E} |\Delta_e|$ in such a way that a given path-flow distribution $\boldsymbol{\lambda}$ is a utilization-attaining Nash-equilibrium, is expressed in the following linear problem:

$$
\begin{aligned}
\min_{\Delta_e, l_i} \quad & h && && (5.14a)\\
\text{subject to} \quad & && && \\
& h \geq \Delta_e && \forall e \in E && (5.14b)\\
& h \geq -\Delta_e && \forall e \in E && (5.14c)\\
& l_i - \sum_{e \in P} \underbrace{\Phi_{e,P} \cdot f_e^{-1}}_{\text{const.}} \cdot \Delta_e = \underbrace{\sum_{e \in P} \Phi_{e,P} \cdot \frac{d_e(u_e^*)}{f_e}}_{\text{const.}} && \forall i, P \in \mathcal{P}_{in} && (5.14d)\\
& \sum_{e \in P} \Phi_{e,P} \cdot \frac{d(u_e^*) + \Delta_e}{f_e} \geq l_i && \forall i, P \in \mathcal{P}_{out} && (5.14e)\\
\text{with} \quad & d(u) = \frac{u}{1-u} && && \\
& l_i = \text{Latency for commodity } i && \forall i &&
\end{aligned}
$$

[68] Albeit reminding that the literature on the design of networks for selfish users (c.f. Section 5.3.4) is frequently based on the assumption of single-commodity networks of parallel links (e.g. Korilis et al. 1995, 1997b).

Note that again both aspects of Wardrop's principles are considered. It is ensured that all paths with a non-zero flow have the same latency per commodity (Eq. (5.14d)) and that unused paths have at least the same latency (Eq. (5.14e)).

Minimizing the deviation w.r.t. $d(u_e)$, instead of u_e directly, is a valid simplification in situations where u_e^* is similar across all machines. Where larger deviations exist in target utilization values, this approach is likely to see large deviations from the desired utilization level for machines with small u_e^* (and hence small $d(u_e^*)$). In such settings, one may substitute Δ_e with a linear approximation of d around u_e^*. Note however that in cases where h is large, this linear approximation is not likely to sufficiently represent the true functional form of d_e.

5.10 A Brief Note on the "Price of Predictability"

The "price of predictability" (Lemma 5.2), i.e. the difference in social cost between a utilization-attaining flow distribution as compared to a flow-distribution "unfettered" by the target-utilization constraint, has not been considered so far. It has to be assumed that a better (in terms of social cost) flow distribution can be attained, where target-utilizations do not have to be met. But the questions is just how much better could it be?

Unfortunately, this research cannot give a clear answer in this regard. Notably, the existing upper bounds on the price of anarchy (c.f. Section 5.3.4) still apply, since the socially optimal utilization-attaining distributions of Section 5.8 were (in the first place) NEs and these persevere the drop of the utilization constraint (Lemma 5.1).

5.11 Discussion

The discussions in this chapter have addressed research question Q_3, as set out in Section 5.2. As such, this work adds to existing efforts to create a scientific understanding of emergent behavior. It can be understood as trying to entice selfish agents to produce aggregate emergent behavior that complies with perceived norms and ideals — here, in particular, the idea of setting fixed levels of utilization across machines to account for various factors — not considered in the fluid approximation used for strategic decision-making.

Where the idea of setting target utilizations is equated to traditional top-down manufacturing system design and control, the analysis performed in this chapter can also be understood as an attempt to maintain certain desirable properties associated with traditional system design methodologies in the context of distributed control. In particular, the discussion on the social cost and the price of anarchy associated with certain capacity distributions (here: capacity distributions derived from utilization-attaining flow distributions) give evidence how the design of systems — in particular the design of process path alternatives — can affect the performance and aggregate behavior of intelligent products. In this context, the results presented here also underline the findings presented in Section 3.3.1, namely that process plants with less flexibility (as the optimal solution to the optimization problems discussed in Section 5.8), are less likely to exhibit negative consequences of myopic behavior.[69]

Progress has been made in understanding the interplay between capacity dimensioning and the resulting product routing decisions in the same way that scientific advancements have been made in other related problems: Starting with an iterative heuristic (Section 5.6), analytical insights were developed (Section 5.8) and important properties (such as existence and social optimality) were shown for the hitherto uninvestigated class of utilization-attaining flow-distributions.

Implications for Practice

The low flexibility associated with the results of Section 5.8 and the subsequent further discussion in Section 5.9 can be attributed to the single-objective target function that makes products to take routing decisions purely based on throughput time. As such, approaches to dimension the system in accordance with agent behavior are bound to eventually "optimize" the flexibility out of the attained solution. The result is a stark reminder of the hitherto unexplored implications of a purely bottom-up constructionist system design approach. When thoroughly followed through, as this chapter has done, the mono-criterial decision function implemented in the fundamental building blocks of the systems, the product agents, and their decision-making, does in fact have ramifications "all the way" to system design, where they imply similarly lopsided network structures.

5.12 Directions for Future Research

It would hence be a worthwhile aim of future research to study, how a multi-criterial decision function would affect the outcome and, in more general, how decision functions of agents can be structured to "appreciate" the flexibility potentials of modern manufacturing systems.

It would also be interesting to investigate, with the toolset set out here, the effect of other forms of influencing agent behavior (as discussed in Chapter 3) on the existence and nature of utilization-attaining capacity distributions. It is very well possible that more flexible utilization-attaining flow distributions can be reached when agent decision making is influenced by tolls on decision alternatives (Gibbens and Kelly 1999) or stigmergy (Armbruster et al. 2006; Van Dyke Parunak 1997) (c.f. also Section 3.5.1).

Finally, the discussion in this research has considered only the WIP, the amount of inventory (=capital) tied in the system at every point in time, when calculating the social cost of systems. Other system characteristics for managerial decision-making have been ignored (c.f. Section 2.1.3). The aspect of investment cost seems tractable with the fluid model set out in this chapter. A discussion of investment cost within our model must consider two competing observations:

[69]Note that agent behavior was not changed here, so only the *impact* of myopic decision-making was altered.

1. For both cost function classes, the analytical discussions in this chapter have provided weighted selection problems, where machine weights increase with u^*. The flow is, therefore, drawn toward machines with low utilization, or, in other words, where extra capacity on top of estimated demand is high. This behavior is not surprising, given the latency-avoiding selfish nature assumed for intelligent products, but could indicate high investment cost when following this approach. Low-utilization machines would see large flows, leading to large investments in capacities, including "tactical over-capacities".

2. At the same time, we see a focus on on small set of machines to be actually implemented, which opens the possibility to reap economies of scale as only a few, "large" (in terms of capacity) machines are acquired.

A good starting point would be the general cost function proposed by Kleinrock (1970) for consideration in capacity-dimensioning methods: cost per machine = $d_e \cdot a_e^{b_e}$, where $d_e, b_e > 0$ (Chandy et al. 1977), where a_e is the machine capacity. With $b_e \leq 1$, in particular, economies of scale can be represented. Future research should seek to investigate the investment cost associated with manufacturing systems designed and dimensioned in a bottom-up way. Such analysis could discover a new aspect to the cost of anarchy by focusing on the additional up-front capital investment necessary, where the bottom-up system design is applied.

Chapter Six

Implications for Science and Practice

"Complexity 'thinking' is the art of maintaining the tension between pretending we know something, and knowing we know nothing for sure."

Richardson (2008, p. 21)

6.1 Summary of Results

This thesis has set out to explore the impact of distributed production control and the constructionist design approach associated with it on the design of manufacturing systems and their control. In particular, explored how *desirable* performance characteristics of hierarchical centralized production control may be maintained, when moving toward distributed control. To make progress in this direction, multiple sub-questions were identified in Section 1.4. In the following, the main findings of this research will be reviewed in their light.

> Q_1: *"Which design decisions concerning both controller and plant impact the duality between hierarchical and distributed control in PPC?"*

In Chapter 3, this thesis has developed a classification model that conceptualizes the quest to design manufacturing systems for distributed PPC as a problem of managing (*not* strictly minimizing) the degree and impact of myopic behavior in the system. The reviewed literature has shown that myopic behavior — while intrinsically linked to the very nature of distributed PPC — exhibits negative consequences that warrant a conscious and careful consideration during plant and controller design. Building upon results from literature in a multitude of disciplines, this research could derive a set of dimensions along which myopic decision-making (and its impact on manufacturing system performance) may be controlled.

By perceiving the design of manufacturing systems as an exercise in myopia management and the goal of balancing the advantages of both distributed and centralized PPC approaches as the problem of determining the "right" amount of myopia to tolerate in a manufacturing system, this thesis can establish a design space classification for distributed PPC systems, which allows the description and comparison of (semi-)heterarchical PPC systems in terms of their efforts into myopia reduction.

> Q_2*: "Which degree of hierarchy in the controller results in the highest logistics performance of a manufacturing system and why?"*

Chapter 4 considers how and to which extent hierarchy in the control network can improve the performance of networks of interacting agents. The chapter develops and analyzes in detail a minimal model of distributed coordination on networks. The chapter offers three important findings relating to the aspects of research question Q_2:

1. The results show a peak in performance for *medium degrees of hierarchy* (involvement of nodes with a "high level" view of the system), confirming so far purely hypothesized beliefs in and beyond the PPC community.
2. An explanation of this phenomenon is found in the analysis of mutual information between "leaders" and agents at lower levels of hierarchy, with additional understanding provided by the comparison with noise sources in simulated annealing processes.
3. Finally, the transferability towards PPC applications was motivated by the application of a forward model, which also confirmed the hypothesis about the relationship between time-scale of optimization and the optimal control architecture.

> Q_3*: "How can a production plant be designed to entice selfish agents to exhibit predictable and desirable emergent system properties?"*

The idea behind Chapter 5 may take a while to be digested at first: Why insist on the attainment of some arbitrarily set target utilization level? But at closer look, this chapter provides an arguably unprecedented glimpse in a world where manufacturing systems are, to the fullest extent possible, designed in a constructionist fashion *bottom-up*.

Chapter 5 makes the consequences of such approach strikingly clear: Where agents are set to maximize a simple single-objective target function, a manufacturing system designed in the constructionist fashion will show little sophistication either. For both affine and $M/M/1$ cost functions the socially optimal and the NE utilization-attaining flow distributions were characterized analytically by subjecting existing optimization problems for the calculation of these flows to the new constraint of a required target utilization. For both cost functions, it could be shown that the utilization-attaining NEs use a small subset of the machines and path alternatives originally provided, while flow distributions that make a broader use of flexibility cannot be reached at all (for affine cost functions) or only at the expense of investment in higher inventory levels (for $M/M/1$ cost functions).

Drawing upon the results from these three chapters, one can formulate a conclusion with respect to the research-guiding question:

> Q_0 (Research-guiding question)*: "How can both the plant and the controller of manufacturing systems be designed to achieve high logistics performance under distributed control?"*

The design of manufacturing systems for distributed control is inherently different from the traditional reductionist approach to manufacturing system design and control. The bottom-up emergent nature of system behavior and the intrinsic connection between distributed decision-making, along with the possibility of myopic behavior, set out new trade-offs to be navigated in system design that may be understood as combining centralized and distributed PPC traits. They, however, essentially describe a trade-off between the adaptiveness and flexibility that comes from quick heuristic-based decision-making based on limited information, and the stability, reliability, and (relative) optimality that traditional PPC methods can induce where conditions are appropriate.

The results of the analytical chapters deepen this perception: The results of the forward model (Section 4.6) provide evidence for the hypothesis that the optimal control architecture is dependent on the goals set out for the controlled system. Also Chapter 5 is essentially the tale of such a trade-off, albeit only visible under the additional constraint of fixed target-utilization: System layouts that exhibit flexibility, will only arise where agents internalize the trade-off (e.g. between performance and flexibility) and are willing to accept lower performance (higher throughput times) in exchange for less fragile production networks.

6.2 Implications for Theory

This thesis has picked up a number of interrelated research streams to derive a holistic understanding of distributed and hybrid PPC systems and to identify ideas relevant to their design and analysis (c.f. Sections 1.5 and 3.2.2). This section reviews the results obtained in this thesis w.r.t. their implications for the underlying research streams.

6.2.1 Design of Hybrid PPC Systems

The design of hybrid PPC systems — systems that combine desirable features of hierarchical and strictly distributed architectures – has been the driving idea behind this thesis (Section 1.2). The main contribution to theory in the design of hybrid PPC systems lies in the additional structure and analytical insights, given in this thesis.

Chapter 3 places the aim to design hybrid PPC systems in a multi-disciplinary effort to understand and shape emergent behavior. While previous attempts have been made to identify and structure measures to eliminate non-performing and/or erratic behavior of distributed PPC systems, the results of this thesis provide substantial evidence that a wider horizon of disciplines should be considered.

The contribution to PPC theory of Chapter 4 lies in the analytical investigation of hierarchy as one measure to reduce the impact of myopic behavior in distributed PPC systems. Not only do the results add to the very small set of literature so far that can claim to have investigated the design spectrum between centralized and distributed control (at least with respect to hierarchy in the control network) in its entirety (c.f. Section 2.3.2), but a mechanistic understanding is provided as to *why* this effect occurs. Designers and researchers of distributed PPC are, therefore, given not only renewed confidence in the hypothesis about an "optimal" degree of hierarchy, but are also provided with ideas as to how the effectiveness of hierarchy can be measured. Together with the discussion and stylized experiments in Chapter 3, this thesis strongly underscores the idea that the optimal performance is most likely to be found in the realm of *hybrid* PPC.

The analytical and experimental insights generated in Chapters 4 and 5 should also further encourage the use of minimal models in the domain of manufacturing system design and analysis (c.f. Section 1.6.2). While questions of transferability could not be fully explored within the scope of this thesis (c.f. recommendations for future research provided in Sections 4.8 and 5.12), the kind of mechanistic understanding that could be developed in the analytical chapters of this work would have likely been beyond reach in more detailed, complex system models.

6.2.2 Complex Leadership Theory

CLT was primarily used in Chapter 4 to develop hypotheses for the mechanisms of leadership that could be tested in the GCD model. The experimental validation of these hypotheses (c.f. Section 4.7) indicates that, while Uhl-Bien et al. (2007) are right to insist that "leadership" cannot simply be equated with "leaders" in complex systems, "leaders" still play a central role in the emergence of leadership, not because of any assigned authority, but owing to their *positions*, central in the organization. They lead (as suggested in Hypotheses H_2 and H_3) by having access to information and by invoking the breakup of local regimes (patterns of behavior). The confirmation of these hypotheses constitutes the first contribution of this thesis to CLT.

The curvilinear performance shape, observed in Section 4.4, also allows to develop *new* hypotheses for leadership mechanisms in CAS that establish another contribution to CLT research: The results suggest that the capabilities of leader nodes to foster the coordination process in the system are constrained by limitations on communication. The results in this thesis have highlighted the out sized role of intra-leader communication avenues to maintain the leader's role as sources of information and orchestrators of the coordination process.

Both contributions also underline (again) the value that analytical investigations of ABMs can bring to the understanding of organizations and the development of organization theory (Hazy 2007; Hazy et al. 2007).

The results should finally be seen as encouragement for organization science researchers to become more deeply involved in neighboring scientific debates. The necessity to design (requiring the ability to analyze) organization schemes based fully on computerized agents has already been pointed out. At the same point, the interaction of such systems with human operators, with their flexibility and information-processing capacity is increasingly recognized in the production control literature as a key challenge for the successful development and adoption of novel forms of production control in practice (Delfmann et al. 2017; Gorecky et al. 2014). A closer cooperation between researchers in the domains of distributed PPC and CLT is not only encouraged by similar ideas and problems, but also mandated by the requirements of future production environments.

6.2.3 (Algorithmic) Game Theory

(Algorithmic) game theory in general and congestion games in particular have proven to be a suitable tool for the analysis of agent behavior and interaction in the context of distributed PPC in Chapter 5. As the literature review has indicated, “logistics-related” applications of congestion games have so far been mainly centered on traffic modeling. In that, this thesis has successfully extended the modeling approach to another (albeit: related) domain of application.

The idea to enforce target-utilization levels is in itself a new facet in the discussion, how capacitated flow networks can be shaped to designer expectations. So far, avoiding a price of anarchy or finding an optimum in the trade-off between flow time and investment has been assumed (c.f. Section 5.3.4).

6.2.4 Complex Adaptive Systems Theory & Constructionist System Design

The fundamental difference between reductionist and constructionist design methodologies, along with the absence of cause-and-effect models to guide constructionist system design (called “emergence engineering” by Trentesaux (2009)), was identified as the research gap in Section 1.3.

This thesis has addressed this lack of understanding through a number of different ways: by means of literature review and classification (c.f. Chapter 3) as well as minimal model experimentation (Chapters 4 and 5). The results substantiate the idea that controlling and designing CAS to attain high performance, may be supported by intentionally “mixing in” traits of traditional hierarchical control approaches. While originally developed for the domain of PPC, the classification model of Chapter 3 seems open to extensions/transfer to a number of engineered systems where system behavior is shaped by a series of hierarchically interconnected design decisions.

The results of Chapter 4 not only further the understanding of hierarchy as one particular design dimension in the context of hybrid PPC, but also provide an abstract measure (mutual information) to assess the impact of “leader nodes”. The results underscore the existing sentiment (Lizier et al. 2008) that information transfer, carefully adopted to the coordination problem at hand, can support the investigation of coordination processes in more general. Finding hypotheses developed from CLT to apply to purely engineered systems, also supports the idea (first expressed by Solow and Szmerekovsky (2006)) that the role of leaders may constitute a more general system design question where organization theory can contribute to general MAS theory and understanding.

Beyond collecting and creating evidence for the preferability of hybrid PPC approaches in earlier chapters, Chapter 5 has advanced the discussion on the larger-scale ramifications of a bottom-up, constructionist system design approach. While (at least in the domain of PPC), distributed control has so far been implicitly assumed to operate within the framework conditions of a system plant designed in a constructionist top-down fashion, this thesis has followed through with a bottom-up design approach “all the way” to strategic decisions of systems design, showing clear implications and correlations between decisions made at the agent level (i.e.: preferences for short throughput time) and observed characteristics at the system level (material-flow networks of little flexibility).

6.3 Implications for Practice

As stated before (Sections 1.3 and 2.2.2), the lack of understanding of the mechanics and uncertainty about the performance of distributed PPC systems — among other things — constitute one significant obstacle in their industrial application. The research-guiding question of this thesis was formulated in an effort to reduce this roadblock to the adoption of agent-based and hybrid PPC approaches. This section will review the implications of this thesis for practitioners contemplating the adoption of agent-based PPC approaches. It will also review the findings in the light of an anticipated "fourth industrial revolution" (c.f. Section 1.1.2) and discuss the possible implications for other domains of application.

6.3.1 Improving the Acceptance of Distributed PPC in Practice

The classification model of Chapter 3 aimed at understanding and describing the design space *between* the poles of hierarchical and strictly distributed PPC. It provides manufacturing system designers with a better understanding of how PPC approaches can *deliberately* be positioned between the poles (and characteristics) of hierarchical and fully distributed PPC and how the system plant influences this trade-off. The results should increase the confidence of manufacturing system designers to shape distributed or hybrid PPC approaches according to their wishes and needs.

While Chapter 3 has implicitly *assumed* that hybrid PPC approaches exhibit the best performance, Chapter 4 has provided experimental evidence. More importantly for applications in industry, it has motivated a concrete measure (mutual information) and a comparison with an established optimization procedure (simulation annealing) that should enhance the understanding and ability to measure and influence the degree of hierarchy in a control system to attain optimal performance.

The ability to predict the emergent behavior of autonomous and selfish products and to form this collective behavior to meet one's own expectations stood at the center of Chapter 5. Beyond their direct applicability to capacity-dimensioning problems, the results, along with the conclusions from Section 3.3.1, support the idea that simple system plants can effectively reduce the impact of myopic behavior. This thesis then helps to delineate, very much in accordance with similar boundaries of applicability defined for PULL production, a range of environmental conditions within which the application of agent-based distributed PPC seems to be associated with little risk.

6.3.2 Implications on the Way to the Fourth Industrial Revolution

If one believes the predictions, the manufacturing systems of the future are likely to find an environment very different from what today's systems were designed for: Highly complex and flexible production plants, low volumes per product, and frequent change in the product mix have been mentioned in Section 1.1.1. The manufacturing systems resulting from the fourth industrial revolution will be expected to strive in these environments (Section 1.1.2). This thesis then has implications for ongoing efforts to promote the digitization of production, at least to the extent that "intelligent" objects and distributed control are part of the vision.

Generally, a production environment that exhibits higher degrees of system plant complexity is — by the logic of Chapter 3 — likely to be in need of more significant interventions in agent behavior in order to avoid experiencing performance and predictability losses due to myopic behavior. A sensible approach to the fourth industrial revolution, which accounts for the substantial financial effects of a breakdown or loss of performance in industrial production facilities, should start with distributed PPC sytems that show rather *little* myopic behavior (e.g. are constrained by significant offline scheduling and with close to global information) and seek optimization potential by allowing more myopic behavior, as frequent changes in production and/or unplanned interruptions call for increased responsiveness. An iterative process that starts at the opposite end (with a high degree of myopic behavior) is more likely to accrue substantial losses in performance, as the designers try to find an optimal configuration in the (still wide) range of hybrid PPC designs.

The results of Chapter 4 not only underline the need to find a balance between distributed and hierarchical PPC, but they can also help to answer one of the hitherto unanswered design questions in the realm of CPS in production: the structure of communication and interaction between agents (c.f. Section 1.3). The results imply that hierarchies should be developed and integrated into the underlying agent network to the extent that the established leaders can support the coordination process (as measurable by mutual information). Connecting agents on a higher level of hierarchy may substantially improve this capability and justify higher degrees of hierarchy.

The findings of Chapters 3 and 5 are also relevant where intelligent products designed by different companies interact within one system plant. The vulnerability, more altruist agents may have to selfish behaviors exhibited by others (c.f. Section 2.1.2) and the profound impact that agents with a simple, one-dimensional target function can have on system design (Chapter 5), both imply that many features, designers could so far build into their manufacturing systems (like flexibility, safety stocks, …) are much harder to achieve, where intelligent products roam through the system. It will take either contractual agreements (and potentially: payments) between supply chain partners, more elaborate decision-making functions, or some form of incentives to make selfish agents e.g. forming a safety stock or populating less attractive production paths. While exact design recommendations for these situations cannot be derived directly from this thesis, the results imply that actions in this direction will be necessary.

6.3.3 Implications Beyond Industrial Production

While developed in the context of manufacturing system design and control, many of the results presented in this thesis have implications for other domains.

The design dimensions developed in Chapter 3 are naturally bound to the domain of this thesis. However, the general ideas of (1) picking up design decisions and (2) evaluating them w.r.t. their influence on the degree of myopic decision-making or the impact thereof on system performance should be transferable to other engineering disciplines where digitization heralds more distributed design approaches.

The interdisciplinary nature of the discussion in Chapter 4 gives evidence of the broad impact and applicability of these results across a wide range of domains: Distributed coordination processes are a common theme in many natural, social, and engineered systems. Understanding the role of hierarchy in these systems is valuable not only to designers of engineered systems, but also to those of social systems (the implications on organization theory have already been discussed) and to the study of biological systems (c.f. Ravasz et al. 2002; Simon 1962, and the discussion on hierarchy in Section 3.3.2).

The model language of 5 was "borrowed" from other disciplines and adapted to a problem motivated in the context of manufacturing system design. However, the motivation to maintain some buffer when making capacity-dimensioning decisions and the inherent contradiction between top-down design decisions and bottom-up emergence of system behavior seem applicable also to those domains that have so far driven the discussion on congestion games (notably: the design of traffic networks and distributed computation).

6.4 Limitations of Minimal Model Experiments

Throughout this thesis (starting with Section 1.6.2), the minimal nature of the models discussed in detail and many of the inputs considered in the development of the classification model, have been stressed. It is then an obvious limitation of these results that they were often not (fully) observed from real manufacturing systems or more highly detailed models (e.g. simulation studies). The models are not only minimal in their semantical capacity to represent the many entities commonly found in manufacturing systems, but also generally assume very simplistic decision rules and system structures. Many of these limitations have already been addressed in the respective chapters. Here some additional and more high-level assumptions will be discussed that may limit the transferability of the results. It should be noted, however, that the research conducted in the analytical chapters of this dissertation should be viewed as *basic or fundamental research*, with the aim to improve the general knowledge and understanding of a phenomenon (National Science Foundation 1953) — not necessarily practical evidence.

For the GCD experiments reported in Chapter 4 and the scheduling application of Section 4.6, the limitations as a scheduling tool are mainly the assumed unit-size processing time across all operations (a commonly made assumption in CA-based scheduling, c.f. Carneiro and Oliveira 2013) and the lack of any technical and environmental constraints (c.f. Bongaerts et al. 2000) that provide additional structure, e.g. by creating jobs out of individual operations. General flexibility and identical capacity of all machines (colors) across all operations were further assumed. While the representativeness of test-instances used to demonstrate scheduling algorithms is generally questionable (Hall and Posner 2001; Hernando et al. 2016; McGeoch 1996), the use of "shortcuts" and the use of a scheduling problem of known structure (a ring graph) to circumvent the $\mathcal{NP}$-hard coloring/scheduling problem is particularly prone to be of little representative value (McGeoch 1996).

In Chapter 5, this research has limited itself to the discussion of two different latency functions; it has generally been assumed that all machines adhere to the same latency function class. It has also been assumed that the processing capacity requirements calculated are generally fulfilled by a single machine (we used e.g. an $M/M/1$, not an $M/M/n$ modeling approach). The assumption of an a priori known and constant demand distribution between products may also clash with the assumptions made about future manufacturing systems (Section 1.1). Here such an assumption is found necessary to provide some planning basis for the allocation of machine capacity.

6.5 Recommendations for Further Research

Model- and problem-specific future research directions have already been discussed in the concluding sections of the main thesis chapters. In particular, it is in the nature of minimal model experiments that the question of result-transferability to more realistic settings cannot be fully answered at present. Such transfer then would be one obvious avenue for further research.

The discussions in Chapter 3 and the conclusions that could be drawn from the experiments in Chapters 4 and 5, however, also justify – in the eye of the author – further such minimal model experiments to investigate individual dimensions of myopia avoidance in the design of manufacturing systems under distributed control. Following the example of Chapter 4, a viable approach would seem to investigate other individual dimensions of myopia control, identified in Chapter 3, through minimal models. Moreover, the analytical experiments in this thesis have focused on single objectives (minimization of makespans and WIP), while similar experiments seem possible that focus on other target dimensions.

Finally, the interplay between multiple dimensions of myopia control within one manufacturing system remains a challenge that has so far withstood attempts to gain insights beyond individual simulation experiments. However, even with this limitation, the classification model of Chapter 3 suggests (and supports) the conceptualization of distributed PPC approaches as forming a design space that may be subjected to more rigorous exploration than has so far been reported. Future research should seek to compare different control approaches that vary across multiple dimensions of myopia control in one system plant to improve our understanding of the interrelationship between myopia control dimensions on the one hand and the interrelationship between manufacturing plant and controller design on the other hand.

Bibliography

Abbott, Russ (2006). "Emergence explained: Abstractions: Getting epiphenomena to do real work." In: *Complexity* 12.1, pp. 13–26. issn: 1099-0526. doi: 10.1002/cplx.20146.

Ackoff, Russell L. (1970). "A concept of corporate planning." In: *Long Range Planning* 3.1, pp. 2 –8. issn: 0024-6301. doi: 10.1016/0024-6301(70)90031-2.

– (1973). "Science in the Systems Age: Beyond IE, OR, and MS." In: *Operations Research* 21.3, pp. 661–671. doi: 10.1287/opre.21.3.661. jstor: 169376.

– (1974). "The systems revolution." In: *Long Range Planning* 7.6, pp. 2 –20. issn: 0024-6301. doi: 10.1016/0024-6301(74)90127-7.

Agre, Philip E. and David Chapman (1990). "What are plans for?" In: *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. Ed. by Pettie Maes. Cambridge, MA: MIT Press, pp. 17 –34. isbn: 0-262-63135-0.

Ahrens, Volker (1996). "New organization concepts for distributed production systems." In: *Kybernetes* 25.2 (Mar. 1996), pp. 24–39. doi: 10.1108/03684929610114628.

– (1998). "Systems Models and Concepts for Disributed Production Planning and Control." In: *Shop Floor Control — A Systems Perspective: From Deterministic Models towards Agile Operations Management*. Ed. by Eric Scherer. Berlin, Heidelberg: Springer, pp. 173–198. isbn: 978-3-642-60313-6. doi: 10.1007/978-3-642-60313-6_7.

Aland, Sebastian, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann (2006). "Exact Price of Anarchy for Polynomial Congestion Games." In: *STACS 2006: 23rd Annual Symposium on Theoretical Aspects of Computer Science, Marseille, France, February 23-25, 2006. Proceedings*. Ed. by Bruno Durand and Wolfgang Thomas. Berlin, Heidelberg: Springer, pp. 218–229. isbn: 978-3-540-32288-7. doi: 10.1007/11672142_17.

Allen, Arnold O. (2014). *Probability, Statistics, and Queueing Theory — with Computer Science Applications*. 2nd Edition. Computer Science and Scientific Computing. Boston, MA: Academic Press Inc. isbn: 0-12-051051-0.

Aloulou, Mohamed Ali and Marie-Claude Portmann (2005). "An Efficient Proactive-Reactive Scheduling Approach to Hedge Against Shop Floor Disturbances." English. In: *Multidisciplinary Scheduling: Theory and Applications*. Ed. by Graham Kendall, EdmundK. Burke, Sanja Petrovic, and Michel Gendreau. Springer, pp. 223–246. isbn: 978-0-387-25266-7. doi: 10.1007/0-387-27744-7_11.

Alshabi, W., S. Ramaswamy, M. Itmi, and H. Abdulrab (2007). "Coordination, Cooperation and Conflict Resolution in Multi-Agent Systems." English. In: *Innovations and Advanced Techniques in Computer and Information Sciences and Engineering*. Ed. by Tarek Sobh. Springer Netherlands, pp. 495–500. isbn: 978-1-4020-6267-4. doi: 10.1007/978-1-4020-6268-1_87.

Altman, E., T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter (2006). "A survey on networking games in telecommunications." In: *Computers & Operations Research* 33.2. Game Theory: Numerical Methods and Applications, pp. 286 –311. issn: 0305-0548. doi: 10.1016/j.cor.2004.06.005.

Anderson, Philip (1999). "Perspective: Complexity Theory and Organization Science." In: *Organization Science* 10.3, pp. 216–232. doi: 10.1287/orsc.10.3.216. jstor: 2640328.

Argoneto, Pierluigi, Giovanni Perrone, Paolo Renna, Giovanna Lo Nigro, Manfredi Bruccoleri, and Sergio Noto La Diega (2008). *Production Planning in Production Networks — Models for Medium and Short-term Planning*. London: Springer Science + Business Media. isbn: 978-1-84800-057-5. doi: 10.1007/978-1-84800-058-2.

Arlbjørn, Jan Stenthoft and Per Vagn Freytag (2013). "Evidence of lean: a review of international peer-reviewed journal articles." In: *European Business Review* 25.2, pp. 174–205. doi: 10.1108/09555341311302675.

Armbruster, D., C. de Beer, M. Freitag, T. Jagalski, and C. Ringhofer (2006). "Autonomous control of production networks using a pheromone approach." In: *Physica A: Statistical Mechanics and its Applications* 363.1. Information and Material Flows in Complex Networks, pp. 104 –114. issn: 0378-4371. doi: 10.1016/j.physa.2006.01.052.

Arnold, Dieter and Kai Furmans (2007). *Materialfluss in Logistiksystemen.* 5th. VDI-Buch. Berlin, Heidelberg: Springer. isbn: 978-3-540-45663-6. doi: 10.1007/978-3-540-45663-6.

Arshad, Muhammad and Marius C Silaghi (2004). "Distributed simulated annealing." In: *Distributed Constraint Problem Solving and Reasoning in Multi-Agent Systems.* Ed. by W. Zhang and V. Sorge. Vol. 112. Frontiers in Artificial Intelligence and Applications. IOS Press. isbn: 978-1-58603-456-6.

Artigues, Christian, Jean-Charles Billaut, and Carl Esswein (2005). "Maximization of solution flexibility for robust shop scheduling." In: *European Journal of Operational Research* 165.2. Project Management and Scheduling, pp. 314 –328. issn: 0377-2217. doi: 10.1016/j.ejor.2004.04.004.

Ashby, William (1961). *An Introduction to Cybernetics.* 4th. London: Chapman & Hall Ltd.

Askin, Ronald G. (2013). "Contributions to the design and analysis of cellular manufacturing systems." In: *International Journal of Production Research* 51.23-24, pp. 6778–6787. doi: 10.1080/00207543.2013.825745.

Askin, Ronald G. and Jeffrey B. Goldberg (2002). *Design and Analysis of Lean Production Systems.* John Wiley & Sons Inc. isbn: 978-0-471-11593-9.

Atzori, Luigi, Antonio Iera, and Giacomo Morabito (2010). "The Internet of Things: A survey." In: *Computer Networks* 54.15, pp. 2787 –2805. issn: 1389-1286. doi: 10.1016/j.comnet.2010.05.010.

Axelrod, Robert (2007). "Simulation in the Social Sciences." In: *Handbook of Research on Nature Inspired Computing for Economics and Management.* Ed. by Jean-Philippe Rennard. Hersey, PA: Idea Group. Chap. 7. isbn: 1591409845. doi: 10.4018/978-1-59140-984-7.ch007.

Babiceanu, Radu F. and F. Frank Chen (2006). "Development and Applications of Holonic Manufacturing Systems: A Survey." In: *Journal of Intelligent Manufacturing* 17.1, pp. 111–131. issn: 1572-8145. doi: 10.1007/s10845-005-5516-y.

Başar, Tamer and Geert Jan Olsder (1998). *Dynamic noncooperative game theory.* 2nd Edition. Vol. 23. Classics in Applied Mathematics. New York, NY: Siam. isbn: 0-89871-429-X.

Bahl, Harish C., Larry P. Ritzman, and Jatinder N. D. Gupta (1987). "Determining Lot Sizes and Resource Requirements: A Review." In: *Operations Research* 35.3, pp. 329–345. doi: 10.1287/opre.35.3.329. jstor: 170535.

Bajracharya, Max, Andrew Howard, Larry H. Matthies, Benyang Tang, and Michael Turmon (2009). "Autonomous off-road navigation with end-to-end learning for the LAGR program." In: *Journal of Field Robotics* 26.1, pp. 3–25. issn: 1556-4967. doi: 10.1002/rob.20269.

Bak, Per, Chao Tang, and Kurt Wiesenfeld (1987). "Self-organized criticality: An explanation of the $1/f$ noise." In: *Phys. Rev. Lett.* 59.4, pp. 381–384. doi: 10.1103/PhysRevLett.59.381.

Baker, Albert D (1998). "A survey of factory control algorithms that can be implemented in a multi-agent heterarchy: Dispatching, scheduling, and pull." In: *Journal of Manufacturing Systems* 17.4, pp. 297 –320. issn: 0278-6125. doi: 10.1016/S0278-6125(98)80077-0.

Bakule, Lubomír (2008). "Decentralized control: An overview." In: *Annual Reviews in Control* 32.1, pp. 87 –98. issn: 1367-5788. doi: 10.1016/j.arcontrol.2008.03.004.

Ballot, Eric, Olivier Gobet, and Benoit Montreuil (2012). "Physical Internet Enabled Open Hub Network Design for Distributed Networked Operations." In: *Service Orientation in Holonic and Multi-Agent Manufacturing Control*. Ed. by Theodor Borangiu, André Thomas, and Damien Trentesaux. Berlin, Heidelberg: Springer, pp. 279–292. isbn: 978-3-642-27449-7. doi: 10.1007/978-3-642-27449-7_21.

Bar-Yam, Yaneer (2002). *General features of complex systems*. In: *Encyclopedia of Life Support Systems (EOLSS), Developed under the Auspices of the UNESCO*. Vol. 1. EOLSS Publishers.

Barabási, Albert-László (2012). "The network takeover." In: *Nature Physics* 8.1 (Jan. 2012), pp. 14–16. doi: 10.1038/nphys2188.

Barabási, Albert-László and Réka Albert (1999). "Emergence of Scaling in Random Networks." In: *Science* 286.5439, pp. 509–512. issn: 0036-8075. doi: 10.1126/science.286.5439.509. pubmed: 10521342.

Barber, K. Suzanne, David C. Han, and Tse-Hsin Liu (2001). "Strategy Selection-Based Meta-level Reasoning for Multi-agent Problem-Solving." In: *Agent-Oriented Software Engineering: First International Workshop, AOSE 2000 Limerick, Ireland, June 10, 2000 Revised Papers*. Ed. by Paolo Ciancarini and Michael J. Wooldridge. Berlin, Heidelberg: Springer, pp. 269–283. isbn: 978-3-540-44564-7. doi: 10.1007/3-540-44564-1_18.

Barbosa, José, Paulo Leitão, Emmanuel Adam, and Damien Trentesaux (2015). "Dynamic self-organization in holonic multi-agent manufacturing systems: The ADACOR evolution." In: *Computers in Industry* 66, pp. 99 –111. issn: 0166-3615. doi: 10.1016/j.compind.2014.10.011.

Bard, Jonathan F. and Thomas A. Feo (1991). "An Algorithm for the Manufacturing Equipment Selection Problem." In: *IIE Transactions* 23.1, pp. 83–92. doi: 10.1080/07408179108963843.

Baron, A., R. Ginosar, and I. Keslassy (2009). "The Capacity Allocation Paradox." In: *IEEE INFOCOM 2009*, pp. 1359–1367. doi: 10.1109/INFCOM.2009.5062051.

Bartholdi, John J., Thomas D. Seeley, Craig A. Tovey, and John H.Vande Vate (1993). "The Pattern and Effectiveness of Forager Allocation Among Flower Patches by Honey Bee Colonies." In: *Journal of Theoretical Biology* 160.1, pp. 23 –40. issn: 0022-5193. doi: 10.1006/jtbi.1993.1002.

Bartlett, Christopher A. and Sumantra Ghoshal (1987). "Managing Across Borders: New Organizational Responses." In: *Sloan Management Review* 29.1 (Fall 1987), pp. 43–53.

Bass, Len, Paul Clements, and Rick Kazman (2003). *Software Architecture in Practice*. 2nd. Addison-Wesley Professional. isbn: 0-321-15495-9.

Batterman, Robert W. (2002). "Asymptotics and the Role of Minimal Models." In: *The British Journal for the Philosophy of Science* 53.1, pp. 21–38. doi: 10.1093/bjps/53.1.21.

Batterman, Robert W. and Collin C. Rice (2014). "Minimal Model Explanations." In: *Philosophy of Science* 81.3, pp. 349–376. issn: 0031-8248. doi: 10.1086/676677.

Bauer, Wilhelm, Sebastian Schlund, Dirk Marrenbach, and Oliver Ganschar (2014). *Industrie 4.0 — Volkswirtschaftliches Potenzial für Deutschland*. Studie. BITKOM Bundesverband Informationswirtschaft, Telekommunikation und neue Medien e.V. and Fraunhofer Institut für Arbeitswirtschaft und Organisation (IAO).

Becker, Raphen, Alan Carlin, Victor Lesser, and Shlomo Zilberstein (2009). "Analyzing Myopic Approaches For Multi-Agent Communication." In: *Computational Intelligence* 25.1, pp. 31 –50. issn: 0824-7935. doi: 10.1111/j.1467-8640.2008.01329.x.

Becker, Till (2012). "Parallels between Manufacturing Systems in Industrial Production and Metabolic Systems in Biological Cells a Network Based Unifying Model from a Production

Logistics Perspective." PhD thesis. Jacobs University Bremen. URN: urn:nbn:de:101:1-201305294988.

Becker, Till (2016). "Application of Topological Network Measures in Manufacturing Systems." In: *Dynamics in Logistics: Proceedings of the 4^{th} International Conference LDIC, Bremen, Germany*. Ed. by Herbert Kotzab, Jürgen Pannek, and Klaus-Dieter Thoben. Cham: Springer International Publishing, pp. 35–45. ISBN: 978-3-319-23512-7. DOI: 10.1007/978-3-319-23512-7_4.

Becker, Till, Moritz E. Beber, Katja Windt, Marc-Thorsten Hütt, and Dirk Helbing (2011). "Flow control by periodic devices: a unifying language for the description of traffic, production, and metabolic systems." In: *Journal of Statistical Mechanics: Theory and Experiment* 2011.05, P05004. DOI: 10.1088/1742-5468/2011/05/P05004.

Becker, Till, Mirja Meyer, and Katja Windt (2014). "A manufacturing systems network model for the evaluation of complex manufacturing systems." In: *International Journal of Productivity and Performance Management* 63.3, pp. 324–340. DOI: 10.1108/IJPPM-03-2013-0047.

Becker, Till, Christoph Illigen, Bill McKelvey, Michael Hülsmann, and Katja Windt (2016). "Using an agent-based neural-network computational model to improve product routing in a logistics facility." In: *International Journal of Production Economics* 174, pp. 156 –167. ISSN: 0925-5273. DOI: 10.1016/j.ijpe.2016.01.003.

Beckmann, Martin, C. B. McGuire, and Christopher B. Winsten (1955). *Studies in the Economics of Transportation*. Tech. rep. RM-1488. RAND Corporation.

Benartzi, Shlomo and Richard H. Thaler (1993). *Myopic Loss Aversion and the Equity Premium Puzzle*. Working Paper 4369. National Bureau of Economic Research. DOI: 10.3386/w4369.

Bendul, Julia C. and Henning Blunck (2018). "The Design Space of Production Planning and Control for Industry 4.0." In: *Computers in Industry*. Under review.

Bennett, David and Rob Dekkers (2005). "Industrial networks of the future: a critical commentary on research and practice." In: *Papers of the 12th International EurOMA Conference on Operational and Global Competitiveness*. Ed. by Krisztina Demeter, pp. 677–686. ISBN: 9632184556.

Berger, Thierry, Yves Sallez, Bruno Valli, Alain Gibaud, and Damien Trentesaux (2010). "Semi-heterarchical Allocation and Routing Processes in FMS Control: A Stigmergic Approach." In: *Journal of Intelligent and Robotic Systems* 58.1, pp. 17–45. ISSN: 1573-0409. DOI: 10.1007/s10846-009-9343-9.

von Bertalanffy, Ludwig (1972). "The History and Status of General Systems Theory." In: *Academy of Management Journal* 15.4, pp. 407–426. DOI: 10.2307/255139. JSTOR: 255139.

Bertrand, J. W. M. and D. R. Muntslag (1993). "Production control in engineer-to-order firms." In: *International Journal of Production Economics* 30. Special Issue Proceeding of the Seventh International Working Seminar on Production Economics, pp. 3–22. ISSN: 0925-5273. DOI: 10.1016/0925-5273(93)90077-X.

Bessedik, Malika, Asma Daoudi, and Karima Benatchba (2014). "A Cooperative Approach Using Ants and Bees for the Graph Coloring Problem." In: *Nature Inspired Cooperative Strategies for Optimization (NICSO 2013): Learning, Optimization and Interdisciplinary Applications*. Ed. by German Terrazas, B. Fernando E. Otero, and D. Antonio Masegosa. Cham: Springer International Publishing, pp. 179–190. ISBN: 978-3-319-01692-4. DOI: 10.1007/978-3-319-01692-4_14.

Bhasin, Sanjay and Peter Burcher (2006). "Lean viewed as a philosophy." In: *Journal of Manufacturing Technology Management* 17.1, pp. 56–72. DOI: 10.1108/17410380610639506.

Billaut, Jean-Claude and François Roubellat (1996). "Characterization of a set of schedules in a multiple resource context." In: *Journal of Decision Systems* 5.1-2, pp. 95–109. DOI: 10.1080/12460125.1996.10511676.

Bitran, Gabriel R. and Reinaldo Morabito (1999). “An Overview of Tradeoff Curves in Manufacturing System Design.” In: *Production and Operations Management* 8.1, pp. 56–75. issn: 1937-5956. doi: 10.1111/j.1937-5956.1999.tb00061.x.

Bitran, Gabriel R. and Devanath Tirupati (1989). “Tradeoff Curves, Targeting and Balancing in Manufacturing Queueing Networks.” In: *Operations Research* 37.4, pp. 547–564. doi: 10.1287/opre.37.4.547. jstor: 171256.

Blackstone, John H., Don T. Phillips, and Gary L. Hogg (1982). “A state-of-the-art survey of dispatching rules for manufacturing job shop operations.” In: *International Journal of Production Research* 20.1, pp. 27–45. doi: 10.1080/00207548208947745.

Blanchet, Max and Thomas Rinn (2016). *The Industrie 4.0 transition quantified.* Ed. by Anne Dujin and Cornelia Geissler. Accessed: 2017-04-11. Roland Berger GmbH. url: https://www.rolandberger.com/publications/publication_pdf/roland_berger_industry_40_20160609.pdf.

Błażewicz, Jacek, Erwin Pesch, and Małgorzata Sterna (2000). “The disjunctive graph machine representation of the job shop scheduling problem.” In: *European Journal of Operational Research* 127.2, pp. 317 –331. issn: 0377-2217. doi: 10.1016/S0377-2217(99)00486-5.

Blunck, Henning and Julia Bendul (2015). “Invariant-Based Production Control Reviewed: Mixing Hierarchical and Heterarchical Control in Flexible Job Shop Environments.” English. In: *Industrial Applications of Holonic and Multi-Agent Systems.* Ed. by Vladimír Mařík, Arnd Schirrmann, Damien Trentesaux, and Pavel Vrba. Vol. 9266. Lecture Notes in Computer Science. Springer International Publishing, pp. 96–107. isbn: 978-3-319-22866-2. doi: 10.1007/978-3-319-22867-9_9.

– (2016). “Controlling Myopic Behavior in Distributed Production Systems — A Classification of Design Choices.” In: *Procedia CIRP* 57.Supplement C. Factories of the Future in the digital environment - Proceedings of the 49th CIRP Conference on Manufacturing Systems, pp. 158 –163. issn: 2212-8271. doi: 10.1016/j.procir.2016.11.028.

Blunck, Henning, Dieter Armbruster, and Julia Bendul (2016). “Simultaneous Workload Allocation and Capacity Dimensioning for Distributed Production Control.” In: *Procedia CIRP* 41. Research and Innovation in Manufacturing: Key Enabling Technologies for the Factories of the Future - Proceedings of the 48th CIRP Conference on Manufacturing Systems, pp. 460 –465. issn: 2212-8271. doi: 10.1016/j.procir.2015.12.117.

– (2017). “Setting production capacities for production agents making selfish routing decisions.” In: *International Journal of Computer Integrated Manufacturing.* Special Issue: Cyber-Physical Product Creation for Industry 4.0. doi: 10.1080/0951192X.2017.1379097.

Boccalatte, Antonio, Andrea Gozzi, Massimo Paolucci, Valentino Queirolo, and Mirko Tamoglia (2004). “A multi-agent system for dynamic just-in-time manufacturing production scheduling.” In: *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No.04CH37583).* Vol. 6, 5548–5553 vol.6. doi: 10.1109/ICSMC.2004.1401077.

Bonabeau, Eric (2002). “Agent-based modeling: Methods and techniques for simulating human systems.” In: *Proceedings of the National Academy of Sciences* 99.suppl 3, pp. 7280–7287. doi: 10.1073/pnas.082080899. pubmed: 12011407.

Bonabeau, Eric, Marco Dorigo, and Guy Theraulaz (1999). *Swarm Intelligence — From Natural to Artificial Systems.* Santa Fe Institute Studies in the Sciences of Complexity. New York, Oxford: Oxford University Press. isbn: 978-0-19-513159-8.

Bond, A.H. (1990). “Distributed decision making in organizations.” In: *Systems, Man and Cybernetics, 1990. Conference Proceedings., IEEE International Conference on.* Jan. 1990, pp. 896–901. doi: 10.1109/ICSMC.1990.142254.

Bongaerts, Luc, László Monostori, Duncan McFarlane, and Botond Kádár (2000). "Hierarchy in distributed shop floor control." In: *Computers in Industry* 43.2, pp. 123 –137. issn: 0166-3615. doi: 10.1016/S0166-3615(00)00062-2.

Bonifaci, Vincenzo, Tobias Harks, and Guido Schäfer (2010). "Stackelberg Routing in Arbitrary Networks." English. In: *Mathematics of Operations Research* 35.2, pp. 330–346. issn: 0364-765X. doi: 10.1287/moor.1100.0442.

Bonney, M. C., Zongmao Zhang, M.A Head, C. C. Tien, and R. J. Barson (1999). "Are push and pull systems really so different?" In: *International Journal of Production Economics* 59.1-3, pp. 53 –64. issn: 0925-5273. doi: 10.1016/S0925-5273(98)00094-2.

Bonney, Maurice (2000). "Reflections on Production Planning and Control (PPC)." In: *Gestão & Produção* 7.3, pp. 181–207. issn: 1806-9649. doi: 10.1590/S0104-530X2000000300002.

Böse, Felix and Katja Windt (2007). "Catalogue of Criteria for Autonomous Control in Logistics." English. In: *Understanding Autonomous Cooperation and Control in Logistics*. Ed. by Michael Hülsmann and Katja Windt. Berlin, Heidelberg: Springer, pp. 57–72. isbn: 978-3-540-47449-4. doi: 10.1007/978-3-540-47450-0_5.

Boston Consulting Group (2016). *The Industry 4.0 Race — Time to accelerate*. Published 2016-04-19. url: http://www.bcg.de/documents/file208017.pdf.

Botti, Vincente and Adriana Giret (2008). *ANEMONA — A Multi-Agent Methodology for Holonic Manufacturing Systems*. Springer Series in Advanced Manufacturing. London: Springer. isbn: 978-1-84800-309-5. doi: 10.1007/978-1-84800-310-1.

Bousbia, S. and D. Trentesaux (2002). "Self-organization in distributed manufacturing control: state-of-the-art and future trends." In: *Systems, Man and Cybernetics, 2002 IEEE International Conference on*. Vol. 5. doi: 10.1109/ICSMC.2002.1176445.

Boyd, Stephen and Lieven Vandenberghe (2004). *Convex Optimization*. Cambridge: Cambridge University Press. isbn: 978-0-521-83378.

Bozarth, Cecil C., Donald P. Warsing, Barbara B. Flynn, and E. James Flynn (2009). "The impact of supply chain complexity on manufacturing plant performance." In: *Journal of Operations Management* 27.1, pp. 78 –93. issn: 0272-6963. doi: 10.1016/j.jom.2008.07.003.

Braess, Dietrich (1968). "Über ein Paradoxon aus der Verkehrsplanung." In: *Unternehmensforschung* 12.1, pp. 258–268. issn: 0042-0573. doi: 10.1007/BF01918335.

Braglia, Marcello and Alberto Petroni (1999). "Data envelopment analysis for dispatching rule selection." In: *Production Planning & Control* 10.5, pp. 454–461. doi: 10.1080/095372899232984.

Brainov, Sviatoslav (1996). "Altruistic cooperation between self-interested agents." In: *Proceedings of the 12th European Conference on Artificial Intelligence, ECAI' 96*. Ed. by Wolfgang Wahlster. Chichester: John Wiley & Sons, Ltd., pp. 519–523. isbn: 0-471-96809-9.

Brandimarte, Paolo (1993). "Routing and scheduling in a flexible job shop by tabu search." English. In: *Annals of Operations Research* 41.3, pp. 157–183. issn: 0254-5330. doi: 10.1007/BF02023073.

Brennan, R. W. (2000). "Performance comparison and analysis of reactive and planning-based control architectures for manufacturing." In: *Robotics and Computer-Integrated Manufacturing* 16.2-3, pp. 191 –200. issn: 0736-5845. doi: 10.1016/S0736-5845(00)00007-7.

Brennan, Robert W. and Douglas H. Norrie (1999). "The performance of partial dynamic hierarchies for manufacturing." In: *Proceedings of the Second International Workshop on Intelligent Manufacturing Systems, Katholieke Universiteit Leuven, Leuven, Belgium*. Ed. by Hendrik Van Brussel. Heverlee: Katholieke Univ. Leuven, Dep. Werktuigkunde, pp. 5–13. isbn: 90-73802-69-5.

– (2001). "Evaluating the performance of reactive control architectures for manufacturing production control." In: *Computers in Industry* 46.3, pp. 235–245. issn: 0166-3615. doi: 10.1016/S0166-3615(01)00108-7.

– (2003). “Metrics for evaluating distributed manufacturing control systems.” In: *Computers in Industry* 51.2. Virtual Enterprise Management, pp. 225 –235. ISSN: 0166-3615. DOI: 10.1016/S0166-3615(03)00038-1.

Brettel, Malte, Niklas Friederichsen, Michael Keller, and Marius Rosenberg (2014). “How Virtualization, Decentralization and Network Building Change the Manufacturing Landscape: An Industry 4.0 Perspective.” In: *International Journal of Mechanical, Aerospace, Industrial, Mechatronic and Manufacturing Engineering* 8.1, pp. 37 –44. ISSN: 1307-6892.

Brintrup, Alexandra, Yu Wang, and Ashutosh Tiwari (2015). “Supply Networks as Complex Systems: A Network-Science-Based Characterization.” In: *Systems Journal, IEEE.* In Press. ISSN: 1932-8184. DOI: 10.1109/JSYST.2015.2425137.

Browning, Tyson R. and Ralph D. Heath (2009). “Reconceptualizing the effects of lean on production costs with evidence from the F-22 program.” In: *Journal of Operations Management* 27.1, pp. 23 –44. ISSN: 0272-6963. DOI: 10.1016/j.jom.2008.03.009.

Brualdi, Richard A. (2009). *Introductory Combinatorics.* 5th. Pearson Higher Ed. ISBN: 0136020402.

Brückner, Sven (2000). “Return from the ant - synthetic ecosystems for manufacturing control.” PhD thesis. Humboldt Universität Berlin. URN: urn:nbn:de:kobv:11-10015033.

Brueckner, Sven A. and H. Van Dyke Parunak (2006). “Information-Driven Phase Changes in Multi-agent Coordination.” In: *Engineering Self-Organising Systems: Third International Workshop, ESOA 2005, Utrecht, The Netherlands, July 25, 2005, Revised Selected Papers.* Ed. by Sven A. Brueckner, Giovanna Di Marzo Serugendo, David Hales, and Franco Zambonelli. Berlin, Heidelberg: Springer, pp. 104–119. ISBN: 978-3-540-33352-4. DOI: 10.1007/11734697_8.

Bunimovich, Leonid A (2001). “Dynamical systems and operations research: a basic model.” In: *Discrete and Continuous Dynamical Systems - Series B* 1.2, pp. 209–218. DOI: 10.3934/dcdsb.2001.1.209.

Burton, Richard M. and Børge Obel (1995). “The validity of computational models in organization science: From model realism to purpose of the model.” In: *Computational & Mathematical Organization Theory* 1.1, pp. 57–71. ISSN: 1572-9346. DOI: 10.1007/BF01307828.

Caliendo, Frank N. and Emin Gahramanov (2013). “Myopia and pensions in general equilibrium.” In: *Journal of Economics and Finance* 37.3, pp. 375–401. ISSN: 1938-9744. DOI: 10.1007/s12197-011-9187-6.

Calvert, Randall (1992). “Leadership and Its Basis in Problems of Social Coordination.” In: *International Political Science Review* 13.1, pp. 7–24. DOI: 10.1177/019251219201300102. JSTOR: 1601435.

Carabelea, Cosmin, Olivier Boissier, and Adina Florea (2004). “Autonomy in Multi-agent Systems: A Classification Attempt.” In: *Agents and Computational Autonomy: Potential, Risks, and Solutions.* Ed. by Matthias Nickles, Michael Rovatsos, and Gerhard Weiss. Berlin, Heidelberg: Springer, pp. 103–113. ISBN: 978-3-540-25928-2. DOI: 10.1007/978-3-540-25928-2_9.

Cardin, Olivier and Pierre Castagna (2009). “Using online simulation in Holonic manufacturing systems.” In: *Engineering Applications of Artificial Intelligence* 22.7. Distributed Control of Production Systems, pp. 1025 –1033. ISSN: 0952-1976. DOI: 10.1016/j.engappai.2009.01.006.

– (2012). “Myopia of Service Oriented Manufacturing Systems: Benefits of Data Centralization with a Discrete-Event Observer.” In: *Service Orientation in Holonic and Multi-Agent Manufacturing Control.* Ed. by Theodor Borangiu, André Thomas, and Damien Trentesaux. Berlin, Heidelberg: Springer, pp. 197–210. ISBN: 978-3-642-27449-7. DOI: 10.1007/978-3-642-27449-7_15.

Cardin, Olivier, Damien Trentesaux, André Thomas, Pierre Castagna, Thierry Berger, and Hind Bril El-Haouzi (2015). “Coupling predictive scheduling and reactive control in manufacturing hybrid control architectures: state of the art and future challenges." In: *Journal of Intelligent Manufacturing*, pp. 1–15. ISSN: 1572-8145. DOI: 10.1007/s10845-015-1139-0.

Caridi, Maria and Andrea Sianesi (2000). "Multi-agent systems in production planning and control: An application to the scheduling of mixed-model assembly lines." In: *International Journal of Production Economics* 68.1, pp. 29 –42. ISSN: 0925-5273. DOI: 10.1016/S0925-5273(99)00097-3.

Carley, Kathleen M. (1995). "Computational and mathematical organization theory: Perspective and directions." In: *Computational & Mathematical Organization Theory* 1.1, pp. 39–56. ISSN: 1572-9346. DOI: 10.1007/BF01307827.

– (1997). "Organizational adaptation." In: *Annals of Operations Research* 75.0, pp. 25–47. ISSN: 1572-9338. DOI: 10.1023/A:1018963630536.

Carley, Kathleen M. and David M. David M. Svoboda (1996). "Modeling Organizational Adaptation as a Simulated Annealing Process." In: *Sociological Methods & Research* 25.1, pp. 138–168. DOI: 10.1177/0049124196025001005.

Carley, Kathleen M. and Les Gasser (1999). "Multiagent Systems." In: ed. by Gerhard Weiss. Cambridge, MA: MIT Press. Chap. Computational Organization Theory, pp. 299–330. ISBN: 0-262-23203-0. ACM ID: 305606.305613.

Carley, Kathleen M. and Yuqing Ren (2001). "Tradeoffs between performance and adaptability for C3I architectures." In: *Proceedings of the 6th International Command and Control Research and Technology Symposium.* eprint: http://dodccrp.org/events/6th_ICCRTS/Tracks/Papers/Track2/082_tr2.pdf.

Carley, Kathleen M, Yuqing Ren, and David Krackhardt (2000). "Measuring and modeling change in C3I architecture." In: *Proceedings of the Command and Control Research and Technology Symposium.* eprint: http://dodccrp.org/events/2000_CCRTS/html/pdf_papers/Track_3/041.pdf.

Carlin, Alan and Shlomo Zilberstein (2009). "Myopic and Non-myopic Communication under Partial Observability." In: *Web Intelligence and Intelligent Agent Technologies, 2009. WI-IAT '09. IEEE/WIC/ACM International Joint Conferences on.* Vol. 2, pp. 331–338. DOI: 10.1109/WI-IAT.2009.174.

Carneiro, Murillo G. and Gina M. B. Oliveira (2013). "Synchronous cellular automata-based scheduler initialized by heuristic and modeled by a pseudo-linear neighborhood." In: *Natural Computing* 12.3, pp. 339–351. ISSN: 1572-9796. DOI: 10.1007/s11047-013-9375-8.

Cassandras, Christos G. and Stéphane Lafortune (2008). *Introduction to Discrete Event Systems.* 2nd. Boston: Springer. ISBN: 978-0-387-33332-8. DOI: 10.1007/978-0-387-68612-7.

Castellani, Brian and Frederic William Hafferty (2009). *Sociology and Complexity Science — A new Field of Inquiry.* Berlin, Heidelberg: Springer. ISBN: 978-3-540-88461-3. DOI: 10.1007/978-3-540-88462-0.

Castellanos, Mayolo Alberto López (2012). *Agent Based Simulation Approach to Assess Supply Chain Complexity and Its Impact on Performance.* Ed. by Wolfgang Kersten. Vol. 12. Supply Chain, Logistics and Operations Management. Lohmar: Josef Eul-Verlag. ISBN: 978-3-8441-0209-3.

Cavalieri, Sergio, Marco Garetti, Marco Macchi, and Marco Taisch (2000). "An experimental benchmarking of two multi-agent architectures for production scheduling and control." In: *Computers in Industry* 43.2, pp. 139 –152. ISSN: 0166-3615. DOI: 10.1016/S0166-3615(00)00063-4.

Cavalieri, Sergio, Marco Macchi, and Paul Valckenaers (2003). "Benchmarking the performance of manufacturing control systems: design principles for a web-based simulated testbed." In: *Journal of Intelligent Manufacturing* 14.1, pp. 43–58. ISSN: 1572-8145. DOI: 10.1023/A:1022287212706.

Chakravorty, Satya S. and J. Brian Atwater (2006). "Bottleneck management: theory and practice." In: *Production Planning & Control* 17.5, pp. 441–447. issn: 0953-7287. doi: 10.1080/09537280600682752.

Chandy, K. M., J. Hogarth, and C. H. Sauer (1977). "Selecting Capacities in Computer Communication Systems." In: *IEEE Transactions on Software Engineering* SE-3.4, pp. 290–295. issn: 0098-5589. doi: 10.1109/TSE.1977.231146.

Chen, Hong and Avi Mandelbaum (1994). "Hierarchical Modeling of Stochastic Networks, Part I: Fluid Models." In: *Stochastic Modeling and Analysis of Manufacturing Systems*. Ed. by David D. Yao. New York, NY: Springer, pp. 47–105. isbn: 978-1-4612-2670-3. doi: 10.1007/978-1-4612-2670-3_2.

Choi, Thomas Y., Kevin J. Dooley, and Manus Rungtusanatham (2001). "Supply networks and complex adaptive systems: control versus emergence." In: *Journal of Operations Management* 19.3, pp. 351 –366. issn: 0272-6963. doi: 10.1016/S0272-6963(00)00068-1.

Christensen, James H. (1994). "Holonic manufacturing systems: initial architecture and standards directions." In: *Proceedings of the 1st European Conference on Holonic Manufacturing Systems*. Hannover, pp. 1–20.

Christodoulou, George and Elias Koutsoupias (2005). "The Price of Anarchy of Finite Congestion Games." In: *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing, STOC '05, Baltimore, MD*. New York, NY, USA: ACM, pp. 67–73. isbn: 1-58113-960-8. doi: 10.1145/1060590.1060600. acm id: 1060600.

Chryssolouris, George (2006). *Manufacturing systems: theory and practice*. Ed. by Frederick F Ling and William Howard Hart. 2nd. Springer-Verlag. isbn: 9780387256832. doi: 10.1007/0-387-28431-1.

Chryssolouris, George, Konstantinos Efthymiou, Nikolaos Papakostas, Dimitris Mourtzis, and Aris Pagoropoulos (2013). "Flexibility and complexity: is it a trade-off?" In: *International Journal of Production Research* 51.23-24, pp. 6788–6802. doi: 10.1080/00207543.2012.761362.

Coase, R. H. (1937). "The Nature of the Firm." In: *Economica* 4.16, pp. 386–405. issn: 00130427, 14680335. doi: 10.1111/j.1468-0335.1937.tb00002.x.

Cocchi, Ron, Scott Shenker, Deborah Estrin, and Lixia Zhang (1993). "Pricing in Computer Networks: Motivation, Formulation, and Example." In: *IEEE/ACM Trans. Netw.* 1.6 (Dec. 1993), pp. 614–627. issn: 1063-6692. doi: 10.1109/90.266050. acm id: 174377.

Cochran, J. K. and H. A. Kaylani (2008). "Optimal design of a hybrid push/pull serial manufacturing system with multiple part types." In: *International Journal of Production Research* 46.4, pp. 949–965. doi: 10.1080/00207540600905547.

Cohen, Joel E. and Frank P. Kelly (1990). "A Paradox of Congestion in a Queuing Network." In: *Journal of Applied Probability* 27.3, pp. 730–734. issn: 0021-9002. jstor: 3214558.

Coleman Jr., Henry J. (1999). "What Enables Self-Organizing Behavior in Businesses." In: *Emergence* 1.1, p. 33. issn: 1521-3250.

Conte, Rosario and Cristiano Castelfranchi (1995). *Cognitive and social action*. London: UCL Press Limited. isbn: 1-85728-186-1.

Correa, José R., Andreas S. Schulz, and Nicolás E. Stier-Moses (2004). "Selfish Routing in Capacitated Networks." English. In: *Mathematics of Operations Research* 29.4, pp. 961–976. issn: 0364765X. doi: 10.1287/moor.1040.0098. jstor: 30035655.

Costa, L. da F., F. A. Rodrigues, G. Travieso, and P. R. Villas Boas (2007). "Characterization of complex networks: A survey of measurements." In: *Advances in Physics* 56.1, pp. 167–242. doi: 10.1080/00018730601170527.

Coviello, Lorenzo, Massimo Franceschetti, Mathew D. McCubbins, Ramamohan Paturi, and Andrea Vattani (2012). "Human Matching Behavior in Social Networks: An Algorithmic Perspective." In: *PLoS ONE* 7.8, e41900. doi: 10.1371/journal.pone.0041900.

Crutchfield, James P. and Melanie Mitchell (1995). "The evolution of emergent computation." In: *Proceedings of the National Academy of Sciences* 92.23, pp. 10742–10746. pubmed: 11607588.

Cummings, Thomas G. (1978). "Self-Regulating Work Groups: A Socio-Technical Synthesis." In: *The Academy of Management Review* 3.3, pp. 625–634. issn: 0363-7425. doi: 10.2307/257551. jstor: 257551.

Czumaj, Artur, Piotr Krysta, and Berthold Vöcking (2010). "Selfish Traffic Allocation for Server Farms." In: *SIAM Journal on Computing* 39.5, pp. 1957–1987. doi: 10.1137/070693862.

Dafermos, Stella C. (1968). "Traffic assignment and resource allocation in transportation networks." PhD thesis. John Hopkins University.

Dafermos, Stella C. and Frederick T. Sparrow (1969). "The traffic assignment problem for a general network." In: *Journal of Research of the National Bureau of Standards, Series B* 73.2, pp. 91–118. doi: 10.6028/jres.073B.010.

Dai, J. G., J. J. Hasenbein, and J. H. Vande Vate (1999). "Stability of a three-station fluid network." In: *Queueing Systems* 33.4, pp. 293–325. issn: 1572-9443. doi: 10.1023/A:1019184331042.

Dal Forno, Arianna and Ugo Merlone (2007). "The Emergence of Effective Leaders: An Experimental and Computational Approach." In: *Complex Systems Leadership Theory.* Ed. by James K. Hazy, Jeffrey A. Goldstein, and Benyamin B. Lichtenstein. Vol. 1. Exploring Organizational Complexity. Mansfield, MA: ISCE Publishing. Chap. 11, pp. 205–225. isbn: 978-0-9791688-6-4.

Dal, Bulent, Phil Tugwell, and Richard Greatbanks (2000). "Overall equipment effectiveness as a measure of operational improvement - A practical analysis." In: *International Journal of Operations & Production Management* 20.12, pp. 1488 –1502. doi: 10.1108/01443570010355750.

Dauzère-Pérès, Stéphane and Jan Paulli (1997). "An integrated approach for modeling and solving the general multiprocessor job-shop scheduling problem using tabu search." In: *Annals of Operations Research* 70.1-4, pp. 281 –306. issn: 02545330. doi: 10.1023/A:1018930406487.

Dauzère-Pérès, Stéphane, W. Roux, and J.B. Lasserre (1998). "Multi-resource shop scheduling with resource flexibility." In: *European Journal of Operational Research* 107.2, pp. 289 –305. issn: 0377-2217. doi: 10.1016/S0377-2217(97)00341-X.

De Carolis, Anna, Giacomo Tavolo, and Marco Tisch (2017). "Gap Analysis on Research and Innovation for Cyber-Physical Systems in Manufacturing." In: *Service Orientation in Holonic and Multi-Agent Manufacturing, Proceedings of SOHOMA 2016.* Ed. by Theodor Borangiu, Damien Trentesaux, André Thomas, Paulo Leitão, and José Barata Oliveira. Vol. 694. Studies in Computational Intelligence. Under Publication. Springer International Publishing, pp. 61–70. isbn: 978-3-319-51099-6. doi: 10.1007/978-3-319-51100-9_6.

De Werra, Dominique and Alain Hertz (2015). "Chromatic scheduling." In: *Topics in Chromatic Graph Theory.* Ed. by Lowell W. Beineke and Robin J. Wilson. Vol. 156. Encyclopedia of Mathematics and Its Applications. Cambridge: Cambridge University Press, May 2015. Chap. 12, pp. 255–276. isbn: 978-1-107-03350-4. doi: 10.1017/CBO9781139519793.015.

De Wolf, Tom and Tom Holvoet (2005). "Emergence Versus Self-Organisation: Different Concepts but Promising When Combined." In: *Engineering Self-Organising Systems.* Ed. by Sven A. Brueckner, Giovanna Marzo Serugendo, Anthony Karageorgos, and Radhika Nagpal. Vol. 3464. Lecture Notes in Computer Science. Springer, pp. 1–15. isbn: 978-3-540-26180-3. doi: 10.1007/11494676_1.

Dechaume-Moncharmont, François-Xavier, Anna Dornhaus, Alasdair I Houston, John M McNamara, Edmund J Collins, and Nigel R Franks (2005). "The hidden cost of information in collective foraging." In: *Proceedings of the Royal Society of London B: Biological Sciences* 272.1573, pp. 1689–1695. issn: 0962-8452. doi: 10.1098/rspb.2005.3137.

Decker, Keith S. (1987). "Distributed problem-solving techniques: A survey." In: *IEEE Transactions on Systems, Man, and Cybernetics* 17.5, pp. 729–740. issn: 0018-9472. doi: 10.1109/TSMC.1987.6499280.

Delfmann, Werner, Michael ten Hompel, Wolfgang Kersten, Thorsten Schmidt, and Wolfgang Stölzle (2017). *Logistik als Wissenschaft — zentrale Forschungsfragen in Zeiten der vierten industriellen Revolution.* Positionspapier des Wissenschaftlichen Beirats. Bremen: Bundesvereinigung Logistik.

Deshmukh, Abhijit V., Joseph J. Talavage, and Moshe M. Barash (1998). "Complexity in manufacturing systems, Part 1: Analysis of static complexity." English. In: *IIE Transactions* 30.7, pp. 645–655. issn: 0740-817X. doi: 10.1023/A:1007542328011.

Deutsch, Andreas and Sabine Dormann (2005). *Cellular Automaton Modeling of Biological Pattern Formation.* Modeling and Simulation in Science, Engineering and Technology. Boston, MA: Birkhäuser. isbn: 978-0-8176-4281-5. doi: 10.1007/b138451.

Dewar, Robert and Jerald Hage (1978). "Size, Technology, Complexity, and Structural Differentiation: Toward a Theoretical Synthesis." In: *Administrative Science Quarterly* 23.1, pp. 111–136. issn: 0001-8392. doi: 10.2307/2392436. jstor: 2392436.

Di Caro, Gianni and Marco Dorigo (1997). *AntNet: A mobile agents approach to adaptive routing.* Technical Report IRIDIA/97-12. IRIDIA, Université Libre de Bruxelles, Belgium.

Dilts, D.M., N.P. Boyd, and H.H. Whorms (1991). "The evolution of control architectures for automated manufacturing systems." In: *Journal of Manufacturing Systems* 10.1, pp. 79 –93. issn: 0278-6125. doi: 10.1016/0278-6125(91)90049-8.

Dittel, Agnes, Simone Göttlich, and Ute Ziegler (2011). "Optimal design of capacitated production networks." In: *Optimization and Engineering* 12.4, pp. 583–602. issn: 1573-2924. doi: 10.1007/s11081-010-9123-1.

Do, Minh Binh and Subbarao Kambhampati (2001). "Planning as constraint satisfaction: Solving the planning graph by compiling it into CSP." In: *Artificial Intelligence* 132.2, pp. 151 –182. issn: 0004-3702. doi: 10.1016/S0004-3702(01)00128-X.

Doran, J. E., S. Franklin, Nicholas R. Jennings, and T. J. Norman (1997). "On cooperation in multi-agent systems." In: *The Knowledge Engineering Review* 12.3, pp. 309–314. issn: 1469-8005.

Dubey, Pradeep (1986). "Inefficiency of Nash Equilibria." In: *Mathematics of Operations Research* 11.1, pp. 1–8. doi: 10.1287/moor.11.1.1.

Duffie, Neil A. (1990). "Synthesis of Heterarchical manufacturing systems." In: *Computers in Industry* 14.1, pp. 167 –174. issn: 0166-3615. doi: 10.1016/0166-3615(90)90118-9.

Duffie, Neil A. and Rex S. Piper (1986). "Nonhierarchical control of manufacturing systems." In: *Journal of Manufacturing Systems* 5.2, p. 141. issn: 0278-6125. doi: 10.1016/0278-6125(86)90036-1.

– (1987). "Non-hierarchical control of a flexible manufacturing cell." In: *Robotics and Computer-Integrated Manufacturing* 3.2, pp. 175 –179. issn: 0736-5845. doi: 10.1016/0736-5845(87)90099-8.

Duffie, Neil A. and V. V. Prabhu (1996). "Heterarchical control of highly distributed manufacturing systems." In: *International Journal of Computer Integrated Manufacturing* 9.4, pp. 270–281. doi: 10.1080/095119296131562.

Duffie, Neil A. and Vittaldas V. Prabhu (1994). "Real-time distributed scheduling of heterarchical manufacturing systems." In: *Journal of Manufacturing Systems* 13.2, pp. 94 –107. issn: 0278-6125. doi: 10.1016/0278-6125(94)90025-6.

Duffie, Neil A. and L. Shi (2010). "Dynamics of WIP Regulation in Large Production Networks of Autonomous Work Systems." In: *IEEE Transactions on Automation Science and Engineering* 7.3, pp. 665–670. issn: 1545-5955. doi: 10.1109/TASE.2009.2036374.

Duffie, Neil, John Fenske, and Madhu Vadali (2012). "Coordination of capacity adjustment modes in work systems with autonomous WIP regulation." In: *Logistics Research* 5.3, pp. 99–104. issn: 1865-0368. doi: 10.1007/s12159-012-0088-7.

Ebeling, Werner, Jan Freund, and Frank Schweitzer (1998). *Komplexe Strukturen: Entropie und Information.* Stuttgart: Teubner. isbn: 9783815430323.

Eglese, R.W. (1990). "Simulated annealing: A tool for operational research." In: *European Journal of Operational Research* 46.3, pp. 271 –281. issn: 0377-2217. doi: 10.1016/0377-2217(90)90001-R.

El-Bouri, Ahmed and Pramit Shah (2006). "A neural network for dispatching rule selection in a job shop." In: *The International Journal of Advanced Manufacturing Technology* 31.3, pp. 342–349. issn: 1433-3015. doi: 10.1007/s00170-005-0190-y.

Enarsson, Leif (2006). *Future Logistics Challenges.* CBS Press. isbn: 8763001705.

Enemark, Daniel P., Mathew D. McCubbins, Ramamohan Paturi, and Nicholas Weller (2011). "Does More Connectivity Help Groups to Solve Social Problems." In: *Proceedings of the 12th ACM Conference on Electronic Commerce.* EC '11. San Jose, California, USA: ACM, pp. 21–26. isbn: 978-1-4503-0261-6. doi: 10.1145/1993574.1993578. acm id: 1993578.

Enemark, Daniel, Mathew D. McCubbins, and Nicholas Weller (2014). "Knowledge and networks: An experimental test of how network knowledge affects coordination." In: *Social Networks* 36. Special Issue on Political Networks, pp. 122 –133. issn: 0378-8733. doi: 10.1016/j.socnet.2012.10.001.

Estrada, Ernesto and Eusebio Vargas-Estrada (2013). "How Peer Pressure Shapes Consensus, Leadership, and Innovations in Social Groups." In: *Scientific Reports* 3. doi: 10.1038/srep02905.

European Commission (2016). *Digitising European Industry.* Accessed: 2017-04-11. url: https://ec.europa.eu/digital-single-market/en/digitising-european-industry.

Evans, Paul A. L. and Yves Doz (1999). "Human Resource Management: Critical Perspectives on Business and Management." In: ed. by Michael Poole. Vol. 3. London, New York: Taylor & Francis. Chap. Dualities: A paradigm for human resource and organizational development in complex multinationals, pp. 83–93. isbn: 0-415-19339-7.

Fabiunke, Marko (1999). "Parallel Distributed Constraint Satisfaction." In: *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications, PDPTA 1999.* CSREA Press, pp. 1585–1591. isbn: 1-892512-15-7.

Fabiunke, Marko and Gerd Kock (2000). "A Connectionist Method to Solve Job Shop Problems." In: *Cybernetics and Systems* 31.5, pp. 491–506. doi: 10.1080/01969720050045994.

Falcone, Rino and Christiano Castelfranchi (2001). "The human in the loop of a delegated agent: the theory of adjustable social autonomy." In: *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 31.5, pp. 406–418. issn: 1083-4427. doi: 10.1109/3468.952715.

Falk, James E. and Karla L. Hoffman (1986). "Concave Minimization Via Collapsing Polytopes." In: *Operations Research* 34.6, pp. 919–929. doi: 10.1287/opre.34.6.919. jstor: 170772.

Fang, Christina, Jeho Lee, and Melissa A. Schilling (2010). "Balancing Exploration and Exploitation Through Structural Design: The Isolation of Subgroups and Organizational Learning." In: *Organization Science* 21.3, pp. 625–642. doi: 10.1287/orsc.1090.0468. jstor: 40792435.

Farid, A.M. and L. Ribeiro (2015). "An Axiomatic Design of a Multiagent Reconfigurable Mechatronic System Architecture." In: *Industrial Informatics, IEEE Transactions on* 11.5 (Oct. 2015), pp. 1142–1155. issn: 1551-3203. doi: 10.1109/TII.2015.2470528.

Fattahi, Parviz, Mohammad Saidi Mehrabad, and Fariborz Jolai (2007). "Mathematical modeling and heuristic approaches to flexible job shop scheduling problems." English. In: *Journal of Intelligent Manufacturing* 18.3, pp. 331–342. issn: 0956-5515. doi: 10.1007/s10845-007-0026-8.

Fedoruk, Alan and Jörg Denzinger (2006). "A general framework for multi-agent search with individual and global goals: Stakeholder search." In: *International Transactions on Systems Science and Applications* 1 (4 2006), pp. 357–362. issn: 1751-1461.

Felstead, Alan (2016). "Tracing the connections: short-termism, training and recession." In: *The International Journal of Human Resource Management*. In Press. doi: 10.1080/09585192.2016.1184176.

Fischer, Klaus, Michael Schillo, and Jörg Siekmann (2003). "Holonic Multiagent Systems: A Foundation for the Organisation of Multiagent Systems." In: *Holonic and Multi-Agent Systems for Manufacturing: First International Conference on Industrial Applications of Holonic and Multi-Agent Systems, HoloMAS 2003, Prague, Czech Republic, September 1-3, 2003. Proceedings.* Ed. by Vladimír Mařík, Duncan McFarlane, and Paul Valckenaers. Berlin, Heidelberg: Springer, pp. 71–80. isbn: 978-3-540-45185-3. doi: 10.1007/978-3-540-45185-3_7.

Fitzpatrick, Stephen and Lambert Meertens (2001). *Soft, Real-Time, Distributed Graph Coloring using Decetraluzed, Synchronous, Stochastic, Iterative-Repair, Anytime Algorithms — A Framework.* Tech. rep. KES.U.01.05. Palo Alto, CA: Kestrel Institute.

– (2002). "Experiments on dense graphs with a stochastic, peer-to-peer colorer." In: *Probabilistic Approaches in Search — Papers from the AAAI Workshop (Technical Report WS-02-14).* Menlo Park, CA: AAAI Press, pp. 24–28. isbn: 1-57735-167-3.

Fleischmann, Bernhard and Herbert Meyr (2003). "Planning Hierarchy, Modeling and Advanced Planning Systems." In: *Supply Chain Management: Design, Coordination and Operation.* Vol. 11. Handbooks in Operations Research and Management Science. Elsevier, pp. 455 –523. doi: 10.1016/S0927-0507(03)11009-2.

Fratta, L., M. Gerla, and L. Kleinrock (1973). "The flow deviation method: An approach to store-and-forward communication network design." In: *Networks* 3.2, pp. 97–133. issn: 1097-0037. doi: 10.1002/net.3230030202.

Freuder, Eugene C. (1995). "The many paths to satisfaction." In: *Constraint Processing: Selected Papers.* Ed. by Manfred Meyer. Berlin, Heidelberg: Springer, pp. 103–119. isbn: 978-3-540-49281-8. doi: 10.1007/3-540-59479-5_21.

Frizelle, Gerry and Eric Woodcock (1995). "Measuring complexity as an aid to developing operational strategy." In: *International Journal of Operations & Production Management* 15.5, pp. 26–39. issn: 0144-3577. doi: 10.1108/01443579510083640.

Fuenmayor, Ramsés (1991). "The roots of reductionism: A counter-ontoepistemology for a systems approach." In: *Systems practice* 4.5, pp. 419–448. issn: 1573-9295. doi: 10.1007/BF01104460.

Gairing, Martin, Burkhard Monien, and Karsten Tiemann (2008). "Selfish Routing with Incomplete Information." English. In: *Theory of Computing Systems* 42.1, pp. 91–130. issn: 1432-4350. doi: 10.1007/s00224-007-9015-8.

Gairing, Martin, Tobias Harks, and Max Klimm (2014). "Complexity and Approximation of the Continuous Network Design Problem." In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014).* Ed. by Klaus Jansen, José D. P. Rolim, Nikhil R. Devanur, and Cristopher Moore. Vol. 28. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 226–241. isbn: 978-3-939897-74-3. doi: 10.4230/LIPIcs.APPROX-RANDOM.2014.226. urn: urn:nbn:de:0030-drops-46998.

Galbraith, Jay R. (1973). *Designing Complex Organizations.* 1st. Boston, MA: Addison-Wesley Longman Publishing Co., Inc. isbn: 0201025590.

– (1974). "Organization Design: An Information Processing View." In: *Interfaces* 4.3, pp. 28–36. doi: 10.1287/inte.4.3.28.

Garey, M. R., D. S. Johnson, and Ravi Sethi (1976). "The Complexity of Flowshop and Jobshop Scheduling." In: *Mathematics of Operations Research* 1.2, pp. 117–129. doi: 10.1287/moor.1.2.117. eprint: http://dx.doi.org/10.1287/moor.1.2.117.

Garg, N., M. Papatriantafilou, and P. Tsigas (1996). "Distributed list coloring: how to dynamically allocate frequencies to mobile base stations." In: *Parallel and Distributed Processing, 1996., Eighth IEEE Symposium on*, pp. 18 –25. doi: 10.1109/SPDP.1996.570312.

Gasser, Les (1991). "Social conceptions of knowledge and action: DAI foundations and open systems semantics." In: *Artificial Intelligence* 47.1, pp. 107–138. issn: 0004-3702. doi: 10.1016/0004-3702(91)90052-L.

Ge, Xiaohua, Fuwen Yang, and Qing-Long Han (2017). "Distributed networked control systems: A brief overview." In: *Information Sciences* 380, pp. 117 –131. issn: 0020-0255. doi: 10.1016/j.ins.2015.07.047.

Gell-Mann, Murray (2002). "What Is Complexity?" In: *Complexity and Industrial Clusters: Dynamics and Models in Theory and Practice*. Ed. by Alberto Quadrio Curzio and Marco Fortis. Heidelberg: Physica-Verlag HD. Chap. 1, pp. 13–24. isbn: 978-3-642-50007-7. doi: 10.1007/978-3-642-50007-7_2.

Gershwin, S., R. Hildebrant, Rajan Suri, and S. Mitter (1986). "A control perspective on recent trends in manufacturing systems." In: *IEEE Control Systems Magazine* 6.2 (Apr. 1986), pp. 3–15. issn: 0272-1708. doi: 10.1109/MCS.1986.1105071.

Gibbens, Richard J. and Frank P. Kelly (1999). "Resource pricing and the evolution of congestion control." In: *Automatica* 35.12, pp. 1969 –1985. issn: 0005-1098. doi: 10.1016/S0005-1098(99)00135-1.

Giret, Adriana and Damien Trentesaux (2015). "Software Engineering Methods for Intelligent Manufacturing Systems: A Comparative Survey." English. In: *Industrial Applications of Holonic and Multi-Agent Systems*. Ed. by Vladimír Mařík, Arnd Schirrmann, Damien Trentesaux, and Pavel Vrba. Vol. 9266. Lecture Notes in Computer Science. Springer International Publishing, pp. 11–21. isbn: 978-3-319-22866-2. doi: 10.1007/978-3-319-22867-9_2.

Glanzer, Klaus, Alexander Hämmerle, and Ralf Geurts (2001). "The application of ZEUS agents in manufacturing environments." In: *Proceedings. 12th International Workshop on Database and Expert Systems Applications*. IEEE, pp. 628–632. doi: 10.1109/DEXA.2001.953129.

Glinton, Robin, Praveen Paruchuri, Paul Scerri, and Katia Sycara (2010). "Self-Organized Criticality of Belief Propagation in Large Heterogeneous Teams." In: *Dynamics of Information Systems: Theory and Applications*. Ed. by Michael J. Hirsch, Panos M. Pardalos, and Robert Murphey. New York, NY: Springer, pp. 165–182. isbn: 978-1-4419-5689-7. doi: 10.1007/978-1-4419-5689-7_8.

Goerner, Sally J., Bernard Lietaer, and Robert E. Ulanowicz (2009). "Quantifying economic sustainability: Implications for free-enterprise theory, policy and practice." In: *Ecological Economics* 69.1. The DPSIR framework for Biodiversity Assessment, pp. 76 –81. issn: 0921-8009. doi: 10.1016/j.ecolecon.2009.07.018.

Goldenberg, Jacob, Sangman Han, Donald R. Lehmann, and Jae Weon Hong (2009). "The Role of Hubs in the Adoption Process." In: *Journal of Marketing* 73.2, pp. 1–13. issn: 0022-2429. jstor: 20619007.

Goldratt, Eliyahu M (1999). *Theory of constraints*. Great Barrington, MA: North River Press. isbn: 0-88427-166-8.

Goldstein, Jeffrey (1999). "Emergence as a Construct: History and Issues." In: *Emergence* 1.1, pp. 49–72. doi: 10.1207/s15327000em0101_4.

Gorecky, Dominic, Mathias Schmitt, Matthias Loskyll, and Detlef Zühlke (2014). "Human-machine-interaction in the industry 4.0 era." In: *12th IEEE International Conference on Industrial Informatics (INDIN)*, pp. 289–294. doi: 10.1109/INDIN.2014.6945523.

Graetz, Fiona and Aaron C.T. Smith (2008). "The role of dualities in arbitrating continuity and change in forms of organizing." In: *International Journal of Management Reviews* 10.3, pp. 265–280. issn: 1468-2370. doi: 10.1111/j.1468-2370.2007.00222.x.

Graham, R. L. (1966). "Bounds for Certain Multiprocessing Anomalies." In: *Bell System Technical Journal* 45.9, pp. 1563–1581. issn: 1538-7305. doi: 10.1002/j.1538-7305.1966.tb01709.x.

Graves, Stephen C. (2011). "Uncertainty and Production Planning." In: *Planning Production and Inventories in the Extended Enterprise: A State of the Art Handbook, Volume 1.* Ed. by G. Karl Kempf, Pınar Keskinocak, and Reha Uzsoy. Boston, MA: Springer, pp. 83–101. isbn: 978-1-4419-6485-4. doi: 10.1007/978-1-4419-6485-4_5.

Green, Shaw, Leon Hurst, Brenda Nangle, Pádraig Cunningham, Fergal Somers, and Richard Evans (1997). *Software agents: A review.* Tech. rep. Trinity College Dublin, Department of Computer Science.

Grewal, Rajdeep, Gary L. Lilien, and Girish Mallapragada (2006). "Location, Location, Location: How Network Embeddedness Affects Project Success in Open Source Systems." In: *Management Science* 52.7, pp. 1043–1056. doi: 10.1287/mnsc.1060.0550. jstor: 256589.

Gronau, Norbert and Hanna Theuer (2015). "Industrie 4.0 in Europa — Ein Vergleich europäischer Industrie 4.0-Initiativen." In: *Industrie 4.0 Management* 31 (5 2015), pp. 31–34.

– (2016). "Determination of the Optimal Degree of Autonomy in a Cyber-Physical Production System." In: vol. 57, pp. 110 –115. doi: 10.1016/j.procir.2016.11.020.

Grundstein, Sebastian, S. Schukraft, B. Scholz-Reiter, and M. Freitag (2015). "Coupling order release methods with autonomous control methods – an assessment of potentials by literature review and discrete event simulation." In: *International Journal of Production Management and Engineering* 3.1, p. 43. doi: 10.4995/ijpme.2015.3199.

Grünig, Rudolf and Richard Kühn (2013). "Decision Problems." In: *Successful Decision-Making: A Systematic Approach to Complex Problems.* Berlin, Heidelberg: Springer, pp. 7–14. isbn: 978-3-642-32307-2. doi: 10.1007/978-3-642-32307-2_2.

Gubbi, Jayavardhana, Rajkumar Buyya, Slaven Marusic, and Marimuthu Palaniswami (2013). "Internet of Things (IoT): A vision, architectural elements, and future directions." In: *Future Generation Computer Systems* 29.7. Including Special sections: Cyber-enabled Distributed Computing for Ubiquitous Cloud and Network Services & Cloud Computing and Scientific Applications — Big Data, Scalable Analytics, and Beyond, pp. 1645 –1660. issn: 0167-739X. doi: 10.1016/j.future.2013.01.010.

Gudehus, Timm and Herbert Kotzab (2009). *Comprehensive Logistics.* Berlin: Springer. isbn: 978-3-540-30722-8. doi: 10.1007/978-3-540-68652-1.

Guetzkow, Harold and Herbert A. Simon (1955). "The Impact of Certain Communication Nets Upon Organization and Performance in Task-Oriented Groups." In: *Management Science* 1.3-4, pp. 233–250. doi: 10.1287/mnsc.1.3-4.233. jstor: 2627162.

Gunasekaran, Angappa and Eric W.T. Ngai (2012). "The future of operations management: An outlook and analysis." In: *International Journal of Production Economics* 135.2. Green Manufacturing and Distribution in the Fashion and Apparel Industries, pp. 687 –701. issn: 0925-5273. doi: 10.1016/j.ijpe.2011.11.002.

Hadzhiev, Borislav, Katja Windt, Werner Bergholz, and Marc-Thorsten Hütt (2009). "A Model of Graph Coloring Dynamics with Attention Waves and Strategic Waiting." In: *Advances in Complex Systems* 12.6, pp. 549–564. doi: 10.1142/S0219525909002386.

Hall, Nicholas G. and Marc E. Posner (2001). "Generating Experimental Data for Computational Testing with Machine Scheduling Applications." In: *Operations Research* 49.6, pp. 854–865. doi: 10.1287/opre.49.6.854.10014. jstor: 3088599.

Halldórsson, Magnús M. and Guy Kortsarz (2004). "Multicoloring: Problems and Techniques." In: *Mathematical Foundations of Computer Science 2004: 29th International Symposium, MFCS 2004, Prague, Czech Republic, August 22-27, 2004. Proceedings.* Ed. by Jiří Fiala, Václav Koubek, and Jan Kratochvíl. Berlin, Heidelberg: Springer, pp. 25–41. isbn: 978-3-540-28629-5. doi: 10.1007/978-3-540-28629-5_2.

Halldórsson, Magnús M., Guy Kortsarz, and Hadas Shachnai (2003). "Sum Coloring Interval and k-Claw Free Graphs with Application to Scheduling Dependent Jobs." English. In: *Algorithmica* 37 (3 2003), pp. 187–209. issn: 0178-4617. doi: 10.1007/s00453-003-1031-8.

Hansen, Pierre, Julio Kuplinsky, and Dominique Werra (1997). "Mixed graph colorings." English. In: *Mathematical Methods of Operations Research* 45.1, pp. 145–160. issn: 1432-2994. doi: 10.1007/BF01194253.

Hatvany, József (1985). "Intelligence and cooperation in heterarchic manufacturing systems." In: *Robotics and Computer-Integrated Manufacturing* 2.2, pp. 101 –104. issn: 0736-5845. doi: 10.1016/0736-5845(85)90065-1.

Haurie, A. and P. Marcotte (1985). "On the relationship between Nash-Cournot and Wardrop equilibria." In: *Networks* 15.3, pp. 295–308. issn: 1097-0037. doi: 10.1002/net.3230150303.

Hax, Arnoldo C. and Harlan C. Meal (1973). *Hierarchical integration of production planning and scheduling.* Sloan School of Management Working Paper 656-73. Cambridge, MA: Massachusetts Institute of Technology. HDL: 1721.1/1868.

Hazy, James K. (2007). "Computer models of leadership: Foundations for a new discipline or meaningless diversion?" In: *The Leadership Quarterly* 18.4. Leadership and Complexity, pp. 391 –410. issn: 1048-9843. doi: 10.1016/j.leaqua.2007.04.007.

– (2008). "Toward a theory of leadership in complex systems: Computational modeling explorations." In: *Nonlinear Dynamics, Psychology, and Life Sciences* 12.3, pp. 281 –310.

Hazy, James K, William P. Millhiser, and Daniel Solow (2007). "Mathematical and Computational Models of Leadership: Past and Present." In: *Complex Systems Leadership Theory.* Ed. by James K. Hazy, Jeffrey A. Goldstein, and Benyamin B. Lichtenstein. Vol. 1. Exploring Organizational Complexity. Mansfield, MA: ISCE Publishing. Chap. 9, pp. 163–193. isbn: 978-0-9791688-6-4.

Hedberg, Bo L. T., Paul C. Nystrom, and William H. Starbuck (1976). "Camping on Seesaws: Prescriptions for a Self-Designing Organization." In: *Administrative Science Quarterly* 21.1, pp. 41–65. issn: 0001-8392. doi: 10.2307/2391877. jstor: 2391877.

Helbing, Dirk, Hendrik Ammoser, and Christian Kühnert (2006a). "Information flows in hierarchical networks and the capability of organizations to successfully respond to failures, crises, and disasters." In: *Physica A: Statistical Mechanics and its Applications* 363.1. Information and Material Flows in Complex Networks, pp. 141 –150. issn: 0378-4371. doi: 10.1016/j.physa.2006.01.055.

Helbing, Dirk, Thomas Seidel, Stefan Lämmer, and Karsten Peters (2006b). "Self-organization Principles in Supply Networks and Production Systems." In: *Econophysics and Sociophysics: Trends and Perspectives.* Wiley-VCH Verlag GmbH & Co. KGaA. Chap. 19, pp. 535–559. isbn: 978-3-527-40670-8. doi: 10.1002/9783527610006.ch19.

Heragu, S.S., R.J. Graves, Byung-In Kim, and A. St Onge (2002). "Intelligent agent based framework for manufacturing systems control." In: *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on* 32.5 (Sept. 2002), pp. 560–573. issn: 1083-4427. doi: 10.1109/TSMCA.2002.804788.

Herbon, Avi, Eugene Khmelnitsky, and Oded Maimon (2004). "Effective information horizon length in measuring off-line performance of stochastic dynamic systems." In: *European Journal of Operational Research* 157.3, pp. 688 –703. issn: 0377-2217. doi: 10.1016/S0377-2217(03)00268-6.

Hernando, Leticia, Alexander Mendiburu, and Jose A. Lozano (2016). "A Tunable Generator of Instances of Permutation-Based Combinatorial Optimization Problems." In: *IEEE Transactions on Evolutionary Computation* 20.2 (Apr. 2016), pp. 165–179. issn: 1089-778X. doi: 10.1109/TEVC.2015.2433680.

Herrmann, Frank (2011). *Operative Planung in IT-Systemen für die Produktionsplanung und -steuerung — Wirkung, Auswahl und Einstellhinweise von Verfahren und Parametern.* Vieweg+Teubner Verlag. isbn: 978-3-8348-1209-4. doi: 10.1007/978-3-8348-8172-4.

Herrmann, Jeffrey W. (2006). "A History of Production Scheduling." In: *Handbook of Production Scheduling.* Ed. by Jeffrey W. Herrmann. Boston, MA: Springer. Chap. 1, pp. 1–22. isbn: 978-0-387-33117-1. doi: 10.1007/0-387-33117-4_1.

Herroelen, Willy and Roel Leus (2005). "Project scheduling under uncertainty: Survey and research potentials." In: *European Journal of Operational Research* 165.2, pp. 289 –306. issn: 0377-2217. doi: 10.1016/j.ejor.2004.04.002.

Heusse, Martin, Dominique Snyers, Sylvain Guérin, and Pascale Kuntz (1998). "Adaptive Agent-Driven Routing and Load Balancing in Communication Networks." In: *Advances in Complex Systems* 1.02n03, pp. 237–254. doi: 10.1142/S0219525998000168.

Himarios, Daniel (2000). "How Forward Looking Are Consumers? Further Evidence for the United States." In: *Southern Economic Journal* 66.4, pp. 991–1000. issn: 0038-4038. doi: 10.2307/1061540. jstor: 1061540.

Hogg, T. and B. A. Huberman (1991). "Controlling chaos in distributed systems." In: *IEEE Transactions on Systems, Man, and Cybernetics* 21.6 (Nov. 1991), pp. 1325–1332. issn: 0018-9472. doi: 10.1109/21.135679.

Holland, John H. (1998). *Emergence: From Chaos to Order.* Cambridge, MA: Perseus Books. isbn: 0-201-14943-5.

– (2002). "Complex Adaptive Systems and Spontaneous Emergence." In: *Complexity and Industrial Clusters: Dynamics and Models in Theory and Practice.* Ed. by Alberto Quadrio Curzio and Marco Fortis. Heidelberg: Physica-Verlag, pp. 25–34. isbn: 978-3-642-50007-7. doi: 10.1007/978-3-642-50007-7_3.

Holvoet, Tom, Danny Weyns, and Paul Valckenaers (2009). "Patterns of Delegate MAS." In: *Self-Adaptive and Self-Organizing Systems, 2009. SASO '09. Third IEEE International Conference on*, pp. 1–9. doi: 10.1109/SASO.2009.31.

Holweg, Matthias (2007). "The genealogy of lean production." In: *Journal of Operations Management* 25.2. Special Issue Evolution of the Field of Operations Management SI/ Special Issue Organisation Theory and Supply Chain Management, pp. 420 –437. issn: 0272-6963. doi: 10.1016/j.jom.2006.04.001.

Holzman, Ron and Nissan Law-Yone (1997). "Strong Equilibrium in Congestion Games." In: *Games and Economic Behavior* 21.1, pp. 85 –101. issn: 0899-8256. doi: 10.1006/game.1997.0592.

Hooker, John N. (1995). "Testing heuristics: We have it all wrong." In: *Journal of Heuristics* 1.1, pp. 33–42. issn: 1572-9397. doi: 10.1007/BF02430364.

Hopp, Wallace J. and Mark L. Spearman (2004). "To Pull or Not to Pull: What Is the Question?" In: *Manufacturing & Service Operations Management* 6.2, pp. 133 –148. issn: 15234614. doi: 10.1287/msom.1030.0028.

Hopp, Wallace J and Mark L Spearman (2008). *Factory physics.* 3rd. McGraw-Hill/Irwin New York. isbn: 978-007-123246-3.

Horling, Bryan and Victor Lesser (2004). "A survey of multi-agent organizational paradigms." In: *Knowl. Eng. Rev.* 19.4 (Dec. 2004), pp. 281–316. issn: 0269-8889. doi: 10.1017/S0269888905000317. acm id: 1103864.

Hurink, Johann, Bernd Jurisch, and Monika Thole (1994). "Tabu search for the job-shop scheduling problem with multi-purpose machines." English. In: *Operations-Research-Spektrum* 15.4, pp. 205–215. issn: 0171-6468. doi: 10.1007/BF01719451.

Ibarra, Herminia (1993). "Network Centrality, Power, and Innovation Involvement: Determinants of Technical and Administrative Roles." In: *Academy of Management Journal* 36.3, pp. 471–501. doi: 10.2307/256589. jstor: 256589.

Igelmund, G. and F. J. Radermacher (1983). “Preselective strategies for the optimization of stochastic project networks under resource constraints.” In: *Networks* 13.1, pp. 1–28. issn: 1097-0037. doi: 10.1002/net.3230130102.

Im, Jin H. and Richard J. Schonberger (1988). “The Pull of Kanban.” In: *Production & Inventory Management Journal* 29.4, pp. 54 –58. issn: 0897-8336.

Innes, Judith E. and David E. Booher (1999). “Consensus Building and Complex Adaptive Systems.” In: *Journal of the American Planning Association* 65.4, pp. 412–423. doi: 10.1080/01944369908976071.

Isern, David, David Sánchez, and Antonio Moreno (2011). “Organizational structures supported by agent-oriented methodologies.” In: *Journal of Systems and Software* 84.2, pp. 169 –184. issn: 0164-1212. doi: 10.1016/j.jss.2010.09.005.

Jackson, James R. (1963). “Jobshop-Like Queueing Systems.” In: *Management Science* 10.1, pp. 131–142. doi: 10.1287/mnsc.10.1.131. jstor: 2627213.

James Jr., Harvey and Farhad Rassekh (2000). “Smith, Friedman, and self-interest in ethical society.” In: *Business Ethics Quarterly* 10.3, pp. 659 –674. issn: 1052-150X. doi: 10.2307/3857897.

Janssens, Maddy and Chris Steyaert (1999). “The world in two and a third way out? The concept of duality in organization theory and practice.” In: *Scandinavian Journal of Management* 15.2, pp. 121 –139. issn: 0956-5221. doi: 10.1016/S0956-5221(98)00010-4.

Jeken, O., N. Duffie, K. Windt, H. Blunck, A. Chehade, and H. Rekersbrink (2012). “Dynamics of autonomously acting products and work systems in production and assembly.” In: *CIRP Journal of Manufacturing Science and Technology* 5.4. Special issue from the 44^{th} CIRP Conference on Manufacturing Systems, pp. 267 –275. issn: 1755-5817. doi: 10.1016/j.cirpj.2012.09.012.

Jennings, Nicholas R. and Stefan Bussmann (2003). “Agent-based control systems: Why are they suited to engineering complex systems?” In: *IEEE Control Systems* 23.3, pp. 61–73. issn: 1066-033X. doi: 10.1109/MCS.2003.1200249.

Jennings, Nicholas R. and J.R. Campos (1997). “Towards a social level characterisation of socially responsible agents.” In: *Software Engineering. IEE Proceedings* 144.1 (Feb. 1997), pp. 11–25. issn: 1364-5080. doi: 10.1049/ip-sen:19971021.

Jennings, Nicholas R. and M. Wooldridge (1998). “Agent Technology: Foundations, Applications, and Markets.” In: ed. by Nicholas R. Jennings and Michael J. Wooldridge. Berlin, Heidelberg: Springer Berlin Heidelberg. Chap. Applications of Intelligent Agents, pp. 3–28. isbn: 978-3-662-03678-5. doi: 10.1007/978-3-662-03678-5_1.

Jennings, Nicholas R. and Michael Wooldridge (1995). “Applying Agent Technology.” In: *Applied Artificial Intelligence* 9.4, pp. 357–369. doi: 10.1080/08839519508945480.

Jennings, Nick R. (1993). “Commitments and conventions: The foundation of coordination in multi-agent systems.” In: *The Knowledge Engineering Review* 8.03 (Sept. 1993), pp. 223–250. issn: 1469-8005. doi: 10.1017/S0269888900000205.

Jensen, Tommy R. and Bjarne Toft (1995). *Graph Coloring Problems*. Series on Discrete Mathematics and Optimization. John Wiley & Sons. isbn: 978-0-471-02865-9.

Johnson, David S., Cecilia R. Aragon, Lyle A. McGeoch, and Catherine Schevon (1991). “Optimization by Simulated Annealing: An Experimental Investigation; Part II, Graph Coloring and Number Partitioning.” In: *Operations Research* 39.3, pp. 378–406. issn: 0030364X. doi: 10.1287/opre.39.3.378. jstor: 171393.

Jones, Albert T. and Charles R. McLean (1986). “A proposed hierarchical control model for automated manufacturing systems.” In: *Journal of Manufacturing Systems* 5.1, pp. 15 –25. issn: 0278-6125. doi: 10.1016/0278-6125(86)90064-6.

Jones, Albert T., Larry H. Reeker, and Abhijit V. Deshmukh (2002). "On information and performance of complex manufacturing systems." In: *Proceedings of the Manufacturing Complexity Network Conference*. Cambridge, Apr. 2002.

Jones, Albert and Abdol Saleh (1990). "A multi-level/multi-layer architecture for intelligent shopfloor control." In: *International Journal of Computer Integrated Manufacturing* 3.1, pp. 60–70. doi: 10.1080/09511929008944433.

Judd, Stephen, Michael Kearns, and Yevgeniy Vorobeychik (2010). "Behavioral dynamics and influence in networked coloring and consensus." In: *Proceedings of the National Academy of Sciences* 107.34, pp. 14978–14982. doi: 10.1073/pnas.1001280107. pubmed: 20696936.

Kagermann, Henning, Wolfgang Wahlster, and Johannes Helbig (2013). *Recommendations for implementing the strategic initiative INDUSTRIE 4.0 — Final report of the Industrie 4.0 Working Group*. Tech. rep. acatech – National Academy of Science and Engineering, Apr. 2013.

Kalenka, Susanne and Nicholas R. Jennings (1999). "Socially Responsible Decision Making by Autonomous Agents." In: *Cognition, Agency and Rationality: Proceedings of the Fifth International Colloquium on Cognitive Science*. Ed. by Kepa Korta, Ernest Sosa, and Xabier Arrazola. Dordrecht: Springer, pp. 135–149. isbn: 978-94-017-1070-1. doi: 10.1007/978-94-017-1070-1_9.

Karakostas, George and Stavros G. Kolliopoulos (2009). "Stackelberg Strategies for Selfish Routing in General Multicommodity Networks." English. In: *Algorithmica* 53.1, pp. 132–153. issn: 0178-4617. doi: 10.1007/s00453-007-9018-5.

Karimi, Hamid Reza, Neil A. Duffie, and Sergey Dashkovskiy (2010). "Local Capacity H_∞ Control for Production Networks of Autonomous Work Systems With Time-Varying Delays." In: *IEEE Transactions on Automation Science and Engineering* 7.4, pp. 849–857. issn: 1545-5955. doi: 10.1109/TASE.2010.2046735.

Karp, Richard M. (1972). "Reducibility among Combinatorial Problems." English. In: *Complexity of Computer Computations*. Ed. by Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger. The IBM Research Symposia Series. Springer, pp. 85–103. isbn: 978-1-4684-2003-6. doi: 10.1007/978-1-4684-2001-2_9.

Karuna, Hadeli, Paul Valckenaers, Bart Saint-Germain, Paul Verstraete, Constantin Bala Zamfirescu, and Hendrik Van Brussel (2005). "Emergent Forecasting Using a Stigmergy Approach in Manufacturing Coordination and Control." In: *Engineering Self-Organising Systems: Methodologies and Applications*. Ed. by Sven A. Brueckner, Giovanna Di Marzo Serugendo, Anthony Karageorgos, and Radhika Nagpal. Berlin, Heidelberg: Springer, pp. 210–226. isbn: 978-3-540-31901-6. doi: 10.1007/11494676_14.

Karuna, Hadeli, Paul Valckenaers, Paul Verstraete, Bart Saint Germain, and Hendrik Van Brussel (2006). "A Study of System Nervousness in Multi-agent Manufacturing Control System." English. In: *Engineering Self-Organising Systems*. Ed. by Sven A. Brueckner, Giovanna Di Marzo Serugendo, David Hales, and Franco Zambonelli. Vol. 3910. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 232–243. isbn: 978-3-540-33342-5. doi: 10.1007/11734697_18.

Kassabalidis, I., M. A. El-Sharkawi, R. J. Marks, P. Arabshahi, and A. A. Gray (2001). "Swarm intelligence for routing in communication networks." In: *Global Telecommunications Conference, 2001. GLOBECOM '01. IEEE*. Vol. 6, pp. 3613–3617. doi: 10.1109/GLOCOM.2001.966355.

Kearns, Michael, Siddharth Suri, and Nick Montfort (2006). "An Experimental Study of the Coloring Problem on Human Subject Networks." In: *Science* 313.5788, pp. 824–827. doi: 10.1126/science.1127207. pubmed: 16902134.

Kim, Chonggun and H. Kameda (1992). "An algorithm for optimal static load balancing in distributed computer systems." In: *IEEE Transactions on Computers* 41.3, pp. 381–384. issn: 0018-9340. doi: 10.1109/12.127455.

Kingman, J. F. C. (1961). “The single server queue in heavy traffic.” In: *Mathematical Proceedings of the Cambridge Philosophical Society* 57.04 (Oct. 1961), pp. 902–904. issn: 1469-8064. doi: 10.1017/S0305004100036094.

Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi (1983). “Optimization by Simulated Annealing.” In: *Science* 220.4598, pp. 671–680. issn: 0036-8075. doi: 10.1126/science.220.4598.671. pubmed: 17813860.

Kirst, Christoph, Marc Timme, and Demian Battaglia (2016). “Dynamic information routing in complex networks.” In: *Nature Communications* 7, p. 11061. doi: 10.1038/ncomms11061. pubmed: 27067257.

Kitano, Hiroaki (2004). “Biological robustness.” In: *Nature Reviews Genetics* 5.11, pp. 826–837. doi: 10.1038/nrg1471.

– (2007). “Towards a theory of biological robustness.” In: *Molecular Systems Biology* 3.1. issn: 1744-4292. doi: 10.1038/msb4100179.

Klein, Robert and Armin Scholl (2011). *Planung und Enscheidung*. 2nd Edition. Vahlens Handbücher. München: Verlag Franz Vahlen GmbH. isbn: 978-3-8006-3884-0.

Kleinrock, Leonard (1970). “Analytic and Simulation Methods in Computer Network Design.” In: *Proceedings of the May 5-7, 1970, Spring Joint Computer Conference*. AFIPS ’70 (Spring). Atlantic City, New Jersey: ACM, pp. 569–579. doi: 10.1145/1476936.1477022. acm id: 1477022.

– (1975). *Queueing Systems, Volume I: Theory*. New York: Wiley Interscience. isbn: 0-471-49110-1.

Klir, George J. (1991). *Facets of Systems Science*. Vol. 7.I. International Series on Systems Science and Engineering. New York, NY: Plenum Press. isbn: 0-306-43959-X.

Koch, Thorsten, Tobias Achterberg, Erling Andersen, Oliver Bastert, Timo Berthold, Robert E. Bixby, Emilie Danna, Gerald Gamrath, Ambros M. Gleixner, Stefan Heinz, Andrea Lodi, Hans Mittelmann, Ted Ralphs, Domenico Salvagnin, Daniel E. Steffy, and Kati Wolter (2011). “MIPLIB 2010.” In: *Mathematical Programming Computation* 3.2, pp. 103–163. issn: 1867-2957. doi: 10.1007/s12532-011-0025-9.

Koestler, Arthur (1970). “Beyond atomism and holism — the concept of the holon.” In: vol. 13. 2. The Johns Hopkins University Press, pp. 131–154.

Kolata, Gina (1990). “What if They Closed 42$^{\text{nd}}$ Street and Nobody Noticed?” In: *New York Times* (Dec. 25, 1990).

Kolisch, Rainer (2001). *Make-to-Order Assembly Management*. 1st ed. Springer Science + Business Media. isbn: 978-3-540-41098-0. doi: 10.1007/978-3-662-04514-5.

Koren, Yoram (2010). *The Global Manufacturing Revolution: Product-Process-Business Integration and Reconfigurable Systems*. Wiley Series in Systems Engineering and Management. Hoboken, NJ: John Wiley & Sons, Inc. isbn: 978-0-470-58377-7.

Korilis, Yannis A., Aurel A. Lazar, and Ariel Orda (1995). “Architecting noncooperative networks.” In: *Selected Areas in Communications, IEEE Journal on* 13.7 (Sept. 1995), pp. 1241–1251. issn: 0733-8716. doi: 10.1109/49.414643.

– (1997a). “Achieving Network Optima Using Stackelberg Routing Strategies.” In: *IEEE/ACM Trans. Netw.* 5.1 (Feb. 1997), pp. 161–173. issn: 1063-6692. doi: 10.1109/90.554730. acm id: 245234.

– (1997b). “Capacity allocation under noncooperative routing.” In: *Automatic Control, IEEE Transactions on* 42.3 (Mar. 1997), pp. 309–325. issn: 0018-9286. doi: 10.1109/9.557575.

– (1999). “Avoiding the Braess paradox in non-cooperative networks.” In: *J. Appl. Probab.* 36.1 (June 1999), pp. 211–222. doi: 10.1239/jap/1032374242.

Korst, Jan, Emile Aarts, Jan Karel Lenstra, and Jaap Wessels (1994). “Periodic assignment and graph colouring.” In: *Discrete Applied Mathematics* 51.3, pp. 291–305. issn: 0166-218X. doi: 10.1016/0166-218X(92)00036-L.

Koutsoupias, Elias and Christos Papadimitriou (1999). “Worst-Case Equilibria.” In: *STACS 99: 16th Annual Symposium on Theoretical Aspects of Computer Science Trier, Germany, March 4–6, 1999 Proceedings.* Ed. by Christoph Meinel and Sophie Tison. Berlin, Heidelberg: Springer, pp. 404–413. isbn: 978-3-540-49116-3. doi: 10.1007/3-540-49116-3_38.

– (2009). “Worst-case equilibria.” In: *Computer Science Review* 3.2, pp. 65 –69. issn: 1574-0137. doi: 10.1016/j.cosrev.2009.04.003.

Krajewski, Lee J., Barry E. King, Larry P. Ritzman, and Danny S. Wong (1987). “Kanban, MRP, and Shaping the Manufacturing Environment.” In: *Management Science* 33.1, pp. 39–57. doi: 10.1287/mnsc.33.1.39. jstor: 2631610.

Kubicek, Herbert (1977). “Heuristische Bezugsrahmen und heuristisch angelegte Forschungsdesigns als Elemente einer Konstruktionsstrategie empirischer Forschung.” In: *Empirische und handlungstheoretische Forschungskonzeptionen in der Betriebswirtschaftslehre / Bericht über die Tagung in Aachen, März 1976.* Ed. by Richard Köhler. 1st ed. Stuttgart: C. E. Poeschel Verlag, pp. 3 –36. isbn: 3-7910-0214-7.

Kuehnle, Hermann (2007). “A system of models contribution to production network (PN) theory.” In: *Journal of Intelligent Manufacturing* 18.5, pp. 543–551. issn: 1572-8145. doi: 10.1007/s10845-007-0062-4.

Kuhn, Axel and Sigrid Wenzel (2008). “Simulation logistischer Systeme.” In: *Handbuch Logistik.* Ed. by Dieter Arnold, Heinz Isermann, Axel Kuhn, Horst Tempelmeier, and Kai Fuhrmans. Berlin: Springer, pp. 73–94. isbn: 978-3-540-72929-7. doi: 10.1007/978-3-540-72929-7.

Kuhn, Fabian and Rogert Wattenhofer (2006). “On the Complexity of Distributed Graph Coloring.” In: *Proceedings of the Twenty-fifth Annual ACM Symposium on Principles of Distributed Computing.* PODC ’06. Denver, Colorado, USA: ACM, pp. 7–15. isbn: 1-59593-384-0. doi: 10.1145/1146381.1146387. acm id: 1146387.

Kumar, P. R. and T. I. Seidman (1990). “Dynamic instabilities and stabilization methods in distributed real-time scheduling of manufacturing systems.” In: *IEEE Transactions on Automatic Control* 35.3, pp. 289–298. issn: 0018-9286. doi: 10.1109/9.50339.

Kuramoto, Yoshiki (1984). *Chemical Oscillations, Waves, and Turbulence.* Vol. 19. Springer Series in Synergetics. Berlin, Heidelberg: Springer. isbn: 978-3-642-69691-6. doi: 10.1007/978-3-642-69689-3.

Kusiak, Andrew (1987). “The production equipment requirements problem.” In: *International Journal of Production Research* 25.3, p. 319. issn: 00207543. doi: 10.1080/00207548708919843.

Langton, Chris G. (1990). “Computation at the edge of chaos: Phase transitions and emergent computation.” In: *Physica D: Nonlinear Phenomena* 42.1, pp. 12 –37. issn: 0167-2789. doi: 10.1016/0167-2789(90)90064-V.

Lasi, Heiner, Peter Fettke, Hans-Georg Kemper, Thomas Feld, and Michael Hoffmann (2014). “Industry 4.0.” In: *Business & Information Systems Engineering* 6.4, pp. 239–242. issn: 1867-0202. doi: 10.1007/s12599-014-0334-4.

Laverty, Kevin J. (2004). “Managerial myopia or systemic short-termism? The importance of managerial systems in valuing the long term.” In: *Management Decision* 42.8, pp. 949–962. doi: 10.1108/00251740410555443.

Lazer, David and Allan Friedman (2007). “The Network Structure of Exploration and Exploitation.” In: *Administrative Science Quarterly* 52.4, pp. 667–694. doi: 10.2189/asqu.52.4.667. jstor: 20109947.

Lee, Edward A. (2008). “Cyber Physical Systems: Design Challenges.” In: *11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing (ISORC)*, pp. 363–369. doi: 10.1109/ISORC.2008.25.

Lee, Heungsoon Felix, Mandyam M. Srinivasan, and Candace Arai Yano (1991). “The optimal configuration and workload allocation problem in flexible manufacturing systems.” English. In:

International Journal of Flexible Manufacturing Systems 3.3-4, pp. 213–230. issn: 0920-6299. doi: 10.1007/BF00170207.

Lee, Heungsoon Felix, Mandyam M. Srinivasan, and Candace Arai Yano (2006). “A framework for capacity planning and machine configuration in flexible assembly systems.” English. In: *International Journal of Flexible Manufacturing Systems* 18.4, pp. 239–268. issn: 0920-6299. doi: 10.1007/s10696-006-9010-4.

Lee, Yong-Han, Soundar R. T. Kumara, and Kalyan Chatterjee (2003). “Multiagent based dynamic resource scheduling for distributed multiple projects using a market mechanism.” In: *Journal of Intelligent Manufacturing* 14.5, pp. 471–484. issn: 1572-8145. doi: 10.1023/A:1025753309346.

Leighton, Frank Thomson (1979). “A graph coloring algorithm for large scheduling problems.” In: *Journal of Research of the National Bureau of Standards* 84.6, pp. 489–506. doi: 10.6028/jres.084.024.

Leitão, Paulo (2009). “Agent-based distributed manufacturing control: A state-of-the-art survey.” In: *Engineering Applications of Artificial Intelligence* 22.7. Distributed Control of Production Systems, pp. 979 –991. issn: 0952-1976. doi: 10.1016/j.engappai.2008.09.005.

Leitão, Paulo and Francisco Restivo (2006). “ADACOR: A holonic architecture for agile and adaptive manufacturing control.” In: *Computers in Industry* 57.2, pp. 121 –130. issn: 0166-3615. doi: 10.1016/j.compind.2005.05.005.

– (2008). “A holonic approach to dynamic manufacturing scheduling.” In: *Robotics and Computer-Integrated Manufacturing* 24.5. BASYS ’06: Balanced Automation Systems 7th IFIP International Conference on Information Technology for Balanced Automation Systems in Manufacturing and Services, pp. 625 –634. issn: 0736-5845. doi: 10.1016/j.rcim.2007.09.005.

Leitão, Paulo, Armando W. Colombo, and Francisco Restivo (2003). “An Approach to the Formal Specification of Holonic Control Systems.” In: *Holonic and Multi-Agent Systems for Manufacturing: First International Conference on Industrial Applications of Holonic and Multi-Agent Systems, HoloMAS 2003, Prague, Czech Republic, September 1st-3rd, 2003.* Ed. by Vladimír Mařík, Duncan McFarlane, and Paul Valckenaers. Berlin, Heidelberg: Springer, pp. 59–70. isbn: 978-3-540-45185-3. doi: 10.1007/978-3-540-45185-3_6.

Leitão, Paulo, Joel Alves, and Ana I. Pereira (2010). “Solving Myopia in Real-time Decision-making using Petri nets Models’ Knowledge for Service-oriented Manufacturing Systems.” In: *IFAC Proceedings Volumes* 43.4. 10th IFAC Workshop on Intelligent Manufacturing Systems, pp. 144 –149. issn: 1474-6670. doi: 10.3182/20100701-2-PT-4011.00026.

Leitão, Paulo, Vladimír Mařík, and Pavel Vrba (2013). “Past, Present, and Future of Industrial Agent Applications.” In: *IEEE Transactions on Industrial Informatics* 9.4, pp. 2360–2372. issn: 1551-3203. doi: 10.1109/TII.2012.2222034.

Lesne, Annick (2006). “Complex Networks: from Graph Theory to Biology.” English. In: *Letters in Mathematical Physics* 78.3, pp. 235–262. issn: 0377-9017. doi: 10.1007/s11005-006-0123-1.

Lesser, Victor R. and Daniel D. Corkill (1981). “Functionally Accurate, Cooperative Distributed Systems.” In: *IEEE Transactions on Systems, Man, and Cybernetics* 11.1, pp. 81–96. issn: 0018-9472. doi: 10.1109/TSMC.1981.4308581.

Leung, Ying-Tat and Rajan Suri (1990). “Performance evaluation of discrete manufacturing systems.” In: *IEEE Control Systems Magazine* 10.4, pp. 77–86. issn: 0272-1708. doi: 10.1109/37.56281.

Levy, David L. (2000). “Applications and Limitations of Complexity Theory in Organization Theory and Strategy.” In: *Handbook of Strategic Management.* Ed. by Jack Rabin, Gerald L. Miller, and W. Bartley Hildreth. 2nd. New York, Basel: Marcel Dekker, Inc., pp. 67–87. isbn: 0-8247-0339-1.

Lewin, Roger, Teresa Parker, and Birute Regine (1998). "Complexity Theory and the Organization: Beyond the Metaphor." In: *Complexity* 3.4, pp. 36–40. issn: 1099-0526. doi: 10.1002/(SICI)1099-0526(199803/04)3:4<36::AID-CPLX7>3.0.CO;2-I.

Lichtenstein, Benyamin B. (2007). "A Matrix of Complexity for Leadership: Fourteen Disciplines of Complex Systems Leadership Theory." In: *Complex Systems Leadership Theory*. Ed. by James K. Hazy, Jeffrey A. Goldstein, and Benyamin B. Lichtenstein. Vol. 1. Exploring Organizational Complexity. Mansfield, MA: ISCE Publishing. Chap. 15, pp. 285–304. isbn: 978-0-9791688-6-4.

Lichtenstein, Benyamin B., Mary Uhl-Bien, Russ Marion, Anson Seers, James Douglas Orton, and Craig Schreiber (2006). "Complexity leadership theory: An interactive perspective on leading in complex adaptive systems." In: *Emergence: Complexity & Organization* 8.4, pp. 2 –12. issn: 1532-7000.

Lichtenstein, Benyamin B., Mary Uhl-Bien, Marion Russ, Anson Seers, James Douglas Orton, and Craig Schreiber (2007). "Complexity Leadership Theory: An Interactive Perspective on Leading Complex Adaptive Systems." In: *Complex Systems Leadership Theory*. Ed. by James K. Hazy, Jeffrey A. Goldstein, and Benyamin B. Lichtenstein. Vol. 1. Exploring Organizational Complexity. Mansfield, MA: ISCE Publishing. Chap. 7, pp. 129–142. isbn: 978-0-9791688-6-4.

Lim, M. K., Z. Zhang, and W. T. Goh (2009). "An iterative agent bidding mechanism for responsive manufacturing." In: *Engineering Applications of Artificial Intelligence* 22.7. Distributed Control of Production Systems, pp. 1068 –1079. issn: 0952-1976. doi: 10.1016/j.engappai.2008.12.003.

Lin, Henry, Tim Roughgarden, Éva Tardos, and Asher Walkover (2011). "Stronger Bounds on Braess's Paradox and the Maximum Latency of Selfish Routing." In: *SIAM Journal on Discrete Mathematics* 25.4, pp. 1667–1686. doi: 10.1137/090769600.

Lindau, Roger A. and Kenth R. Lumsden (1995). "Actions taken to prevent the propagation of disturbances in manufacturing systems." In: *International Journal of Production Economics* 41.1-3. Proceedings of the 12th International Conference on Production Research, pp. 241 –248. issn: 0925-5273. doi: 10.1016/0925-5273(94)00068-9.

Little, John D. C. (1961). "A Proof for the Queuing Formula: $L = \lambda W$." In: *Operations Research* 9.3, pp. 383–387. doi: 10.1287/opre.9.3.383. jstor: 167570.

Lizier, Joseph T., Mikhail Prokopenko, and Albert Y. Zomaya (2008). "Local information transfer as a spatiotemporal filter for complex systems." In: *Phys. Rev. E* 77.2, p. 026110. doi: 10.1103/PhysRevE.77.026110.

Lombardi, Michele and Michela Milano (2012). "A min-flow algorithm for Minimal Critical Set detection in Resource Constrained Project Scheduling." In: *Artificial Intelligence* 182, pp. 58 –67. issn: 0004-3702. doi: 10.1016/j.artint.2011.12.001.

Lovejoy, William S. (1992). "Stopped Myopic Policies in Some Inventory Models with Generalized Demand Processes." In: *Management Science* 38.5, pp. 688–707. issn: 0025-1909, 1526-5501. jstor: 2632685.

Malone, Thomas W. and Kevin Crowston (1994). "The Interdisciplinary Study of Coordination." In: *ACM Comput. Surv.* 26.1 (Mar. 1994), pp. 87–119. issn: 0360-0300. doi: 10.1145/174666.174668. acm id: 174668.

Malone, Thomas W. and Stephen A. Smith (1988). "Modeling the Performance of Organizational Structures." In: *Operations Research* 36.3, pp. 421–436. doi: 10.1287/opre.36.3.421. jstor: 170985.

March, James G. (1991). "Exploration and Exploitation in Organizational Learning." In: *Organization Science* 2.1, pp. 71 –87. issn: 1047-7039. doi: 10.1287/orsc.2.1.71. jstor: 2634940.

– (1999). *The pursuit of organizational intelligence: Decisions and learning in organizations*. Cambridge, MA: Blackwell Publishers, Inc. isbn: 0-631-21101-2.

Marcotte, Patrice and Michael Patriksson (2007). "Traffic Equilibrium." In: *Transportation*. Ed. by Cynthia Barnhart and Gilbert Laporte. Vol. 14. Handbooks in Operations Research and Management Science. Elsevier. Chap. 10, pp. 623 –713. doi: 10.1016/S0927-0507(06)14010-4.

Marden, Jason R. (2016). "The Role of Information in Distributed Resource Allocation." In: *IEEE Transactions on Control of Network Systems*. Under Publication. issn: 2325-5870. doi: 10.1109/TCNS.2016.2553363.

Marden, Jason R. and Jeff S. Shamma (2015). "Game Theory and Distributed Control." In: *Handbook of Game Theory with Economic Applications*. Ed. by H. Peyton Young and Shmuel Zamir. Vol. 4. Elsevier. Chap. 16, pp. 861 –899. doi: 10.1016/B978-0-444-53766-9.00016-1.

Marginson, David and Laurie McAulay (2008). "Exploring the debate on short-termism: a theoretical and empirical analysis." In: *Strategic Management Journal* 29.3, pp. 273–292. issn: 1097-0266. doi: 10.1002/smj.657.

Mařík, Vladimír and Jiří Lažanský (2007). "Industrial applications of agent technologies." In: *Control Engineering Practice* 15.11. Special Issue on Manufacturing Plant Control: Challenges and Issues INCOM 2004. 11th IFAC INCOM'04 Symposium on Information Control Problems in Manufacturing, pp. 1364 –1380. issn: 0967-0661. doi: 10.1016/j.conengprac.2006.10.001.

Mařík, Vladimír and Duncan McFarlane (2005). "Industrial adoption of agent-based technologies." In: *Intelligent Systems, IEEE* 20.1, pp. 27–35. issn: 1541-1672. doi: 10.1109/MIS.2005.11.

Marion, Russ and Mary Uhl-Bien (2001). "Leadership in complex organizations." In: *The Leadership Quarterly* 12.4, pp. 389 –418. issn: 1048-9843. doi: 10.1016/S1048-9843(01)00092-3.

Marr, Carsten and Marc-Thorsten Hütt (2009). "Outer-totalistic cellular automata on graphs." In: *Physics Letters A* 373.5, pp. 546–549. issn: 0375-9601. doi: 10.1016/j.physleta.2008.12.013.

Marr, Carsten, Mark Müller-Linow, and Marc-Thorsten Hütt (2007). "Regularizing capacity of metabolic networks." In: *Phys. Rev. E* 75.4, p. 041917. doi: 10.1103/PhysRevE.75.041917.

Marx, Dániel (2004). "Graph Coloring Problems and their Applications in Scheduling." In: *Periodica Polytechnica Electrial Engineering* 48.1-2, pp. 11–16.

Marzo Serugendo, Giovanna, Noria Foukia, Salima Hassas, Anthony Karageorgos, Soraya Kouadri Mostéfaoui, Omer F. Rana, Mihaela Ulieru, Paul Valckenaers, and Chris Aart (2004). "Self-Organisation: Paradigms and Applications." In: *Engineering Self-Organising Systems: Nature-Inspired Approaches to Software Engineering*. Ed. by Giovanna Marzo Serugendo, Anthony Karageorgos, Omer F. Rana, and Franco Zambonelli. Berlin, Heidelberg: Springer, pp. 1–19. isbn: 978-3-540-24701-2. doi: 10.1007/978-3-540-24701-2_1.

Mason, Winter and Duncan J. Watts (2012). "Collaborative learning in networks." In: *Proceedings of the National Academy of Sciences* 109.3, pp. 764–769. doi: 10.1073/pnas.1110069108. pubmed: 22184216.

Mataric, Maja J. (1992). "Minimizing complexity in controlling a mobile robot population." In: *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*, 830–835 vol.1. doi: 10.1109/ROBOT.1992.220192.

Matni, Nikolai, Ao Tang, and John C. Doyle (2015). "A Case Study in Network Architecture Tradeoffs." In: *Proceedings of the 1st ACM SIGCOMM Symposium on Software Defined Networking Research*. SOSR '15. Santa Clara, California: ACM, 18:1–18:7. isbn: 978-1-4503-3451-8. doi: 10.1145/2774993.2775011. acm id: 2775011.

Maturana, Francisco P. and Douglas H. Norrie (1996). "Multi-agent Mediator architecture for distributed manufacturing." In: *Journal of Intelligent Manufacturing* 7.4, pp. 257–270. issn: 1572-8145. doi: 10.1007/BF00124828.

Mayo, Elton (1933). *The Human Problems of an Industrial Civilization*. New York, NY: Macmillan Company.

McCubbins, Mathew D., Ramamohan Paturi, and Nicholas Weller (2009). “Connected Coordination: Network Structure and Group Coordination.” In: *American Politics Research* 37.5, pp. 899–920. doi: 10.1177/1532673X09337184.

McDonnell, Patrick, Glen Smith, Sanjay Joshi, and Soundar R.T. Kumara (1999). “A Cascading Auction Protocol as a Framework for Integrating Process Planning and Heterarchical Shop Floor Control.” In: *International Journal of Flexible Manufacturing Systems* 11.1, pp. 37–62. issn: 1572-9370. doi: 10.1023/A:1008092606526.

McFarlane, Duncan C. and Stefan Bussmann (2000). “Developments in holonic production planning and control.” In: *Production Planning & Control* 11.6, pp. 522–536. doi: 10.1080/095372800414089.

– (2003). “Holonic Manufacturing Control: Rationales, Developments and Open Issues.” In: *Agent-Based Manufacturing: Advances in the Holonic Approach.* Ed. by S. M. Deen. Berlin, Heidelberg: Springer Berlin Heidelberg. Chap. 13, pp. 303–326. isbn: 978-3-662-05624-0. doi: 10.1007/978-3-662-05624-0_13.

McGeoch, Catherine C. (1996). “Toward an Experimental Method for Algorithm Simulation.” In: *INFORMS Journal on Computing* 8.1, pp. 1–15. doi: 10.1287/ijoc.8.1.1.

McKay, Kenneth N. (2011). “The Historical Foundations of Manufacturing Planning and Control Systems.” English. In: *Planning Production and Inventories in the Extended Enterprise.* Ed. by Karl G. Kempf, Pinar Keskinocak, and Reha Uzsoy. Vol. 151. International Series in Operations Research & Management Science. New York, Dordrecht, Heidelberg, London: Springer, pp. 21–31. isbn: 978-1-4419-6484-7. doi: 10.1007/978-1-4419-6485-4_2.

McKelvey, Bill (1999). “Complexity Theory in Organization Science: Seizing the Promise or Becoming a Fad?” In: *Emergence* 1.1, pp. 5–32. doi: 10.1207/s15327000em0101_2.

– (2004). “Complexity science as order-creation science: New theory, new method.” In: *Emergence: Complexity & Organization* 6.4, pp. 2 –27. issn: 1532-7000.

McKinsey & Company (2015). *Industry 4.0 — How to navigate digitization of the manufacturing sector.* McKinsey & Company.

McMillan, Elizabeth (2008). *Complexity, Management and the Dynamics of Change: Challenges for Practice.* London, New York: Routledge. isbn: 978-0-415-41721-1.

McNally, Michael G. (2000). “The Four Step Model.” In: *Handbook of Transport Modelling.* Ed. by David A. Hensher and Kenneth J. Button. 3rd. Pergamon. Chap. 3, pp. 35–52.

Merriam-Webster, ed. (2016). *Myopia.* Accessed: 2017-04-11. Springfield, MA: Merriam-Webster.com.

Mesarovic, M. D., D. Macko, and Y. Takahara (1970). “Two coordination principles and their application in large scale systems control.” In: *Automatica* 6.2, pp. 261 –270. issn: 0005-1098. doi: 10.1016/0005-1098(70)90097-X.

Meyer, Gerben G., Kary Främling, and Jan Holmström (2009). “Intelligent Products: A survey.” In: *Computers in Industry* 60.3, pp. 137 –148. issn: 0166-3615. doi: 10.1016/j.compind.2008.12.005.

Meyer, Mirja (2016). “Redundancy Investments in Manufacturing Systems - The role of redundancies for manufacturing system robustness.” PhD thesis. Jacobs University Bremen. urn: urn:nbn:de:gbv:579-opus-1006482.

Meyers, Carol A. and Andreas S. Schulz (2012). “The complexity of welfare maximization in congestion games.” In: *Networks* 59.2, pp. 252–260. issn: 1097-0037. doi: 10.1002/net.20439.

Miller, David M. and Robert P. Davis (1977). “The machine requirements problem.” In: *International Journal of Production Research* 15.2, pp. 219–231. doi: 10.1080/00207547708943119.

Miller, Tan C. (2002). *Hierarchical Operations and Supply Chain Planning.* 1st. London: Springer-Verlag. isbn: 978-1-85233-592-2. doi: 10.1007/978-1-4471-0305-9.

Minton, Steven, Mark D. Johnston, Andrew B. Philips, and Philip Laird (1992). “Minimizing conflicts: a heuristic repair method for constraint satisfaction and scheduling problems.” In:

Artificial Intelligence 58.1, pp. 161 –205. issn: 0004-3702. doi: 10.1016/0004-3702(92)90007-K.

Möhring, Rolf H. and Frederik Stork (2000). "Linear preselective policies for stochastic project scheduling." English. In: *Mathematical Methods of Operations Research* 52.3, pp. 501–515. issn: 1432-2994. doi: 10.1007/s001860000095.

Moldaschl, Manfred (1998). "Rationality, Culture and Politics of Production." English. In: *Shop Floor Control — A Systems Perspective: From Deterministic Models towards Agile Operations Management.* Ed. by Eric Scherer. Berlin, Heidelberg: Springer, pp. 67–90. isbn: 978-3-642-64349-1. doi: 10.1007/978-3-642-60313-6_3.

Moldaschl, Manfred and Wolfgang G. Weber (1998). "The "Three Waves" of Industrial Group Work: Historical Reflections on Current Research on Group Work." In: *Human Relations* 51.3, pp. 347–388. doi: 10.1177/001872679805100307.

Mönch, Lars (2005). *Agentenbasierte Produktionssteuerung komplexer Produktionssysteme.* 1st. Wiesbaden: Deutscher Universitäts Verlag. isbn: 3-8350-0249-X. doi: 10.1007/978-3-8350-9079-8.

Mönch, Lars and René Driessel (2005). "A distributed shifting bottleneck heuristic for complex job shops." In: *Computers & Industrial Engineering* 49.3, pp. 363 –380. issn: 0360-8352. doi: 10.1016/j.cie.2005.06.004.

Monderer, Dov and Lloyd S. Shapley (1996). "Potential Games." In: *Games and Economic Behavior* 14.1, pp. 124 –143. issn: 0899-8256. doi: 10.1006/game.1996.0044.

Monostori, L., J. Váncza, and S. R. T. Kumara (2006). "Agent-Based Systems for Manufacturing." In: *CIRP Annals - Manufacturing Technology* 55.2, pp. 697–720. issn: 0007-8506. doi: 10.1016/j.cirp.2006.10.004.

Monostori, L., B. Cs. Csáji, B. Kádár, A. Pfeiffer, E. Ilie-Zudor, Zs. Kemény, and M. Szathmári (2010). "Towards adaptive and digital manufacturing." In: *Annual Reviews in Control* 34.1, pp. 118 –128. issn: 1367-5788. doi: 10.1016/j.arcontrol.2010.02.007.

Monostori, László, Paul Valckenaers, Alexandre Dolgui, Hervé Panetto, Mietek Brdys, and Balázs Csanád Csáji (2015). "Cooperative control in production and logistics." In: *Annual Reviews in Control* 39, pp. 12 –29. issn: 1367-5788. doi: 10.1016/j.arcontrol.2015.03.001.

Montreuil, Benoit (2011). "Toward a Physical Internet: meeting the global logistics sustainability grand challenge." English. In: *Logistics Research* 3.2-3, pp. 71–87. issn: 1865-035X. doi: 10.1007/s12159-011-0045-x.

Montreuil, Benoit, Russell D. Meller, and Eric Ballot (2013). "Physical Internet Foundations." English. In: *Service Orientation in Holonic and Multi Agent Manufacturing and Robotics.* Ed. by Theodor Borangiu, Andre Thomas, and Damien Trentesaux. Vol. 472. Studies in Computational Intelligence. Berlin, Heidelberg: Springer, pp. 151–166. isbn: 978-3-642-35851-7. doi: 10.1007/978-3-642-35852-4_10.

Moreira, André A., Abhishek Mathur, Daniel Diermeier, and Luís A. N. Amaral (2004). "Efficient system-wide coordination in noisy environments." In: *Proceedings of the National Academy of Sciences of the United States of America* 101.33, pp. 12085–12090. doi: 10.1073/pnas.0400672101. pubmed: 15297617.

Moretti, Paolo and Miguel A. Muñoz (2013). "Griffiths phases and the stretching of criticality in brain networks." In: *Nature Communications* 4. doi: 10.1038/ncomms3521.

Morgan, Gareth (2006). *Images of Organization.* Updated Edition. Thousand Oaks, London, New Delhi: Sage Publications. isbn: 978-1-4129-3979-9.

Moulin, Bernard and Brahim Chaib-Draa (1996). "An overview of distributed artificial intelligence." In: *Foundations of distributed artificial intelligence.* Ed. by G. M. P. O'Hare and Nicholas R. Jennings. Vol. 1. New York: Wiley & Sons. Chap. 1, pp. 3–55. isbn: 978-0-471-00675-6.

Mula, J., R. Poler, J.P. García-Sabater, and F.C. Lario (2006). "Models for production planning under uncertainty: A review." In: *International Journal of Production Economics* 103.1, pp. 271 –285. issn: 0925-5273. doi: 10.1016/j.ijpe.2005.09.001.

Müller-Linow, Mark, Claus C. Hilgetag, and Marc-Thorsten Hütt (2008). "Organization of Excitable Dynamics in Hierarchical Biological Networks." In: *PLoS Comput Biol* 4.9 (Sept. 2008), e1000190. doi: 10.1371/journal.pcbi.1000190.

Müller, H. Jürgen (1997). "Distributed expertise Towards agent systems engineering." In: *Data & Knowledge Engineering* 23.3, pp. 217 –245. issn: 0169-023X. doi: 10.1016/S0169-023X(97)00013-X.

Myszkowski, Pawel B. (2008). "Solving Scheduling Problems by Evolutionary Algorithms for Graph Coloring Problem." In: *Metaheuristics for Scheduling in Industrial and Manufacturing Applications*. Ed. by Fatos Xhafa and Ajith Abraham. Vol. 128. Studies in Computational Intelligence. Berlin, Heidelberg: Springer, pp. 145–167. isbn: 978-3-540-78984-0. doi: 10.1007/978-3-540-78985-7_7.

Namatame, Akira and Takanori Sasaki (1998). "Self-organization of complex adaptive systems as a society of rational agents." In: *Artificial Life and Robotics* 2.4, pp. 189–195. issn: 1614-7456. doi: 10.1007/BF02471180.

National Research Council (1998). *Visionary Manufacturing Challenges for 2020*. Washington, D.C.: National Academy Press. isbn: 0-309-52290-0.

National Science Foundation (1953). *The third annual report of the National Science Foundation*. Annual Report. Washington, D.C.: National Science Foundation.

Negri da Silva, Claudio Rogerio and Reinaldo Morabito (2009). "Performance evaluation and capacity planning in a metallurgical job-shop system using open queueing network models." In: *International Journal of Production Research* 47.23, pp. 6589–6609. doi: 10.1080/00207540802350732.

von Neumann, John (1963). "The General and Logical Theory of Automata." In: *Design of Computers, Theory of Automata and Numerical Analysis*. Ed. by John von Neumann and Abraham Haskel Taub. Vol. V. John von Neumann Collected Works. Oxford: Pergamon Press, pp. 288–326.

Niazi, Muaz and Amir Hussain (2011). "Agent-based computing from multi-agent systems to agent-based models: a visual survey." In: *Scientometrics* 89.2, pp. 479–499. issn: 1588-2861. doi: 10.1007/s11192-011-0468-9.

Nilsson, Fredrik and Vince Darley (2006). "On complex adaptive systems and agent-based modelling for improving decision-making in manufacturing and logistics settings: Experiences from a packaging company." In: *International Journal of Operations & Production Management* 26.12, pp. 1351–1373. doi: 10.1108/01443570610710588.

Nisan, Noam and Amir Ronen (1999). "Algorithmic Mechanism Design (Extended Abstract)." In: *Proceedings of the Thirty-first Annual ACM Symposium on Theory of Computing, STOC '99, Atlanta, GA*. New York, NY: ACM, pp. 129–140. isbn: 1-58113-067-8. doi: 10.1145/301250.301287. acm id: 301287.

Nisan, Noam, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, eds. (2007). *Algorithmic Game Theory*. New York, NY: Cambridge University Press. isbn: 978-0-521-87282-9.

Nof, S.Y., G. Morel, L. Monostori, A. Molina, and F. Filip (2006). "From plant and logistics control to multi-enterprise collaboration." In: *Annual Reviews in Control* 30.1. 2005 IFAC Milestone Reports, pp. 55 –68. issn: 1367-5788. doi: 10.1016/j.arcontrol.2006.01.005.

Nwana, H. S., L. Lee, and N. R. Jennings (1997). "Co-ordination in multi-agent systems." In: *Software Agents and Soft Computing Towards Enhancing Machine Intelligence: Concepts and Applications*. Ed. by Hyacinth S. Nwana and Nader Azarmi. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 42–58. isbn: 978-3-540-68079-6. doi: 10.1007/3-540-62560-7_37.

Nyhuis, Peter and Hans-Peter Wiendahl (2009). *Fundamentals of Production Logistics.* Berlin, Heidelberg: Springer. isbn: 978-3-540-34210-6. doi: 10.1007/978-3-540-34211-3.

Okubo, Hiroki, Weng Jiahua, and Hisashi Onari (2000). "Characteristics of distributed autonomous production control." In: *International Journal of Production Research* 38.17, pp. 4205–4215. doi: 10.1080/00207540050205028.

Olhager, Jan and Björn Östlund (1990). "OR for Engineers Expert Systems and Decision-Aid An integrated push-pull manufacturing strategy." In: *European Journal of Operational Research* 45.2, pp. 135 –142. issn: 0377-2217. doi: 10.1016/0377-2217(90)90180-J.

Olhager, Jan and Joakim Wikner (2000). "Production planning and control tools." In: *Production Planning & Control* 11.3, pp. 210–222. doi: 10.1080/095372800232180.

Oliveira, Gina M.B. and Paulo M. Vidica (2012). "A Coevolutionary Approach to Cellular Automata-Based Task Scheduling." English. In: *Cellular Automata.* Ed. by Georgios Ch. Sirakoulis and Stefania Bandini. Vol. 7495. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 111–120. isbn: 978-3-642-33349-1. doi: 10.1007/978-3-642-33350-7_12.

Orlicky, Joseph (1975). *MRP: Material Requirements Planning.* 1st. Mc Graw Hill. isbn: 0-07-047708-6.

Ossowski, Sascha (1999). *Co-ordination in Artificial Agent Societies.* Lecture Notes in Computer Science 1535. Berlin, Heidelberg: Springer. isbn: 978-3-540-65495-7. doi: 10.1007/3-540-49212-7.

Ossowski, Sascha and Ana García-Serrano (1999). "Social Structure in Artificial Agent Societies: Implications for Autonomous Problem-Solving Agents." English. In: *Intelligent Agents V: Agents Theories, Architectures, and Languages.* Ed. by Jörg P. Müller, Anand S. Rao, and Munindar P. Singh. Vol. 1555. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 133–148. isbn: 978-3-540-65713-2. doi: 10.1007/3-540-49057-4_9.

Ossowski, Sascha and Andrea Omicini (2002). "Coordination knowledge engineering." In: *The Knowledge Engineering Review* 17.4, pp. 309–316. doi: 10.1017/S0269888903000596.

Ouelhadj, Djamila and Sanja Petrovic (2009). "A survey of dynamic scheduling in manufacturing systems." In: *Journal of Scheduling* 12.4, pp. 417 –431. issn: 10946136. doi: 10.1007/s10951-008-0090-8.

Pach, Cyrille, Thierry Berger, Thérèse Bonte, and Damien Trentesaux (2014). "ORCA-FMS: a dynamic architecture for the optimized and reactive control of flexible manufacturing scheduling." In: *Computers in Industry* 65.4, pp. 706 –720. issn: 0166-3615. doi: 10.1016/j.compind.2014.02.005.

Packard, Norman H. and Stephen Wolfram (1985). "Two-dimensional cellular automata." In: *Journal of Statistical Physics* 38.5, pp. 901–946. issn: 1572-9613. doi: 10.1007/BF01010423.

Panwalkar, S. S. and Wafik Iskander (1977). "A Survey of Scheduling Rules." In: *Operations Research* 25.1, pp. 45–61. doi: 10.1287/opre.25.1.45. jstor: 169546.

Papadimitriou, Christos (2001). "Algorithms, Games, and the Internet." In: *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing.* STOC '01. Hersonissos, Greece: ACM, pp. 749–753. isbn: 1-58113-349-9. doi: 10.1145/380752.380883. acm id: 380883.

Papadimitriou, Christos and Gregory Valiant (2010). "A New Look at Selfish Routing." In: *Innovations in Computer Science (ICS2010).* Ed. by Andrew Chi-Chih Yao. Tsinghua University Press, pp. 178–187. isbn: 978-7-302-21752-7.

Papadopolous, H. T., Cathal Heavey, and Jim Browne (1993). *Queueing Theory in Manufacturing Systems Analysis and Design.* 1st ed. London: Chapman & Hall. isbn: 0-412-38720-4.

Papadopoulos, H. T. and Cathal Heavey (1996). "Queueing theory in manufacturing systems analysis and design: A classification of models for production and transfer lines." In: *European Journal of Operational Research* 92.1, pp. 1 –27. issn: 0377-2217. doi: 10.1016/0377-2217(95)00378-9.

Paslack, Rainer (1991). *Urgeschichte der Selbstorganisation: Zur Archäologie eines wissenschaftlichen Paradigmas*. Vol. 32. Wissenschaftstheorie, Wissenschaft und Philosophie. Braunschweig, Wiesbaden: Vieweg+Teubner. isbn: 978-3-528-06423-5. doi: 10.1007/978-3-322-88776-4.

Passino, Kevin M (2005). *Biomimicry for optimization, control, and automation*. London, New York: Springer. isbn: 9781852338046. doi: 10.1007/b138169.

Paul, Valckenaers, Van Brussel Hendrik, Bongaerts L., and Wyns J. (1997). "Holonic Manufacturing Systems." In: *Integrated Computer-Aided Engineering* 4.3, pp. 191–201. issn: 1069-2509. doi: 10.3233/ICA-1997-4304.

Pearl, Judea (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Ed. by Ronald J. Brachman. The Morgan Kaufmann Series in Representation and Reasoning. San Mateo, CA: Morgan Kaufmann. isbn: 0-934613-73-7.

Peeters, Patrick, Hendrik Van Brussel, Paul Valckenaers, Jo Wyns, Luc Bongaerts, Martin Kollingbaum, and Tapio Heikkilä (2001). "Pheromone based emergent shop floor control system for flexible flow shops." In: *Artificial Intelligence in Engineering* 15.4. Methodology of Emergent Synthesis, pp. 343 –352. issn: 0954-1810. doi: 10.1016/S0954-1810(01)00026-7.

Perez-Gonzalez, Paz and Jose M. Framinan (2015). "Assessing scheduling policies in a permutation flowshop with common due dates." In: *International Journal of Production Research* 53.19, pp. 5742–5754. doi: 10.1080/00207543.2014.994077.

Pettersen, Jan-Arne and Anders Segerstedt (2009). "Restricted work-in-process: A study of differences between Kanban and CONWIP." In: *International Journal of Production Economics* 118.1. Special section on problems and models of inventories. Selected papers of the 14^{th} International symposium on inventories, pp. 199 –207. issn: 0925-5273. doi: 10.1016/j.ijpe.2008.08.043.

Pettigrew, Andrew M. and Evelyn M. Fenton (2000). "Complexities and dualities in innovative forms of organizing." In: *The Innovating Organization*. Ed. by Andrew M. Pettigrew and Evelyn M. Fenton. London: Sage Publications Ltd. Chap. 9, pp. 279–300. isbn: 0-7619-6433-9. doi: 10.4135/9781446219379.n10.

Philipp, T., F. Böse, and K. Windt (2006). "Evaluation of Autonomously Controlled Logistic Processes." In: *Proceedings of 5^{th} CIRP International Seminar on Intelligent Computation in Manufacturing Engineering*. Ed. by Roberto Teti. Ischia, Italy: University of Naples, pp. 347–352.

Philipp, Thorsten, Christoph Beer, Katja Windt, and Bernd Scholz-Reiter (2007). "Evaluation of Autonomous Logistic Processes – Analysis of the Influence of Structural Complexity." English. In: *Understanding Autonomous Cooperation and Control in Logistics*. Ed. by Michael Hülsmann and Katja Windt. Berlin, Heidelberg: Springer, pp. 303–324. isbn: 978-3-540-47449-4. doi: 10.1007/978-3-540-47450-0_20.

Pickardt, Christoph W., Torsten Hildebrandt, Jürgen Branke, Jens Heger, and Bernd Scholz-Reiter (2013). "Evolutionary generation of dispatching rule sets for complex dynamic scheduling problems." In: *International Journal of Production Economics* 145.1, pp. 67 –77. issn: 0925-5273. doi: 10.1016/j.ijpe.2012.10.016.

Pinedo, Michael L. (2008). *Scheduling*. 3rd. New York, NY: Springer. isbn: 978-0-387-78934-7. doi: 10.1007/978-0-387-78935-4.

Plattform Industrie 4.0 (2015). *Industrie 4.0 — Whitepaper FuE-Themen*. Accessed: 2017-04-11. Apr. 2015. url: http://www.din.de/blob/67744/de1c706b159a6f1baceb95a6677ba497/whitepaper-fue-themen-data.pdf.

Plowman, Donde Ashmos, Stephanie Solansky, Tammy E. Beck, LaKami Baker, Mukta Kulkarni, and Deandra Villarreal Travis (2007). "The role of leadership in emergent, self-organization." In: *The Leadership Quarterly* 18.4. Leadership and Complexity, pp. 341 –356. issn: 1048-9843. doi: 10.1016/j.leaqua.2007.04.004.

Pochet, Yves and Laurence A Wolsey (2006). *Production Planning by Mixed Integer Programming*. Springer Series in Operations Research and Financial Engineering. New York: Springer. isbn: 978-0-387-29959-4. doi: 10.1007/0-387-33477-7.

Policella, Nicola, Amedeo Cesta, Angelo Oddi, and Stephen F. Smith (2007). "From precedence constraint posting to partial order schedules: A CSP approach to Robust Scheduling." In: *AI Communications* 20.3, pp. 163 –180. issn: 09217126.

Powell, Warren B. and Yosef Sheffi (1982). "The Convergence of Equilibrium Algorithms with Predetermined Step Sizes." In: *Transportation Science* 16.1, pp. 45–55. doi: 10.1287/trsc.16.1.45.

Prigogine, Ilya and Isabelle Stengers (1984). *Order out of chaos: Man's new dialogue with nature*. Toronto: Bantam Books. isbn: 0-553-34082-4.

Puchkova, Alena, Julien Le Romancer, and Duncan McFarlane (2016). "Balancing Push and Pull Strategies within the Production System." In: *IFAC-PapersOnLine* 49.2. 7^{th} IFAC Conference on Management and Control of Production and Logistics (MCPL) 2016, Bremen, Germany, February 22^{nd}-24^{th} 2016, pp. 66 –71. issn: 2405-8963. doi: 10.1016/j.ifacol.2016.03.012.

Queiroz, Marcelo and Carlos Humes Jr. (2003). "A heuristic for the continuous capacity and flow assignment." In: *European Journal of Operational Research* 146.3, pp. 444 –459. issn: 0377-2217. doi: 10.1016/S0377-2217(02)00219-9.

Quijano, Nicanor and Kevin M. Passino (2010). "Honey bee social foraging algorithms for resource allocation: Theory and application." In: *Engineering Applications of Artificial Intelligence* 23.6, pp. 845 –861. issn: 0952-1976. doi: 10.1016/j.engappai.2010.05.004.

Rahimifard, S. (2004). "Semi-heterarchical production planning structures in the support of team-based manufacturing." In: *International Journal of Production Research* 42.17, pp. 3369 –3382. issn: 0020-7543. doi: 10.1080/00207540410001717128.

Raith, Andrea, Judith Y. T. Wang, Matthias Ehrgott, and Stuart A. Mitchell (2014). "Solving multi-objective traffic assignment." English. In: *Annals of Operations Research* 222.1, pp. 483–516. issn: 0254-5330. doi: 10.1007/s10479-012-1284-1.

Rajkumar, Ragunathan, Insup Lee, Lui Sha, and John Stankovic (2010). "Cyber-physical Systems: The Next Computing Revolution." In: *Proceedings of the 47^{th} Design Automation Conference, DAC' 10, Anaheim California*. New York, NY, USA: ACM, pp. 731–736. isbn: 978-1-4503-0002-5. doi: 10.1145/1837274.1837461. acm id: 1837461.

Ravasz, E., A. L. Somera, D. A. Mongru, Z. N. Oltvai, and A.-L. Barabási (2002). "Hierarchical Organization of Modularity in Metabolic Networks." In: *Science* 297.5586, pp. 1551–1555. doi: 10.1126/science.1073374. pubmed: 12202830.

Ravasz, Erzsébet and Albert-László Barabási (2003). "Hierarchical organization in complex networks." In: *Phys. Rev. E* 67.2 (Feb. 2003), p. 026112. doi: 10.1103/PhysRevE.67.026112.

Reveliotis, Spyros A. (2005). *Real-Time Management of Resource Allocations Systems*. Vol. 79. International Series in Operations Research & Management Science. Boston: Kluwer Academic Publishers. isbn: 0-387-23960-X. doi: 10.1007/b104057.

Richardson, Kurt A. (2008). "Managing Complex Organizations: Complexity Thinking and the Science and Art of Management." In: *Emergence: Complexity & Organization* 10.2, pp. 13 –26. issn: 1532-7000.

Ringhofer, Christian (2012). "Traffic Flow Models and Service Rules for Complex Production Systems." English. In: *Decision Policies for Production Networks*. Ed. by Dieter Armbruster and Karl G. Kempf. London: Springer, pp. 209–233. isbn: 978-0-85729-643-6. doi: 10.1007/978-0-85729-644-3_9.

Roethlisberger, Fritz J. and William Dickson (1939). *Management and the Worker*. Harvard University Press.

Rogers, Paul and Robert W. Brennan (1997). "A simulation testbed for comparing the performance of alternative control architectures." In: *Proceedings of the 29th Winter Simulation Conference*. Ed. by S. Andradorrit, K. J. Healy, D. H. Withers, and B. L. Nelson. IEEE Computer Society. Atlanta, GA, USA, pp. 880–887. isbn: 0-7803-4278-X. doi: 10.1145/268437.268669. acm id: 268669.

Rolón, Milagros and Ernesto Martínez (2012). "Agent-based modeling and simulation of an autonomic manufacturing execution system." In: *Computers in Industry* 63.1, pp. 53 –78. issn: 0166-3615. doi: 10.1016/j.compind.2011.10.005.

Rosenthal, Robert W. (1973). "A class of games possessing pure-strategy Nash equilibria." In: *International Journal of Game Theory* 2.1, pp. 65–67. issn: 1432-1270. doi: 10.1007/BF01737559.

Roughgarden, Tim (2001). "Designing networks for selfish users is hard." In: *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*. Oct. 2001, pp. 472–481. doi: 10.1109/SFCS.2001.959923.

– (2003). "The price of anarchy is independent of the network topology." In: *Journal of Computer and System Sciences* 67.2. Special Issue on STOC 2002, pp. 341–364. issn: 0022-0000. doi: 10.1016/S0022-0000(03)00044-8.

– (2004). "Stackelberg Scheduling Strategies." In: *SIAM Journal on Computing* 33.2, pp. 332–350. doi: 10.1137/S0097539701397059.

– (2005). *Selfish routing and the price of anarchy*. English. Vol. 174. Cambridge, MA: MIT press. isbn: 0-262-18243-2.

– (2007). "Routing Games." In: *Algorithmic Game Theory*. Ed. by Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V Vazirani. New York, NY: Cambridge University Press. isbn: 978-0-521-87282-9. doi: 10.1017/CBO9780511800481.020.

Roughgarden, Tim and Éva Tardos (2002). "How Bad is Selfish Routing?" In: *Journal of the ACM* 49.2 (Mar. 2002), pp. 236–259. issn: 0004-5411. doi: 10.1145/506147.506153. acm id: 506153.

Roy, Bernard and B Sussmann (1964). "Les problèmes d'ordonnancement avec contraintes disjonctives." In: *Note DS* 9.

Russell, Stuart and Eric Wefald (1991). "Principles of metareasoning." In: *Artificial Intelligence* 49.1, pp. 361 –395. issn: 0004-3702. doi: 10.1016/0004-3702(91)90015-C.

Ryan, Sarah M., Bruno Baynat, and F. Fred Choobineh (2000). "Determining inventory levels in a CONWIP controlled job shop." In: *IIE Transactions* 32.2, pp. 105–114. issn: 1573-9724. doi: 10.1023/A:1007602028992.

Sallez, Y., T. Berger, S. Raileanu, S. Chaabane, and D. Trentesaux (2010). "Semi-heterarchical control of FMS: From theory to application." In: *Engineering Applications of Artificial Intelligence* 23.8, pp. 1314 –1326. issn: 0952-1976. doi: 10.1016/j.engappai.2010.06.013.

Sandholm, Tuomas W. (2000). "Automated contracting in distributed manufacturing among independent companies." In: *Journal of Intelligent Manufacturing* 11.3, pp. 271–283. issn: 1572-8145. doi: 10.1023/A:1008971326005.

Sandholm, William H. (2001). "Potential Games with Continuous Player Sets." In: *Journal of Economic Theory* 97.1, pp. 81 –108. issn: 0022-0531. doi: 10.1006/jeth.2000.2696.

Santos, F. C. and J. M. Pacheco (2005). "Scale-Free Networks Provide a Unifying Framework for the Emergence of Cooperation." In: *Phys. Rev. Lett.* 95.9 (Aug. 2005), p. 098104. doi: 10.1103/PhysRevLett.95.098104.

Sarraj, Rochdi, Eric Ballot, Shenle Pan, Driss Hakimi, and Benoit Montreuil (2014). "Interconnected logistic networks and protocols: simulation-based efficiency assessment." In: *International Journal of Production Research* 52.11, pp. 3185–3208. doi: 10.1080/00207543.2013.865853.

Sbayti, Hayssam, Chung-Cheng Lu, and Hani S. Mahmassani (2007). "Efficient Implementation of Method of Successive Averages in Simulation-Based Dynamic Traffic Assignment Models for Large-Scale Network Applications." In: *Transportation Research Record: Journal of the Transportation Research Board* 2029.1, pp. 22 –30. doi: 10.3141/2029-03.

Schenk, Michael, Siegfried Wirth, and Egon Müller (2010). *Factory Planning Manual - Situation-Driven Production Facility Planning*. 1st. Berlin, Heidelberg: Springer Science + Business Media. doi: 10.1007/978-3-642-03635-4.

Scherer, Eric (1998). "Models, Systems and Reality: Knowledge Generation and Strategies for Systems Design." In: *Shop Floor Control — A Systems Perspective: From Deterministic Models towards Agile Operations Management*. Ed. by Eric Scherer. Berlin, Heidelberg: Springer, pp. 93–125. isbn: 978-3-642-60313-6. doi: 10.1007/978-3-642-60313-6_4.

Schneeweiss, Christoph (2003a). *Distributed Decision Making*. 2nd. Berlin, Heidelberg, New York: Springer. isbn: 978-3-642-07289-5. doi: 10.1007/978-3-540-24724-1.

– (2003b). "Distributed decision making — a unified approach." In: *European Journal of Operational Research* 150.2, pp. 237 –252. issn: 0377-2217. doi: 10.1016/S0377-2217(02)00501-5.

– (2003c). "Distributed decision making in supply chain management." In: *International Journal of Production Economics* 84.1, pp. 71 –83. issn: 0925-5273. doi: 10.1016/S0925-5273(02)00381-X.

Schneider, Marguerite and Mark Somers (2006). "Organizations as complex adaptive systems: Implications of Complexity Theory for leadership research." In: *The Leadership Quarterly* 17.4, pp. 351 –365. issn: 1048-9843. doi: 10.1016/j.leaqua.2006.04.006.

Scholz-Reiter, Bernd, Michael Freitag, C. de Beer, and T. Jagalski (2006). "The Influence of Production Networks Complexity on the Performance of Autonomous Control Methods." In: *Intelligent Computation in Manufacturing Engineering. Proceedings of the 5th CIRP International Seminar on Computation in Manufacturing Engineering (CIRP ICME 06)*. Ed. by Roberto Teti. Naples, Italy: University of Naples, pp. 317–320.

Scholz-Reiter, Bernd, Thomas Jagalski, and Julia C. Bendul (2008). "Autonomous Control of a Shop Floor Based on Bee's Foraging Behaviour." English. In: *Dynamics in Logistics*. Ed. by Hans-Jörg Kreowski, Bernd Scholz-Reiter, and Hans-Dietrich Haasis. Berlin, Heidelberg: Springer, pp. 415–423. isbn: 978-3-540-76861-6. doi: 10.1007/978-3-540-76862-3_41.

Scholz-Reiter, Bernd, Michael Görges, and Thorsten Philipp (2009a). "Autonomously controlled production systems — Influence of autonomous control level on logistic performance." In: *CIRP Annals - Manufacturing Technology* 58.1, pp. 395 –398. issn: 0007-8506. doi: 10.1016/j.cirp.2009.03.011.

Scholz-Reiter, Bernd, J. Kolditz, and Torsten Hildebrandt (2009b). "Engineering autonomously controlled logistic systems." In: *International Journal of Production Research* 47.6, pp. 1449 –1468. issn: 0020-7543. doi: 10.1080/00207540701581791.

Schönsleben, Paul (2012). *Integral logistics management*. Boca Raton, FL: CRC Press ET - 4. isbn: 9781439878231.

Schuh, Günther, ed. (2006). *Produktionsplanung und -steuerung*. 3rd. Berlin, Heidelberg: Springer. isbn: 978-3-540-40306-7. doi: 10.1007/3-540-33855-1.

Sengupta, S. S. and Russell L. Ackoff (1965). "Systems Theory from an Operations Research Point of View." In: *IEEE Transactions on Systems Science and Cybernetics* 1.1, pp. 9–13. issn: 0536-1567. doi: 10.1109/TSSC.1965.300053.

Seredynski, F. and A.Y. Zomaya (2002). "Sequential and parallel cellular automata-based scheduling algorithms." In: *Parallel and Distributed Systems, IEEE Transactions on* 13.10, pp. 1009–1023. issn: 1045-9219. doi: 10.1109/TPDS.2002.1041877.

Sethi, AndreaKrasa and SureshPal Sethi (1990). "Flexibility in manufacturing: A survey." English. In: *International Journal of Flexible Manufacturing Systems* 2.4, pp. 289–328. issn: 0920-6299. doi: 10.1007/BF00186471.

Shaw, Michael J. (1987). "A distributed scheduling method for computer integrated manufacturing: the use of local area networks in cellular systems." In: *International Journal of Production Research* 25.9, pp. 1285–1303. doi: 10.1080/00207548708919915.

Sheffi, Yosef and Warren Powell (1981). "A comparison of stochastic and deterministic traffic assignment over congested networks." In: *Transportation Research Part B: Methodological* 15.1, pp. 53 –64. issn: 0191-2615. doi: 10.1016/0191-2615(81)90046-1.

Shen, Jian, Fuyong Xu, and Peng Zheng (2005). "A tabu search algorithm for the routing and capacity assignment problem in computer networks." In: *Computers & Operations Research* 32.11, pp. 2785 –2800. issn: 0305-0548. doi: 10.1016/j.cor.2004.04.004.

Shen, Weiming (2002). "Distributed manufacturing scheduling using intelligent agents." In: *Intelligent Systems, IEEE* 17.1, pp. 88–94. issn: 1541-1672. doi: 10.1109/5254.988492.

Shen, Weiming, Lihui Wang, and Qi Hao (2006a). "Agent-based distributed manufacturing process planning and scheduling: a state-of-the-art survey." In: *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* 36.4, pp. 563–577. issn: 1094-6977. doi: 10.1109/TSMCC.2006.874022.

Shen, Weiming, Qi Hao, Hyun Joong Yoon, and Douglas H. Norrie (2006b). "Applications of agent-based systems in intelligent manufacturing: An updated review." In: *Advanced Engineering Informatics* 20.4, pp. 415 –431. issn: 1474-0346. doi: 10.1016/j.aei.2006.05.004.

Shenker, S. J. (1995). "Making greed work in networks: a game-theoretic analysis of switch service disciplines." In: *IEEE/ACM Transactions on Networking* 3.6, pp. 819–831. issn: 1063-6692. doi: 10.1109/90.477727.

Shore, Jesse, Ethan Bernstein, and David Lazer (2015). "Facts and Figuring: An Experimental Investigation of Network Structure and Performance in Information and Solution Spaces." In: *Organization Science* 26.5, pp. 1432–1446. doi: 10.1287/orsc.2015.0980.

Siggelkow, Nicolaj and Daniel A. Levinthal (2003). "Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation." In: *Organization Science* 14.6, pp. 650–669. doi: 10.1287/orsc.14.6.650.24840. jstor: 4135126.

Siggelkow, Nicolaj and Jan W. Rivkin (2005). "Speed and Search: Designing Organizations for Turbulence and Complexity." In: *Organization Science* 16.2, pp. 101–122. doi: 10.1287/orsc.1050.0116. jstor: 25145954.

Simon, Herbert A. (1962). "The Architecture of Complexity." In: *Proceedings of the American Philosophical Society* 106.6, pp. 467–482. issn: 0003-049X. jstor: 985254.

– (1996). *The sciences of the artificial.* 3rd Edition. MIT press. isbn: 9780262193740.

Smith, Michael D., Norman Ramsey, and Glenn Holloway (2004). "A Generalized Algorithm for Graph-coloring Register Allocation." In: *SIGPLAN Not.* 39.6, pp. 277–288. issn: 0362-1340. doi: 10.1145/996893.996875. acm id: 996875.

Smith, Reid G. (1980). "The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver." In: *IEEE Transactions on Computers* 29.12, pp. 1104–1113. issn: 0018-9340. doi: 10.1109/TC.1980.1675516.

Solot, Philippe and Mario van Vliet (1994). "Analytical models for FMS design optimization: A survey." English. In: *International Journal of Flexible Manufacturing Systems* 6.3, pp. 209–233. issn: 0920-6299. doi: 10.1007/BF01328812.

Solotorevsky, Gadi and Ehud Gudes (1997). "Solving a real-life time tabling and transportation problem using distributed CSP techniques." In: *Constraints & Agents: Papers from the 1997 AAAI Workshop*, pp. 148–153. isbn: 978-1-57735-032-3.

Solotorevsky, Gadi, Ehud Gudes, and Amnon Meisels (1996). “Modeling and solving distributed constraint satisfaction problems (DCSPs).” In: *Principles and Practice of Constraint Programming — CP’96: Second International Conference, CP96 Cambridge, MA, USA, August 19–22, 1996 Proceedings*. Ed. by Eugene C. Freuder. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 561–562. isbn: 978-3-540-70620-5. doi: 10.1007/3-540-61551-2_116.

Solow, Daniel and Chartchai Leenawong (2003). “Mathematical Models for Studying the Value of Cooperational Leadership in Team Replacement.” In: *Computational & Mathematical Organization Theory* 9.1, pp. 61–81. issn: 1572-9346. doi: 10.1023/B:CMOT.0000012309.42988.2f.

Solow, Daniel and Joseph G. Szmerekovsky (2006). “The role of leadership: What management science can give back to the study of complex systems.” In: *Emergence: Complexity & Organization* 8.4, pp. 52 –60. issn: 1532-7000.

Solow, Daniel, Sandy Piderit, Apostolos Burnetas, and Chartchai Leenawong (2005). “Mathematical Models for Studying the Value of Motivational Leadership in Teams.” In: *Computational & Mathematical Organization Theory* 11.1, pp. 5–36. issn: 1572-9346. doi: 10.1007/s10588-005-1752-0.

Spearman, Mark L. and Michael A. Zazanis (1992). “Push and Pull Production Systems: Issues and Comparisons.” In: *Operations Research* 40.3, pp. 521–532. issn: 0030-364X. doi: 10.1287/opre.40.3.521. jstor: 171414.

Spearman, Mark L., David L. Woodruff, and Wallace J. Hopp (1990). “CONWIP: a pull alternative to kanban.” In: *International Journal of Production Research* 28.5, pp. 879–894. doi: 10.1080/00207549008942761.

Spencer, Michael S. (1991). “Using "The Goal" in an MRP System.” In: *Production & Inventory Management Journal* 32.4, pp. 22 –28. issn: 0897-8336.

Stacey, Ralph D. (1993). “Strategy as order emerging from chaos.” In: *Long Range Planning* 26.1, pp. 10 –17. issn: 0024-6301. doi: 10.1016/0024-6301(93)90228-8.

– (1995). “The Science of Complexity: An Alternative Perspective for Strategic Change Processes.” In: *Strategic Management Journal* 16.6, pp. 477–495. issn: 01432095, 10970266. jstor: 2486790.

Steegmans, Elke, Kurt Schelfthout, Tom Holvoet, Yolande Berbers, Paul Valckenaers, and Bart Saint Germain (2004). “A Basic Taxonomy for Role Composition.” In: *Software Engineering for Multi-Agent Systems II: Research Issues and Practical Applications*. Ed. by Carlos Lucena, Alessandro Garcia, Alexander Romanovsky, Jaelson Castro, and Paulo S. C. Alencar. Berlin, Heidelberg: Springer, pp. 92–110. isbn: 978-3-540-24625-1. doi: 10.1007/978-3-540-24625-1_6.

Stephan, Klaas Enno, Lars Kasper, Lee M. Harrison, Jean Daunizeau, Hanneke E.M. den Ouden, Michael Breakspear, and Karl J. Friston (2008). “Nonlinear dynamic causal models for fMRI.” In: *NeuroImage* 42.2, pp. 649 –662. issn: 1053-8119. doi: 10.1016/j.neuroimage.2008.04.262. pubmed: 18565765.

Stepney, Susan, Fiona A. C. Polack, and Turner Heather R. (2006). “Engineering emergence.” In: *11th IEEE International Conference on Engineering of Complex Computer Systems (ICECCS’06)*. doi: 10.1109/ICECCS.2006.1690358.

Stevenson, M., L. C. Hendry, and B. G. Kingsman (2005). “A review of production planning and control: the applicability of key concepts to the make-to-order industry.” In: *International Journal of Production Research* 43.5, pp. 869–898. issn: 00207543.

Stevenson, William J. (2009). *Operations Management*. 9th. McGraw Hill. isbn: 978-0-07-066821-8.

Stommel, Heinz Josef and Dieter Kunz (1973). *Untersuchungen über Durchlaufzeiten in Betrieben der metallverarbeitenden Industrie mit Einzel- und Kleinserienfertigung*. Springer. doi: 10.1007/978-3-322-88578-4.

Stone, Peter and Manuela Veloso (2000). "Multiagent Systems: A Survey from a Machine Learning Perspective." In: *Autonomous Robots* 8.3, pp. 345–383. issn: 1573-7527. doi: 10.1023/A:1008942012299.

Subba Rao, S., A. Gunasekaran, S. K. Goyal, and T. Martikainen (1998). "Waiting line model applications in manufacturing." In: *International Journal of Production Economics* 54.1, pp. 1 –28. issn: 0925-5273. doi: 10.1016/S0925-5273(97)00129-1.

Surana, Amit, Soundar Kumara, Mark Greaves, and Usha Nandini Raghavan (2005). "Supply-chain networks: a complex adaptive systems perspective." In: *International Journal of Production Research* 43.20, pp. 4235 –4265. issn: 00207543. doi: 10.1080/00207540500142274.

Swiecicka, A. and F. Seredynski (2000). "Cellular automata approach to scheduling problem." In: *Parallel Computing in Electrical Engineering, 2000. PARELEC 2000. Proceedings. International Conference on*, pp. 29–33. doi: 10.1109/PCEE.2000.873596.

Tantawi, Asser N. and Don Towsley (1985). "Optimal Static Load Balancing in Distributed Computer Systems." In: *J. ACM* 32.2 (Apr. 1985), pp. 445–465. issn: 0004-5411. doi: 10.1145/3149.3156. acm id: 3156.

Tay, Joc Cing and Nhu Binh Ho (2008). "Evolving dispatching rules using genetic programming for solving multi-objective flexible job-shop problems." In: *Computers & Industrial Engineering* 54.3, pp. 453 –473. issn: 0360-8352. doi: 10.1016/j.cie.2007.08.008.

Taylor, Frederick Winslow (1911). *The principles of scientific management.* New York, London: Harper & Brothers.

Tero, Atsushi, Seiji Takagi, Tetsu Saigusa, Kentaro Ito, Dan P. Bebber, Mark D. Fricker, Kenji Yumiki, Ryo Kobayashi, and Toshiyuki Nakagaki (2010). "Rules for Biologically Inspired Adaptive Network Design." In: *Science* 327.5964, pp. 439–442. issn: 0036-8075. doi: 10.1126/science.1177894. pubmed: 20093467.

Tetzlaff, Ulrich A.W. (1995). "A model for the minimum cost configuration problem in flexible manufacturing systems." English. In: *International Journal of Flexible Manufacturing Systems* 7.2, pp. 127–146. issn: 0920-6299. doi: 10.1007/BF01358906.

Tharumarajah, A. (2001). "Survey of resource allocation methods for distributed manufacturing systems." In: *Production Planning & Control* 12.1, pp. 58–68. doi: 10.1080/09537280150203988.

Tharumarajah, A., A. J. Wells, and L. Nemes (1996). "Comparison of the bionic, fractal and holonic manufacturing system concepts." In: *International Journal of Computer Integrated Manufacturing* 9.3, pp. 217–226. doi: 10.1080/095119296131670.

The Economist Intelligence Unit (2014). *Networked Manufacturing: The Digital Future.* Briefing Paper. The Economist Intelligence Unit Ltd.

Theuer, Hanna (2012). "Analysis of Correlation of Lean and Autonomous Production Processes." In: *Proceeding of the 22nd International Conference of Flexible Automation and Intelligent Manufacturing (FAIM 2012), Helsinki, Finnland.* Ed. by Hasse Nylund, Satu Kantti, Ville Toivonen, and Seppo Torvinen. Tampere, pp. 520–530. isbn: 978-952-15-2783-8.

Thomas, Catherine, Renata Kaminska-Labbé, and Bill McKelvey (2005a). *Managing the Control/Autonomy Dilemma: From Impossible Balance to Irregular Oscillation Dynamics.* Research Paper. University of Nice.

– (2005b). "Managing the MNC and Exploitation/Exploration Dilemma: From Static Balance to Dynamic Oscillation." In: *Strategy Process.* Ed. by Gabriel Szulanski, Joe Porac, and Yvez Doz. Vol. 22. Advances in Strategic Management. Amsterdam, Boston: Elsevier. Chap. 10, pp. 213–247. isbn: 978-0-7623-1200-9. doi: 10.1016/S0742-3322(05)22008-9.

T'kindt, Vincent and Jean-Charles Billaut (2006). *Multicriteria scheduling — Theory, models and algorithms.* Berlin: Springer. isbn: 9783540282303. doi: 10.1007/b106275.

Toffler, Alvin (1984). In: Prigogine, Ilya and Isabelle Stengers. *Order out of chaos: Man's new dialogue with nature.* Toronto: Bantam Books. Chap. Foreword, pp. xi –xxvi. isbn: 0-553-34082-4.

Tokoro, Mario (1996). "Agents: Towards a society in which humans and computers cohabitate." In: *Distributed Software Agents and Applications: 6th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW '94 Odense, Denmark, August 3–5, 1994 Proceedings*. Ed. by John W. Perram and Jean-Pierre Müller. Berlin, Heidelberg: Springer, pp. 1–10. isbn: 978-3-540-68335-3. doi: 10.1007/3-540-61157-6_17.

Trentesaux, Damien (2009). "Distributed control of production systems." In: *Engineering Applications of Artificial Intelligence* 22.7. Distributed Control of Production Systems, pp. 971 –978. issn: 0952-1976. doi: 10.1016/j.engappai.2009.05.001.

Trentesaux, Damien, Cyrille Pach, Abdelghani Bekrar, Yves Sallez, Thierry Berger, Thérèse Bonte, Paulo Leitão, and José Barbosa (2013). "Benchmarking flexible job-shop scheduling and control systems." In: *Control Engineering Practice* 21.9, pp. 1204 –1225. issn: 0967-0661. doi: 10.1016/j.conengprac.2013.05.004.

Tsoukas, Haridimos (1996). "The Firm as a Distributed Knowledge System: A Constructionist Approach." In: *Strategic Management Journal* 17. Special Issue: Knowledge and the Firm, pp. 11–25. issn: 0143-2095. jstor: 2486988.

Uhl-Bien, Mary and Russ Marion (2009). "Complexity leadership in bureaucratic forms of organizing: A meso model." In: *The Leadership Quarterly* 20.4. Meso-Modeling of Leadership: Integrating Micro- and Macro-Perspectives of Leadership, pp. 631 –650. issn: 1048-9843. doi: 10.1016/j.leaqua.2009.04.007.

Uhl-Bien, Mary, Russ Marion, and Bill McKelvey (2007). "Complexity Leadership Theory: Shifting leadership from the industrial age to the knowledge era." In: *The Leadership Quarterly* 18.4. Leadership and Complexity, pp. 298 –318. issn: 1048-9843. doi: 10.1016/j.leaqua.2007.04.002.

Uzsoy, Reha, Chung-Yee Lee, and Louis A. Martin-Vega (1994). "A Review of Production Planning and Scheduling Models in the Semiconductor Industry Part II: Shop-Floor Control." In: *IIE Transactions* 26.5, pp. 44–55. doi: 10.1080/07408179408966627.

Valckenaers, Paul, Hendrik van Brussel, and Tom Holvoet (2008). "Fundamentals of Holonic Systems and Their Implications for Self-Adaptive and Self-Organizing Systems." In: *Self-Adaptive and Self-Organizing Systems Workshops, 2008. SASOW 2008. Second IEEE International Conference on*, pp. 168–173. doi: 10.1109/SASOW.2008.29.

Van Brussel, Hendrik, Jo Wyns, Paul Valckenaers, Luc Bongaerts, and Patrick Peeters (1998). "Reference architecture for holonic manufacturing systems: PROSA." In: *Computers in Industry* 37.3, pp. 255 –274. issn: 0166-3615. doi: 10.1016/S0166-3615(98)00102-X.

Van Dyke Parunak, H. (1987). "Manufacturing experience with the contract net." In: *Distributed artificial intelligence*. Ed. by Michael N. Huhns. Research Notes in Artificial Intelligence. London: Pitman Publishing. Chap. 10, pp. 285–310. isbn: 0-273-08778-9.

– (1996). "Applications of distributed artificial intelligence in industry." In: *Foundations of Distributed Artificial Intelligence*. Ed. by G. M. P. O'Hare and N. R. Jennings. New York, NY: Wiley Interscience. Chap. 4, pp. 139–164. isbn: 0-471-00675-0.

– (1997). "'Go to the ant': Engineering principles from natural multi-agent systems." English. In: *Annals of Operations Research* 75.0, pp. 69–101. issn: 0254-5330. doi: 10.1023/A:1018980001403.

– (1999). "From chaos to commerce: Practical issues and research opportunities in the nonlinear dynamics of decentralized manufacturing systems." In: *Proceedings of Second International Workshop on Intelligent Manufacturing Systems 1999 (IMS 1999)*. KU Leuven, k15–k25. isbn: 90-73802-69-5.

– (2000). "Agents in Overalls: Experiences and Issues in the Development and Deployment of Industrial Agent-Based Systems." In: *International Journal of Cooperative Information Systems* 09.03, pp. 209–227. doi: 10.1142/S0218843000000119.

Van Dyke Parunak, H. and Sven Brueckner (2001). “Entropy and Self-organization in Multi-agent Systems.” In: *Proceedings of the Fifth International Conference on Autonomous Agents*. AGENTS ’01. Montreal, Quebec, Canada: ACM, pp. 124–130. isbn: 1-58113-326-X. doi: 10.1145/375735.376024. acm id: 376024.

Van Dyke Parunak, H., Sven A. Brueckner, Robert Matthews, and John Sauter (2003). “How to Calm Hyperactive Agents.” In: *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*. AAMAS ’03. Melbourne, Australia: ACM, pp. 1092–1093. isbn: 1-58113-683-8. doi: 10.1145/860575.860811. acm id: 860811.

Vánca, József (2014). “Agent Theory.” In: *CIRP Encyclopedia of Production Engineering*. Ed. by Luc Laperrière and Gunther Reinhart. Berlin, Heidelberg: Springer, pp. 33–39. isbn: 978-3-642-20616-0. doi: 10.1007/978-3-642-20617-7_100016.

Vattani, Andrea (2012). “Algorithmic modeling of decision making over networks.” PhD thesis. University of California, San Diego.

Vec, Miloš, Alexandra Margarete Freund, and Marc-Thorsten Hütt, eds. (2006). *Selbstorganisation: Ein Denksystem für Natur und Gesellschaft*. Köln: Böhlau Verlag. isbn: 978-3412221058.

van der Vecht, Bob, André P. Meyer, Martijn Neef, Frank Dignum, and John-Jules Ch. Meyer (2007). “Influence-Based Autonomy Levels in Agent Decision-Making.” English. In: *Coordination, Organizations, Institutions, and Norms in Agent Systems II*. Ed. by Pablo Noriega, Javier Vázquez-Salceda, Guido Boella, Olivier Boissier, Virginia Dignum, Nicoletta Fornara, and Eric Matson. Vol. 4386. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 322–337. isbn: 978-3-540-74457-3. doi: 10.1007/978-3-540-74459-7_21.

Verstraete, Paul, Bart Saint Germain, Paul Valckenaers, Hendrik Van Brussel, Jan Belle, and Hadeli Hadeli (2008a). “Engineering manufacturing control systems using PROSA and delegate MAS.” In: *International Journal of Agent-Oriented Software Engineering* 2.1, pp. 62–89. doi: 10.1504/IJAOSE.2008.0168.

Verstraete, Paul, P. Valckenaers, H. Van Brussel, B. Saint Germain, K. Hadeli, and J. Van Belle (2008b). “Towards robust and efficient planning execution.” In: *Engineering Applications of Artificial Intelligence* 21.3, pp. 304 –314. issn: 0952-1976. doi: 10.1016/j.engappai.2007.09.002.

Vicari, Salvatore, Georg von Krogh, Johan Roos, and Volker Mahnke (1996). “Knowledge Creation through Cooperative Experimentation.” In: *Managing Knowledge: Perspectives on Cooperation and Competition*. Ed. by Georg von Krogh and Johan Roos. London: SAGE Publications. Chap. 9, pp. 184–202. isbn: 0-7619-5180-6. doi: 10.4135/9781446280195.n10.

Viswanadham, N., Y. Narahari, and T.L. Johnson (1992). “Stochastic modelling of flexible manufacturing systems.” In: *Mathematical and Computer Modelling* 16.3, pp. 15 –34. issn: 0895-7177. doi: 10.1016/0895-7177(92)90043-K.

Vrabič, Rok and Peter Butala (2012). “Assessing operational complexity of manufacturing systems based on statistical complexity.” In: *International Journal of Production Research* 50.14, pp. 3673 –3685. issn: 00207543. doi: 10.1080/00207543.2011.575098.

Vrabič, Rok, Damir Husejnagić, and Peter Butala (2012). “Discovering autonomous structures within complex networks of work systems.” In: *CIRP Annals - Manufacturing Technology* 61.1, pp. 423 –426. issn: 0007-8506. doi: 10.1016/j.cirp.2012.03.083.

Wacker, John G (1998). “A definition of theory: research guidelines for different theory-building research methods in operations management.” In: *Journal of Operations Management* 16.4, pp. 361 –385. issn: 0272-6963. doi: 10.1016/S0272-6963(98)00019-9.

Wang, Judith Y.T., Andrea Raith, and Matthias Ehrgott (2010). “Tolling Analysis with Bi-objective Traffic Assignment.” English. In: *Multiple Criteria Decision Making for Sustainable Energy and Transportation Systems*. Ed. by Matthias Ehrgott, Boris Naujoks, Theodor J. Stewart, and Jyrki Wallenius. Vol. 634. Lecture Notes in Economics and Mathematical

Systems. Berlin, Heidelberg: Springer, pp. 117–129. isbn: 978-3-642-04044-3. doi: 10.1007/978-3-642-04045-0_10.

Wang, Tianyi, Henning Blunck, and Julia Bendul (2017). "Exploring the Design Space for Myopia-Avoiding Distributed Control Systems Using a Classification Model." In: *Service Orientation in Holonic and Multi-Agent Manufacturing, Proceedings of SOHOMA 2016*. Ed. by Theodor Borangiu, Damien Trentesaux, André Thomas, Paulo Leitão, and José A. Barata Oliveira. Vol. 694. Studies in Computational Intelligence. Springer International Publishing. Chap. 26, pp. 295–304. isbn: 978-3-319-51099-6. doi: 10.1007/978-3-319-51100-9_26.

Wang, Xiao Fan and Guanrong Chen (2002). "Synchronization in scale-free dynamical networks: robustness and fragility." In: *Circuits and Systems I: Fundamental Theory and Applications, IEEE Transactions on* 49.1 (Jan. 2002), pp. 54–62. issn: 1057-7122. doi: 10.1109/81.974874.

Wang, Yunzeng (2001). "The optimality of myopic stocking policies for systems with decreasing purchasing prices." In: *European Journal of Operational Research* 133.1, pp. 153 –159. issn: 0377-2217. doi: 10.1016/S0377-2217(00)00189-2.

Wardrop, John Glen (1952). "Some Theoretical Aspects of Road Traffic Research." In: *ICE Proceedings: Engineering Divisions* 1.3 (Jan. 1952), pp. 325–362. doi: 10.1680/ipeds.1952.11259.

Watts, Duncan J. and Steven H. Strogatz (1998). "Collective dynamics of 'small-world' networks." In: *Nature* 393.6684, pp. 440–442. doi: 10.1038/30918.

Wenzel, Sigrid, Matthias Weiss, Simone Collisi-Böhmer, Holger Pitsch, and Oliver Rose (2008). *Qualitätskriterien für die Simulation in Produktion und Logistik*. VDI-Buch. Berlin, Heidelberg: Springer. isbn: 978-3-540-35272-3. doi: 10.1007/978-3-540-35276-1.

Weyns, Danny, Tom Holvoet, and Alexander Helleboogh (2007). "Anticipatory Vehicle Routing using Delegate Multi-Agent Systems." In: *2007 IEEE Intelligent Transportation Systems Conference*, pp. 87–93. doi: 10.1109/ITSC.2007.4357809.

Wiendahl, H.-P. and P. Scholtissek (1994). "Management and Control of Complexity in Manufacturing." In: *CIRP Annals - Manufacturing Technology* 43.2, pp. 533 –540. issn: 0007-8506. doi: 10.1016/S0007-8506(07)60499-5.

Wiendahl, Hans-Peter (1997). *Fertigungsregelung*. Hanser Munich. isbn: 978-3-446-19084-9.

Wiendahl, Hans-Peter, Jürgen Reichardt, and Peter Nyhuis (2009). *Handbuch Fabrikplanung — Konzept, Gestaltung und Umsetzung wandlungsfähiger Produktionsstätten*. München, Wien: Carl Hanser Verlag. isbn: 978-3-446-22477-3.

Windt, K., O. Jeken, and A. Arbabzadah (2010a). "Improved logistics performance through the use of locked flexibility potentials." In: *Proceedings of the 43rd CIRP International Conference on Manufacturing Systems. Sustainable Production and Logistics in Global Networks*. Ed. by W. Sihn and P. Kuhlang. Vienna, Austria, pp. 892–899.

Windt, Katja and Michael Hülsmann (2007). "Changing Paradigms in Logistics — Understanding the Shift from Conventional Control to Autonomous Cooperation and Control." English. In: *Understanding Autonomous Cooperation and Control in Logistics*. Ed. by Michael Hülsmann and Katja Windt. Berlin, Heidelberg: Springer. Chap. 1, pp. 1–16. isbn: 978-3-540-47449-4. doi: 10.1007/978-3-540-47450-0_1.

Windt, Katja and Marc-Thorsten Hütt (2010). "Graph coloring dynamics: A simple model scenario for distributed decisions in production logistics." In: *CIRP Annals - Manufacturing Technology* 59.1 (Jan. 2010), pp. 461–464. doi: 10.1016/j.cirp.2010.03.082.

Windt, Katja, Felix Böse, and Thorsten Philipp (2008a). "Autonomy in production logistics: Identification, characterisation and application." In: *Robotics and Computer-Integrated Manufacturing* 24.4. ICMR2005: Third International Conference on Manufacturing Research, pp. 572 –578. issn: 0736-5845. doi: 10.1016/j.rcim.2007.07.008.

Windt, Katja, Thorsten Philipp, and Felix Böse (2008b). "Complexity cube for the characterization of complex production systems." In: *International Journal of Computer Integrated Manufacturing* 21.2, pp. 195 –200. issn: 0951192X. doi: 10.1080/09511920701607725.

Windt, Katja, Till Becker, Oliver Jeken, and Achim Gelessus (2010b). "A classification pattern for autonomous control methods in logistics." English. In: *Logistics Research* 2 (2 2010), pp. 109–120. issn: 1865-035X. doi: 10.1007/s12159-010-0030-9.

Windt, Katja, Till Becker, Dimitar Asenov, and Farhad Arbabzadah (2010c). "A generic implementation approach of autonomous control methods in production logistics." In: *8th IEEE International Conference on Control and Automation (ICCA)*, pp. 629 –633. doi: 10.1109/ICCA.2010.5524133.

Wolff, Ronald W. (1982). "Poisson Arrivals See Time Averages." In: *Operations Research* 30.2, pp. 223–231. issn: 0030364X, 15265463. doi: 10.1287/opre.30.2.223. jstor: 170165.

Wolfram, Stephen (1983). "Statistical mechanics of cellular automata." In: *Rev. Mod. Phys.* 55.3, pp. 601–644. doi: 10.1103/RevModPhys.55.601.

– (1984). "Cellular automata as models of complexity." In: *Nature* 311.5985 (Oct. 1984), pp. 419–424. doi: 10.1038/311419a0.

Womack, James P, Daniel T Jones, and Daniel Roos (1990). *The Machine that changed the world.* New York, NY: Simon and Schuster.

Wong, T.N., C.W. Leung, K.L. Mak, and R.Y.K. Fung (2006). "Dynamic shopfloor scheduling in multi-agent manufacturing systems." In: *Expert Systems with Applications* 31.3, pp. 486 –494. issn: 0957-4174. doi: 10.1016/j.eswa.2005.09.073.

Wooldridge, Michael and Nicholas R. Jennings (1995). "Intelligent agents: theory and practice." In: *The Knowledge Engineering Review* 10.02 (June 1995), pp. 115–152. issn: 1469-8005. doi: 10.1017/S0269888900008122.

Wu, Jianshe, Licheng Jiao, Rui Li, and Weisheng Chen (2011). "Clustering dynamics of nonlinear oscillator network: Application to graph coloring problem." In: *Physica D: Nonlinear Phenomena* 240.24, pp. 1972 –1978. issn: 0167-2789. doi: 10.1016/j.physd.2011.09.010.

Wu, S. David, Eui-Seok Byeon, and Robert H. Storer (1999). "A Graph-Theoretic Decomposition of the Job Shop Scheduling Problem to Achieve Scheduling Robustness." In: *Operations Research* 47.1, pp. 113–124. doi: 10.1287/opre.47.1.113. jstor: 222898.

Yang, Hai and Michael G. H. Bell (1998). "Models and algorithms for road network design: a review and some new developments." In: *Transport Reviews* 18.3, pp. 257–278. doi: 10.1080/01441649808717016.

Yeo, Jaehyun, Heesoo Lee, and Sehun Kim (2002). "An efficient broadcast scheduling algorithm for TDMA ad-hoc networks." In: *Computers & Operations Research* 29.13, pp. 1793 –1806. issn: 0305-0548. doi: 10.1016/S0305-0548(01)00057-0.

Yokoo, Makoto (2001). *Distributed Constraint Satisfaction — Foundations of Cooperation in Multi-agent Systems.* 1st Edition. Springer Series on Agent Technology. Berlin, Heidelberg: Springer. isbn: 3-540-67596-5. doi: 10.1007/978-3-642-59546-2.

Yokoo, Makoto, Edmund H. Durfee, Toru Ishida, and Kazuhiro Kuwabara (1998). "The distributed constraint satisfaction problem: formalization and algorithms." In: *Knowledge and Data Engineering, IEEE Transactions on* 10.5, pp. 673–685. issn: 1041-4347. doi: 10.1109/69.729707.

Youn, Hyejin, Michael T. Gastner, and Hawoong Jeong (2008). "Price of Anarchy in Transportation Networks: Efficiency and Optimality Control." In: *Phys. Rev. Lett.* 101.12 (Sept. 2008), p. 128701. doi: 10.1103/PhysRevLett.101.128701.

Zambrano Rey, Gabriel (2014). "Reducing Myopic Behavior in FMS Control: A Semi-Heterarchical SimulationOptimization Approach." PhD thesis. Université de Valenciennes et du Hainaut-Cambresis.

Zambrano Rey, Gabriel, C. Pach, N. Aissani, T. Berger, and D. Trentesaux (2011). "An approach for temporal myopia reduction in Heterarchical Control Architectures." In: *Industrial Electronics (ISIE), 2011 IEEE International Symposium on.* June 2011, pp. 1767–1772. doi: 10.1109/ISIE.2011.5984424.

Zambrano Rey, Gabriel, Nassima Aissani, Abdelghani Bekrar, and Damien Trentesaux (2012). "A Holonic Approach to Myopic Behavior Correction for the Allocation Process in Flexible-Job Shops Using Recursiveness." English. In: *Service Orientation in Holonic and Multi-Agent Manufacturing Control.* Ed. by Theodor Borangiu, André Thomas, and Damien Trentesaux. Vol. 402. Studies in Computational Intelligence. Berlin, Heidelberg: Springer, pp. 115–128. isbn: 978-3-642-27448-0. doi: 10.1007/978-3-642-27449-7_9.

Zambrano Rey, Gabriel, Cyrille Pach, Nassima Aissani, Abdelghani Bekrar, Thierry Berger, and Damien Trentesaux (2013). "The control of myopic behavior in semi-heterarchical production systems: A holonic framework." In: *Engineering Applications of Artificial Intelligence* 26.2, pp. 800 –817. issn: 0952-1976. doi: 10.1016/j.engappai.2012.08.011.

Zambrano Rey, Gabriel, Thérèse Bonte, Vittaldas Prabhu, and Damien Trentesaux (2014). "Reducing myopic behavior in FMS control: A semi-heterarchical simulation-optimization approach." In: *Simulation Modelling Practice and Theory* 46.0. Simulation-Optimization of Complex Systems: Methods and Applications, pp. 53 –75. issn: 1569-190X. doi: 10.1016/j.simpat.2014.01.005.

Zapf, Michael and Thomas Weise (2007). *Offline emergence engineering for agent societies.* Paper presented at EUMAS'07. Fifth European Workshop on Multi-Agent Systems. Hammamet, Tunesia December 13-14, 2007. urn: urn:nbn:de:hebis:34-2007120719844.

Zhang, Weixiong, Guandong Wang, and Lars Wittenburg (2002). "Distributed stochastic search for constraint satisfaction and optimization: Parallelism, phase transitions and performance." In: *Proceedings of AAAI Workshop on Probabilistic Approaches in Search, Technical Report WS-02-14.* AAAI Press. isbn: 978-1-57735-167-2.

Ziaee, Mohsen (2014). "A heuristic algorithm for solving flexible job shop scheduling problem." English. In: *The International Journal of Advanced Manufacturing Technology* 71.1-4, pp. 519–528. issn: 0268-3768. doi: 10.1007/s00170-013-5510-z.

Zimmermann, Armin (2008). *Stochastic Discrete Event Systems.* Berlin, Heidelberg: Springer. isbn: 978-3-540-74172-5. doi: 10.1007/978-3-540-74173-2.

Zuehlke, Detlef (2010). "SmartFactory — Towards a factory-of-things." In: *Annual Reviews in Control* 34.1, pp. 129 –138. issn: 1367-5788. doi: 10.1016/j.arcontrol.2010.02.008.

Index