Transforming Complex Business Challenges into Opportunities for Innovative Change - An Application for Planning and Scheduling Technology

Brian Drabble
DMM Ventures, Inc
Yorktown, United States

Bernd Schattenberg
Büro für intelligente Technologie-Beratung
Laupheim, Germany

Abstract

The new driving force of today's market has many names. Some call it Digitalization, Internet of Things, Industry 4.0, and many more. Based on technological advancements, the impact on the physical world with actionable information technology has been pushed to new levels. This in turn enabled small and agile players to create so-called disruptive innovations and to increase pressure on established players. As a consequence, all industries face a tremendous change. Both established as well as new players thus face a constantly growing complexity. Traditional decision support systems, which focus on the operational level, turn out to be ineffective to address that complexity. Additionally, businesses and organizations are facing new challenges such as the issues associated with the "Circular Economy" and the implications it has for the way they conduct their businesses. This presentation gives an introduction to application domains in this realm and points out solutions to current industries' most pressing issues, in which Artificial Intelligence Planning and Scheduling technology plays a central role.

Introduction – The Age of Digitization

Over recent years, the *digitization* wave mutated into a substantial challenge for enterprises of any size. What started out as a notion of "utilizing information technology" became one of the central driving forces in today's markets. The speed and versatility of the digital economy increases the pressure on virtually every traditional industry and service provider. The main aspects of this phenomenon are coined by Clayton M. Christensen by the term *disruptive innovations*, which basically means that there exist innovations, that create new markets and value networks, effectively displacing the players in existing markets (Bower and Christensen 1995). With the rise of the digital economy, the frequency and amplitude of such disruptive changes has increased immensely. In addition, recent economic developments like Circular Economy (CE) with its integral causal feedback loops represents an economical challenge but none the less a reasonable business opportunity (Ellen MacArthur Foundation 2015). Furthermore, it is an excellent use case for Artificial Intelligence Planning and Scheduling (AI P&S) technology, because strategic decisions in CE are inherently more complex than traditional ones.

In summary, we nowadays find disruptiveness as the architect of new business models. This implies for especially large players the need to adapt faster and with increasingly substantial change to stay synchronized with market dynamics (Kagermann 2015). Organizational scale, which was the former guarantor for success, now turns into a big disadvantage. Without having appropriate support at the operational level as well as at the strategic level we will see more organizations being displaced from the market.

The classical digitization strategies of major companies, like the closely related initiatives "Industry 4.0" and "Internet of Things" do on a broader scale, heavily focus on the technical issues of collaboration on the operational level. Hence, they focus on solutions to interoperability- and infrastructure-related problems together with the respective standards for their implementation. This represents a traditional, mostly on complication focused approach. Implications arising from complexity are not sufficiently covered by these approaches, as they do not take into account the dynamics of the processes involved and organizational entities that are required to explore upcoming market opportunities.

In order to address these business challenges in a constructive manner, we propose a paradigm that follows a goal-driven approach: First, the involved stakeholders agree on a (potentially new) value proposition and corresponding business models. Then they assess their existing processes, their own capabilities, and that of potential collaborators. From this information they develop the steps necessary to deliver that value proposition and the structure of the value-added chain that implements the desired business model – and they do so with support by AI P&S technology. In short: driven by the goal of serving a new market or delivering an innovative product, the stakeholders employ planning models and algorithms to decide how to realize their goals with the assistance of whom in which way. Since all these business dynamics are basically goal-oriented, the AI P&S community's technology portfolio constitutes a perfect match for the implicitly present demand.

We believe that – at this point of the economic development – the proposed AI P&S approach is considerably more adequate for the digital age than other paradigms, in particular manual planning techniques. How exactly AI P&S fits

Submitted to the 30th Workshop "Planen/Scheduling und Konfigurieren/Entwerfen" (PuK), Klagenfurt, Austria, September 26, 2016.

into the picture of the digitization and disruption movements portrayed above will be motivated in the following sections.

Transformation: Business Challenges Can Become Opportunities for Innovation

We begin with some reflections on the characteristics of the problems that we face when implementing a digitization-induced change in an organization. To this end, we can make the following four observations:

1. Every business is defined by its business model and a corresponding value-added chain. Every link of that chain, that means, every compartment of the company like *production* or *marketing*, is governed in terms of its workflows by well-defined business processes.
2. In general, these business processes are locally both well understood and highly optimized. This is, however, usually not the case for the implicit dependencies and causalities (or so-called *n-order effects*) on the level of a global value-added chain.
3. Established and, in particular, large businesses have set-up a highly specialized and efficient, workflows. However, these huge organiztions are usually inflexible which hampers their process change efforts. They become increasingly mandatory in cooperation and CE scenarios. This makes them particularly reluctant to adopting new *business strategies*. In addition, it renders them unable to react to disruptive stimuli in an adequate manner, because change in the organization has to propagate up the structural hierarchies and back. Adapting established business processes is only possible as a long-term initiative with numerous occasions being reported of failed company mergers.
4. Small businesses, for example start-up firms, do work completely different for that matter. Thus, cooperating with another business of similar size – whether within the same market segment and technology strain or not – is typically "only one phone-call away". All business models and value-added chains are explicitly known to the stakeholders and usually negotiable. Joint ventures are born by integrating individual contributions to a joint value proposition.

As stated above, economic scale was the key attribute to grant access to high volumes and possibly margins. Organizations therefore are structured around the idea of setting up static value chains to orchestrate actions of a large set of stakeholders. Functional decomposition together with low level of vertical integration resulted in inflexible structures. Change, in any form, is risk and will be impacted from both the chain and the involved functions.

A solution to this problem needs to ensure both alignment, to address efficiency, and to allow for autonomy, to address upcoming changes. Statically defined structures, processes and technologies won't be able to do this job. Our proposed goal-based approach ensures a convergent and aligned value chain that can instantaneously implement required opportunistic local transformation impulses. Required actions are developed and compliance is sustained by AI P&S and are integrated directly in to day-to-day operations.

The key strategic element of our proposed effective business decision support is to facilitate "controlled flexibility", that means, flexibly setting up alliances with business partners in order to realize new products and value propositions that were not possible by the individual firms alone. But this means instead of cooperating on a purely operational level, cooperations have to be installed on the organizational level, too. With an adequately detailed model of the stakeholders' value-added chains and business processes, AI P&S is able to produce plans that reflect:

1. Ad-hoc alignment of strategic goals, that means, cooperating in a trans-organizational way along the horizontal value-added chain. Improved products at considerably lower overhead costs make this scenario commercially interesting even for smaller markets.
2. Ad-hoc assignment of tasks, that is, utilizing new opportunities by creating and optimizing processes along the vertical value-added chain. This basically means, that the company can establish new processes and behaviors by utilizing its internal resources.

AI P&S is thereby employed to facilitate running dynamic organizations by designing aligned autonomous stakeholders. Eventually, both of the above notions of plans have to be utilized in order to gain the full benefit of the methodology and they work on both the large and the small scale:

- Setting up virtual organizations from the functional and organizational cells of the stakeholders' organizations. This can also be viewed as an implementation of important aspects of the *Rule of 150* (Gladwell 2006) and *Agile project management* (Highsmith 2009) in that planning conceptually takes defined organizational structures as a source for re-groupings in view of a new purpose.
- Controlling and orchestrating technical systems inside businesses. This domain has been studied for some time by the planning community, among others under the aspects of scheduling, mission-planning, and multi-agency. Although it resembles the application area of classical business process models (BPM), these processes are highly dynamic in nature and are most certainly never executed on a regular basis. Many techniques from the BPM literature will therefore not be available in this setting.

It becomes clear that, on a global, organization-wide scale, more than one technology is needed. A comprehensive and universal approach has to address collaboration issues and the differences of industry-specific applications at basically every single level of detail and ranging from an initial knowledge acquisition phase to the final process monitoring. Currently, the economy almost exclusively demands for information processing methods to harvest implicit information from inside their own business and from their (potential) customers. Techniques including data analytics and data mining provide access to this information and provide it in a form that can be directly used by AI P&S technologies. This raises another dimension of complexity to this problem, namely, what to do with all this knowledge [1]. We note that most of the relevant knowledge concerning the stakeholders strategic and operational options (viz., the very knowledge required for effectively employing AI P&S methods as sketched above) is available though is often overlooked in current approaches to AI P&S.

By adopting a more pragmatic approach, we therefore suggest concentrating on the first part of the overall scenario in order to establish AI P&S methods as our customers' primary means to manage and operationalize knowledge about their business dynamics. From that, future support methods can emerge, which integrate more stages of a comprehensive issue treatment approach and that introduce new facets for specific market situations, organization types, and so forth.

Our first application domain isolates an issue of the manufacturing systems engineering industry, which faces a significant change in their business model, namely the problem of responding to batch size 1 requests. requests. While we have used the batch 1 as an illustrative example, the same type of goal and plan driven approach applies to design, analysis and issues of the circuar economy. That means, their customers become increasingly interested in individualized products and therefore order custom-made units. Currently, established procedures are tailored for mass-production, so our proposed paradigm will make a difference by applying planning to an on-demand setup of production workflows, production resources, and proper handshakes with subcontractors, external suppliers, and internal service providers.

Two key components of one such solution are specific instances of AI P&S technology, deployed following in the Software-as-a-Service paradigm. The next section gives some details for these two aspects.

Adequate System Support for Transformations

Artificial Intelligence Planning and Scheduling Technology

In recent years, an increasing demand for more personalized products finally reached those industries which provide the actual product makers with their production tools; among them, the manufacturing systems engineering industry. These companies face the problem of receiving a decreasing number of large quantity orders, say, for supplying a manufacturing plant with dozens of virtually identical machines, and instead getting more and more orders for single, custom-tailored machines. Their evolutionary balanced business models begin to suffer for mainly two reasons:

1. Re-configuring the factory site such that the order can be processed is inherently costly, because it has been set-up in structurally fixed ways. And in addition, with every product being typically processed through a large number of passes on an equally large number of work stations, any deviation from an established routine challenges the organization's productivity.
2. Established batch processes are inadequate and have to be re-designed into more project-oriented *bidireactional* interactions that coordinate all in-house as well as subcontracted stakeholders like order processing, purchasing department, construction/production of components, sales department, and so forth.

Both factors are mutually amplified up to the point where nowadays the entrepreneurial risk becomes vital.

We propose the employment of *hybrid planning* (Schattenberg 2009; Kambhampati, Mali, and Srivastava 1998; Castillo, Fernández-Olivares, and González 2000) in these scenarios. This paradigm integrates building plans based on the individual action's causal structure (also referred to as *partial-order causal-link planning* (Penberthy and Weld 1992; McAllester and Rosenblitt 1991)) with obtaining plans from iteratively implementing abstract actions by pre-defined partial solutions (the so-called *hierarchical task-network planning* (Erol, Hendler, and Nau 1994; Yang 1998)). The hierarchical domain model aspects thereby represent regular solutions provided by domain experts, while the causality-based techniques complete underspecified procedures and address exceptional cases. Although any state-based approaches may be used in application areas of this kind (Rodríguez-Moreno et al. 2007; Hoffmann, Weber, and Kraft 2012), we prefer the hybrid methodology for their ability to adequately reflect procedural knowledge in combination with addressing variance by reasoning from first principles. As described above, the situation for the stakeholders maps onto the hybrid modeling method as follows:

- State of the art procedures are key to the industrial businesses. Back in the days, when running a business was just *complicated*, the stakeholders figured out very effectively all relevant processes and how to implement them. In terms of hybrid planning, the procedures are mapped on task networks, which are basically plans, consisting of tasks as process steps, orderings between the tasks, and causal dependencies between them. The procedures can be organized in a hierarchical fashion, such that different process variants can be subsumed as a set of tasks networks that *implement* an abstract place-holder task or process step, which may in turn be part of some more abstract process specifications, and so on. The emerging hierarchy of processes/tasks reflects the organizations' options on the strategic level (abstract tasks) as well as on the operational (primitive tasks).
- Solid processes may be key, but they are for stable organizations only. As long as a complicated situation is coped with by effective routines, everything is fine. But at some point the size-dimension of organizations has been surpassed by the dimension of their structural dependencies and thus turned business into *complex* ventures. In most real-world situations, the procedures will be underspecified and plainly not applicable anymore in any change situation. This is where causal reasoning on the task level, primitive or abstract, comes into play. By applying a means-end analysis and proving causal interactions and dependencies, planning builds goal-specific

[1] c.f. Quentin Hardy, "The Peril of Knowledge Everywhere", The New York Times, May 10, 2014

sequences of tasks and processes that implement proper workflows across organization boundaries, if required.

The adequacy of the hybrid planning technology and its domain modeling formalisms has been proven in similar application domains (Estlin, Chien, and Wang 1997; Castillo, Fernández-Olivares, and González 2001). We also note that capturing the (planning) models that represent their organizations and options is already an extremely valuable information for the stakeholders, for this knowledge most often lies unused or is not available in a sufficiently formal and, hence, unambiguous format.

While scheduling can already be found in numerous commercial application scenarios, ranging from staffing (Stylianou, Gerasimou, and Andreou 2012) to manufacturing scheduling (Framinan, Leisten, and García 2014), it becomes necessary in this wider sense to integrate planning and scheduling on basically every abstraction level of the model. Obviously, resource reasoning is relevant to all economic problems, but in the transformation scenario not all of the valid action options are known in advance. We learn that optimization issues in ad-hoc alliances are of lower-ranking interest (these will be addressed later when the cooperations are more consolidated), but a reliable feasibility analysis is most definitely vital.

Last, but not least, some additional aspects become important when the solution is to be fielded successfully, namely plan repair, interactive planning, and plan explanation (Schattenberg 2015). These functionalities are requested by practical considerations, because in particular at this early stage, all domain models and mutual understandings about the solution spaces are at any point in time still "work in progress". Recovering from not customer-approved solutions as well as guiding search more or less directly by domain experts is the economically most efficient way to begin with. On the other hand, if solutions are to be presented to decision-makers, the solution has to be trusted and finally accepted as the verdict of a virtual expert. This becomes in particular an issue when stakeholders from different organizations and backgrounds/disciplines meet.

Supporting Solution Generation with a Platform Strategy

A monolithic approach, in which domain management, AI P&S execution and human interaction management is integrated into a single software artifact turned out to be overly restrictive and not well suited to explore the arising market opportunities as outlined above. Our proposed architecture is based on loosely coupled micro services. The overall system can be decomposed into these four building blocks:

1. *Domain modeling and management*, that is a collection of tools to store, version, and collaborate on domain models. In our case we make use of text file based representations of domains, including PDDL [2].

   The PDDL representation is augmented and fused with a dependency network based model (Drabble 2014). This adopts a "system of systems" approach to domain modelling allowing for direct, indirect, cascading and cumulative effects of actions to be represented and reasoned with. A dependency network model consists of nodes (persons, groups, organizations, resources, locations, concepts, etc.) with pairs of nodes linked by arcs. The arcs specify strength of the dependency (range 1 – 10) and the nature of the dependency (actor operates physical, physical supplies physical, etc.). These direct dependencies are used to calculate transitive dependency scores for each node in the network. These scores are then ranked to identify the most important nodes in network. The ranking provides a user with the ability to identify bottlenecks (over dependency on a specified node), lack of robustness and opportunities for network function and output improvement. These identified improvements are used to generate goals and tasks from which intelligent tools can generate appropriate plans and schedules. Plans are converted into additional nodes and links which are added to the dependency network. These additional links and nodes may identify new dependencies or increase the score of known dependencies. This allows users to identify weaknesses in their plans and to iteratively improve them over time.

   As the creation and maintenance of domains is a highly iterative and collaborative process, providing the conceptual guidance during the complete lifecycle of the application, we went for the integration of GitHub as one of the widely supported and verified collaboration platforms (Longo and Kelley 2015). Versioning, branching, and pull-request support helps us to streamline the whole process and it is closely related to established technologies and development workflows.

2. *Plan execution*, that is specifically a question of which AI P&S strategy and philosophy is best suited to solve a certain problem. We found that there is a great demand for pluggable architectures, because application domains and therewith the effectiveness of AI P&S technology varies. Our planning as a service platform makes use of the Docker virtualization infrastructure, in which planner *configurations* are available as reusable components (Anderson 2015). They are immutable and versioned, such that at every time point any previously available image can be incorporated. Standardization helps to use commercial off-the-shelf cloud service instances to run these components at reasonable cost and with almost linear horizontal scaling support.

3. *Human interaction management*, that is a standardized way for humans to interact with the specific AI P&S setup that is prepared for a given application. We consider both development and maintenance as well as in-process interaction as use cases for our application front-end. The front-end is implemented as a state-of-the-art web interface. We integrate the pieces of our micro-service archi-

[2] In its latest version 3.1, the Planning Domain Definition Language does not support hierarchical domain model aspects. For those application scenarios in which non-hierarchical planning suffices, we propose planning systems that comply with the standard summarized in Daniel L. Kovacs' BNF description. In most cases, however, we employ a proprietary hybrid modeling language.

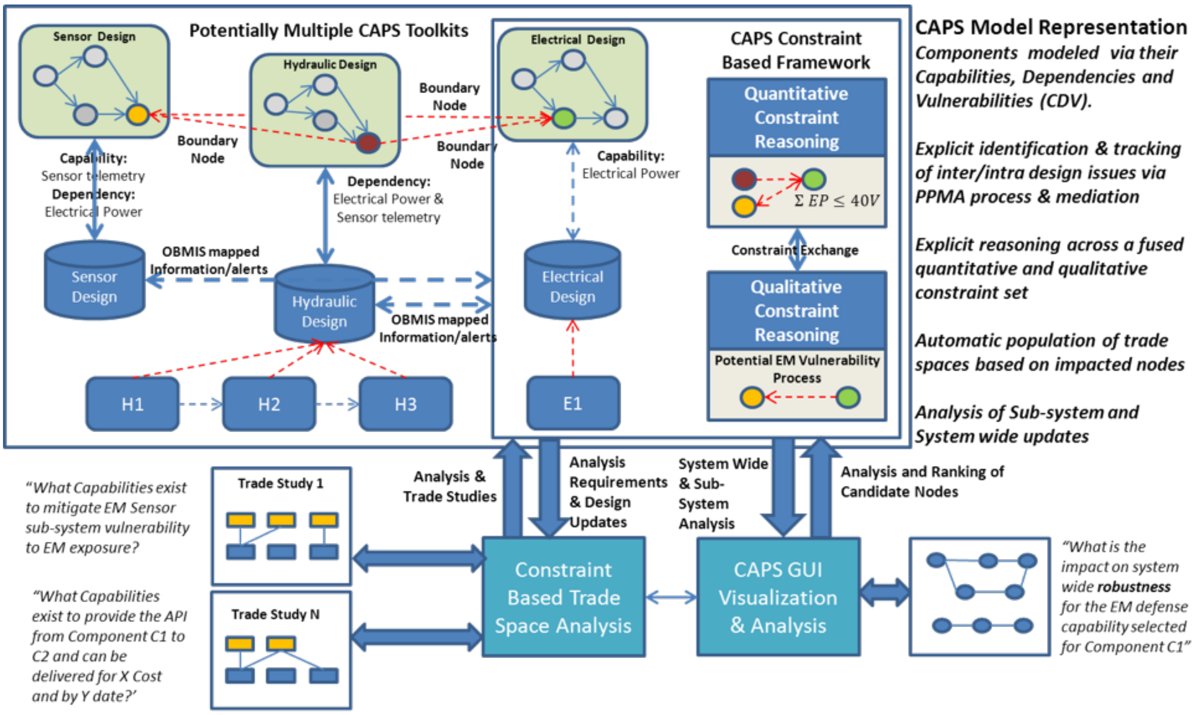

Figure 1: Cassandra Analysis and Planning System (CAPS) Architecture Overview

tecture into a seamless user experience. The main use cases are (1) environment configuration and maintenance, (2) tracking of the plan generation processes, (3) historical results archive, (4) error tracking and monitoring, and (5) simple human-machine interactions.

4. *Application Programming Interface (API)*, that is a machine-to-machine interface for sensors and actors of the application domain to directly interact with the AI P&S platform. We support both design/maintenance and operational use-cases. On the one hand, for example, a simple command line utility can be used by domain architects to trigger and monitor planning processes from a local copy of the domain repository whereas on the other hand a process execution support system can trigger and consume the results of the plan and schedule generation based on current sensor data and references to existing domains.

The dependency model described earlier was used as the primary representation for the CAPS (Drabble 2015). CAPS provides a distributed collaborative design and planning toolkit which allows individual designers to see the consequences of their own decisions and the impacts on other dependent designs. An overview of the CAPS system is provided in Figure 1 which shows the basic components of the architecture.

The right side of Figure 1 shows an instantiation of CAPS to support an EP sub-system designer and comprises the GUI through which the designer can view their design as a dependency graph, the information sources (for example, manuals, documents, specifications, etc.) the designers is using and details of the design held in the quantitative and qualitative systems respectively. This allows the EP designer to "drill down" on design nodes in the GUI to identify what options are available for a node (component or subsystem), instigate a trade study, assess the effects on their design of a decision they are considering, prioritize design decisions based on the effect on the design, etc. The Trade Space analysis component allows the EP designer to instigate trade studies which are auto-populated by CAPS based on the capability, dependency and vulnerability (CDV) of the node(s) selected by the designer. For example, the dependency of a hydraulic sensor on EP identifies it as a node potentially impacted by the trade study, the capability of generators to output different amounts of EP identifies them as candidates and the potential vulnerability of one or more of the impacted sensors to EM radiation also makes them a candidate. This allows CAPS to rank potential node decomposition and/or instantiations based on how well their capabilities meet the overall specification, their impact on other open decisions and identify potential trade-offs. For example, the choice of a generator in the EP sub-system design which has a slightly higher EM output than the "cheapest cost" choice would avoid imposing a "single choice" option on several components in the hydraulic sub-system design. The automatic population of the trade study matrix and identification of trade-offs firstly ensures only available options are considered, secondly it frees the user from having to search for trade-off options and thirdly decreases the overall design time by only requiring designers to interact when the mitigation step in the Propose $\rightarrow$ Propogate $\rightarrow$ Mediate $\rightarrow$ Assert process cannot find a suitable set of trade-offs.

Having this integrative platform in place is a key asset to focus on the development of specific AI P&S technologies whilst ensuring that domain architects can focus on their work. With emerging interoperable domain languages the platform will be key to further speed up development and value creation of AI P&S technology in domains as stated above.

Conclusion

Artificial Intelligence Planning and Scheduling technology turns out to be a key resource to solve essential transformational challenges in the context of the ongoing digitization. However, it has to be understood more as a tool or capability than as an actual product. As a powerful capability, it fosters the ability of organizations to integrate new knowledge into their operational perspective more quickly and more safely. As a product, if you will, it delivers advice in form of actionable and valid plans.

At the end of our presentation we have to note, that the actual degree of automation depends on the actual business case. While some organization might be able to utilize fully automated transformation processes and workflows, others will benefit exclusively from cooperation and interaction plans in a negotiation-safe format. The proposed methodology is able to address any scenario within this spectrum at basically any level of abstraction. The described challenge of "batch size 1" to the manufacturing industry serves as an example for an instantiation of the AI P&S methodology in this domain.

Our vision can be summarized as follows: Like database systems sustained the scale of information that came with the growth of businesses and web/cloud technology sustained the scatterdness of that information in increasingly distribution organizations, so will AI P&S sustain the complexity of (re-) acting within huge organizational bodies. Eventually, it will be perceived as an integral layer in the technology stack of modern software products. In order to foster this development, future work will focus on making the technology an Open-Source movement, including comprehensive repositories of domain models; this will help to pick up the momentum by early adopters.

Acknowledgement

We would like to thank Jan-Jakob Wachs of *Bapo UG* and Markus Kuhnt of *disrupt consulting e.G.* for their valuable input and support in an earlier version of this paper.

References

[Anderson 2015] Anderson, C. 2015. Docker [software engineering]. *IEEE Software* 32(3):102–c3. doi:10.1109/MS.2015.62.

[Bower and Christensen 1995] Bower, J. L., and Christensen, C. M. 1995. Disruptive technologies: Catching the wave. *Harvard Business Review* 73(1):43–53.

[Castillo, Fernández-Olivares, and González 2000] Castillo, L. A.; Fernández-Olivares, J.; and González, A. 2000. A hybrid hierarchical/operator-based planning approach for the design of control programs. In *ECAI Workshop on Planning and Configuration: New results in planning, scheduling and design*, 1–10.

[Castillo, Fernández-Olivares, and González 2001] Castillo, L. A.; Fernández-Olivares, J.; and González, A. 2001. On the adequacy of hierarchical planning characteristics for real-world problem solving. In *Proceedings of VI European Conference of Planning*.

[Drabble 2014] Drabble, B. 2014. Modeling c2 networks as dependencies: Understanding what the real issues are. In Grant, T.; Janssen, R.; and Monsuur, H., eds., *Network Topology in Command and Control: Organization, Operation, and Evolution*. Hershey, PA: Information Science Reference. 125–151. doi:10.4018/978-1-4666-6058-8.ch006.

[Drabble 2015] Drabble, B. 2015. Dependency-based collaborative design: a comparison of modeling methods. *Journal of Concurrency and Computation: Practice and Experience*. John Wiley & Sons, Ltd, doi:10.1002/cpe.3445.

[Ellen MacArthur Foundation 2015] 2015. Towards a circular economy: Business rationale for an accelerated transition. Ellen MacArthur Foundation. https://www.ellenmacarthurfoundation.org.

[Erol, Hendler, and Nau 1994] Erol, K.; Hendler, J.; and Nau, D. S. 1994. UMCP: A sound and complete procedure for hierarchical task-network planning. In *Proceedings of the 2nd International Conference on Artificial Intelligence Planning Systems (AIPS 1994)*, 249–254.

[Estlin, Chien, and Wang 1997] Estlin, T. A.; Chien, S. A.; and Wang, X. 1997. An argument for a hybrid HTN/operator-based approach to planning. In *Proceedings of the 4th European Conference on Planning: Recent Advances in AI Planning*, 182–194.

[Framinan, Leisten, and García 2014] Framinan, J. M.; Leisten, R.; and García, R. R. 2014. *Manufacturing Scheduling Systems*. Springer London. doi:10.1007/978-1-4471-6272-8.

[Gladwell 2006] Gladwell, M. 2006. *The tipping point: How little things can make a big difference*. Little, Brown.

[Highsmith 2009] Highsmith, J. 2009. *Agile Project Management: Creating Innovative Products*. Agile Software Development Series. Pearson Education.

[Hoffmann, Weber, and Kraft 2012] Hoffmann, J.; Weber, I.; and Kraft, F. M. 2012. SAP speaks PDDL: exploiting a software-engineering model for planning in business process management. *J. Artif. Intell. Res. (JAIR)* 44:587–632. doi:10.1613/jair.3636.

[Kagermann 2015] Kagermann, H. 2015. Change through digitization—value creation in the age of industry 4.0. In Albach, H.; Meffert, H.; Pinkwart, A.; and Reichwald, R., eds., *Management of Permanent Change*. Springer Fachmedien Wiesbaden. 23–45. doi:10.1007/978-3-658-05014-6_2.

[Kambhampati, Mali, and Srivastava 1998] Kambhampati, S.; Mali, A.; and Srivastava, B. 1998. Hybrid planning for partially hierarchical domains. In *Proceedings of the 15th National Conference on Artificial Intelligence*, 882–888.

American Association for Artificial Intelligence (AAAI Press).

[Longo and Kelley 2015] Longo, J., and Kelley, T. M. 2015. Use of github as a platform for open collaboration on text documents. In *Proceedings of the 11th International Symposium on Open Collaboration*, OpenSym '15, 22:1–22:2. New York, NY, USA: ACM. doi:10.1145/2788993.2789838.

[McAllester and Rosenblitt 1991] McAllester, D., and Rosenblitt, D. 1991. Systematic nonlinear planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, 634–639.

[Penberthy and Weld 1992] Penberthy, J. S., and Weld, D. S. 1992. UCPOP: A sound, complete, partial order planner for ADL. In *Proceedings of the third International Conference on Knowledge Representation and Reasoning*, 103–114.

[Rodríguez-Moreno et al. 2007] Rodríguez-Moreno, M. D.; Borrajo, D.; Cesta, A.; and Oddi, A. 2007. Integrating planning and scheduling in workflow domains. *Expert Systems with Applications* 33(2):389–406. http://dx.doi.org/10.1016/j.eswa.2006.05.027.

[Schattenberg 2009] Schattenberg, B. 2009. *Hybrid Planning And Scheduling*. Ph.D. Dissertation, Ulm University, Institute of Artificial Intelligence. urn:nbn:de:bsz:289-vts-68953.

[Schattenberg 2015] Schattenberg, B. 2015. Hybrid planning and scheduling. *KI - Künstliche Intelligenz*. Springer Berlin Heidelberg, doi:10.1007/s13218-015-0390-z.

[Stylianou, Gerasimou, and Andreou 2012] Stylianou, C.; Gerasimou, S.; and Andreou, A. 2012. A novel prototype tool for intelligent software project scheduling and staffing enhanced with personality factors. In *2012 IEEE 24th International Conference on Tools with Artificial Intelligence (ICTAI)*, volume 1, 277–284. doi:10.1109/ICTAI.2012.45.

[Yang 1998] Yang, Q. 1998. *Intelligent Planning. A Decomposition and Abstraction Based Approach*. Springer.