Online Planning: Challenges and Lessons Learned

René Schumann[1] and Jürgen Sauer[2]

[1] Goethe University, Information Systems and Simulation,
Robert-Mayer-Str. 10, 60325 Frankfurt am Main, Germany
reschu@informatik.uni-frankfurt.de

[2] Carl von Ossietzky University, Business Engineering,
Ammerländer Heerstr. 114-118 26129 Oldenburg
juergen.sauer@uni-oldenburg.de

Abstract. Shifting the focus from offline to online planning implies changes that increase complexity tremendously. First of all the optimality can only be proven ex-post. Events that trigger planning actions must be abstracted from primitive data that can be collected, e.g. by RFID bulk scans. Moreover effects like shop floor nervousness, where no stable plan exist, must be avoided. A key aspect in this challenge is the notion of robustness. On the one hand robustness reduces the number of needed repair steps but on the other hand robustness is in conflict with the objective function and can only be computed ex-post. In this article we point out the aspects mentioned above and survey how these issues are tackled in fields of planning and scheduling in artificial intelligence, so far.

1 Introduction

In this paper we discuss specific challenges that arise when dealing with dynamic scenarios in planning and scheduling. Thereby we focus especially on the challenges in transportation planning and shop floor scheduling. Even if these application domains differ the resulting problems are shifting the scope from a static environment to a dynamic one are quite similar. In both cases it is assumed that a plan exists and that the plan execution deviates from the plan. The deviation is typically indicated by an event that arises and changes the environment. If the existing plan is not feasible for the changed environment it has to be repaired, that is typically the task of online planning or reactive planning. In the following the terms online planning and reactive planning / scheduling will be used synonymously. The goals of reactive planning are typically threefold [6]:

1. The reaction to an event should be fast. The current plan should become feasible quickly. This is motivated for two reasons. First the plan execution should not be stopped (to long) until a new valid plan is available. Ideally the execution is going on, while the plan for future activities is adopted. Second in a dynamic environment, like the shop floor or the transportation domain, events are rather frequent. If the system would react to slow the "new" plan is overtaken by reality as it is computed.

2. The existing plan should be widely conserved. This is required as it is intended that steps that have already been performed, should be considered in the new plan, if possible. So work done so far should not be discarded. Moreover this enforces plan stability that is important for the coordination of the planning with other planning problems, like procurement, and for the reputation of the planning systems.
3. The third goal is to maintain the plan's quality. A key motivation for using enhanced planning systems is, of course, the ability to compute plans of good quality in a short time. Lower evaluated plans lead to inefficient transportation or production and in consequence to operational loss of the company. It goes without saying, that this has to be avoided. Therefore to be practical applicable planning systems have to perform at least acceptable and comparable to human planners in dynamic environments and this is measured in plan quality.

As mentioned above it becomes clear now, that these requirements are partially contradicting. Computing a good plan, according to an objective function, and having a fast reaction to an event is obviously hard to achieve. And conserving large parts of a plan while maintaining its quality can become hard, as well.

Currently the research about online planning, even if its established for some years (see [21] for a review of early work in transportation or [29] for job shop scheduling) is still in an initial constructive phase of research. Approaches based on techniques like local search [29, 23], metaheuristics [7, 17], multiagent systems [8, 12, 16] or hybrid approaches [31] are presented. In consequence a variety of different approaches exist for online planning.

If one have to select a technique for a given problem the selection of one approach from given set of techniques becomes a hard task. This is because of different scopes of what online planning compromises, a lack of benchmarks and metrics.

In fact what is needed are means that can help to compare different online planning approaches. This comparison should be a multidimensional one. It is not enough, and in dynamic environment typically not feasible, to find that method a is superior to method b. Which method performs better typically depends on characteristics of the application domain and on characteristics of the changes within the environment. Thus for a more engineering oriented approach it would be necessary to identify those characteristics and evaluate the abilities of existing methods according to them.

In the following we are first clarifying the terminology used in different fields. This is especially important as different terms are used interchangeable even if they may have different semantics. In section 3 we discuss current challenges that have to be faced towards a more engineering approach for the design and application of online planning techniques. Finally we summarize our hypothesis and outline future work.

2 Existing terminology and techniques

2.1 Terminology

The problem of keeping existing plans in a dynamic environment valid is addressed in different fields of research and applications. The fields addressed here are mainly operations research and artificial intelligence. The applications mentioned here are vehicle routing problems and (job-shop) scheduling problems. Different terminologies have been established and terms are used partially interchangeable. Therefore it is necessary to clarify the terminology. As already mentioned the terms online planning and reactive planning are used here synonymously. It can be argued that online planning is the more general term, as it compromises least commitment planning. Such approaches were presented e.g. by [20] and [27]. Interestingly methods based on least commitment strategies do not play an important role in current research, even if they are used in practical applications [9]. But for the purpose of this paper we exclude this field and use those terms interchangeable.

Technically there exist two major approaches how reactive planning can be implemented [10]:

- Plan repair: An existing plan is going to be adopted to a changed situation. Approaches are presented in [29, 2].
- Replanning: If an event occurs the existing plan is discarded and a new plan is computed from scratch [5].

Both approaches have their strengthes and weaknesses. Especially concerning the goals of reactive planning, mentioned above. From the perspective of complexity it has been shown that both face the same structure problems [15]. A comparison emphasizing stability can be found in [10].

It has to be mentioned that this terminology outlined here is not exclusiv. There are different naming schemes in use. For example in [33] the most general term used is rescheduling. The authors present an interesting framework for rescheduling approaches, that can be classified in their framework. Thereby they distinguish between dynamic scheduling and predictive-reactive scheduling. Dynamic scheduling is characterized where no plan exist a priori, like in least commitment strategies. Furthermore the authors mention that other names for dynamic scheduling are online scheduling or reactive scheduling.

2.2 Techniques for online planning

As previously outlined we can identify different techniques for online planning. Therefore we can identify different technologies for plan repair and replanning, even if there exist technologies that can be applied for both approaches, e.g. metaheuristics.

In the following we are going to discuss techniques that have been applied to plan repair. Approaches based on local search techniques apply plan modifications, like shifting dispatched operations in the future (right shift) or pull them backwards in time (left shift). A typical approach using those techniques is implemented in the OPIS system presented in [29].
Based on local search techniques, metaheuristics can be applied to plan repair as well. A counterargument therefore is the longer reaction time. But it has to be stated, that metaheuristics strive fast to good solutions, so that computation can be interrupted at a given time and results can be of good quality. A typical trend of plan's quality over time is sketched in figure 1. An early approach using tabu-search for instance is [7].

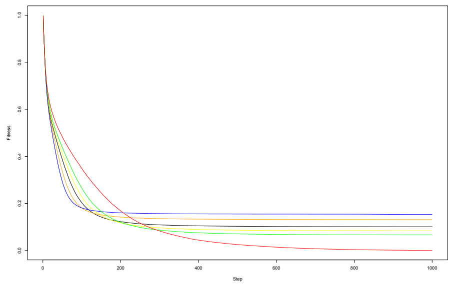

Fig. 1. Characteristic plan's quality trend over time, using different parameter sets

Other approaches, that has been applied to plan repair are multiagent systems. Using the characteristics that agents can adopt their behavior and their plans to their environment. Often orders and resources are represented by agents. Agents negotiate about the assignments, while typically orders want to be processed as quickly as possible while resources want to maximize their utilization. Examples of agents in dynamic environments in the manufacturing domain can be found in [4, 18, 25, 34]. Application in the transportation domain can be found in [8, 20] for example.
Of course hybride approaches are possible as well. Where techniques from multiagent systems and evolutionary approaches are combined, e.g. [30].

For the field of replanning different options are possible, and applied, as well. A complete replanning compromises the generation of a new plan for all non executed activities. In contrast in partial replanning only a subset of activities is rescheduled. Other parts of the plan remain fixed and needed resources are marked as blocked. After a new schedule has been computed, which can be done faster as the initial computation, because the number of activities to be scheduled have been reduced, the unchanged elements of the plan and the newly computed plan can be joined. The quite common technique to remove the orders that are

affected by an event, update those data and the integrate those events as new orders thus is a special case of partial replanning. According to [19] complete replanning leads to better plan's quality than partial replanning.
Of course looking at aspects, like plan stability, complete replanning is problematic, as the plan can change with each event completely. One can try to soften those changes by adding new constraints to the planning problem, that fixes parts of the plan in favor of stability. This decreases typically the ability to find feasible plans at all, because it is unknown to what degree similarity can be archived at all.

Another approach that is currently on its way into practical application is based on a less intelligent or optimizing way, but is a more interactive approach. Simulation techniques integrated in manufacturing control stations, typically incorporated in manufacturing execution systems (MES), allow the human operator to evaluate possible reaction options quickly. Sadly the aspect of interactivity is often ignored in research done in operational research and artificial intelligence. But it has to be mentioned that it is an important aspect to gain confidence in the system and its reaction to disturbing events.

Another discussion about applied methods for reactive planning with further references can be found e.g. in [19].

3 Challenges in online planning

In this section we outline a number of points that arise in planning in dynamic environments independently from the terminology and nearly independently from the application domain.

3.1 Optimizing without knowing where to go

A first fact that has to be admitted is that optimizing in a dynamic environment is not possible. The optimal solution can only be computed ex-post. Thus during the execution finding the optimal solution is comparable to guessing. This is of course a formulation that exaggerate the actual state. But it is in fact true that a lot of plan repair strategies, especially those who try to optimize the new generated plan, rely on the ceteris paribus assumption. That is that the existence of events in the future is neglected. This is for nearly all events (expect the last one) not true. And therefore the resulting plan is typically suboptimal. This makes a perspective shift necessary, especially from the classical perspective of operations research, here one typically wants to find the optimal solution.

3.2 What is dynamic and how can it be recognized?

In nearly all models for reactive planning the events are triggered by events, that change the environment, and may invalidate the current plan[3]. Those events are often abstractions from data collected in reality. Data collection is nowadays widely automated. Techniques based on GPS or RFID are used. For instance a bulk scan of RFID tagged products at the outgoing gate of a crossdocking station can only be interpreted as a certain number of products, leaving a certain place at a certain time. The conclusion that an order is late, is an abstraction of those information, that can be derived by data collected by the bulk scan and further context information. What is needed is a vertical data integration [26]. Information gathered by PDA (production data acquisition) systems or corresponding systems in the transportation domain is very fine grained. A hierarchy of different abstraction levels have to be derived regarding the context of the given production situation at hand.
As it becomes clear so far, the reactive planning systems have to be integrated in the existing IT structure of companies. Especially the integration with the aforementioned PDA systems and a sound abstraction of data collected by this system is a necessity to implement an applicable reactive planning system. Because otherwise deviations in the execution from a given plan could not be identified. On the other hand specialized triggers have to be realized that allow an integration within the existing ERP system. Because those systems are another source of dynamic information. Information concerning new, changed or canceled orders or resources are typically maintained in the ERP system.
Another very promising approach is complex event processing, which is designed to automated process a high number of events, modeled as an event-flow, with up to some 100.000 events per second, as it is generated by systems based on PDA or RFID techniques [14]. In complex event processing those raw, fine-grained events can be aggregated, efficiently.

3.3 React but not overreact

Assume we can surely identify events that change the environment. An event can either have the potential to influence the plan in a positive or negative way. Note that an event does not have to effect a plan at all. For instance a resource break down of a resource that is not scheduled in the current plan will not effect the plan at all; even if the environment has changed.

An example of a negative effect on an existing plan is the arrival of a new order. Typically it is encoded as a hard constraint that all existing orders have to be concerned in the plan. Thus this event has obviously negative effects to the plan, as it becomes invalid. Nevertheless, the new plan might be of lower quality, as well.

[3] This is not the only way to handle plan adaption in dynamic environments, of course. In their framework Vieira et al. [33] classify the approaches into event-driven or periodic. In a periodic approach the plan is regularly updated regarding a rolling time horizon.

An example of an event that can have positive effects on a plan is a withdraw of an existing order. There are at least three options how to react on this event.

- Ignore the event completely and do not change the plan, at all. Of course there are unnecessary actions within the plan, but commonly the plan will not be decreased[4].
- Delete the assignments needed to fulfill the withdrawn order. Thereby you can not decrease your plan. It will become eventually better, if the fulfillment of this order has evaluated negatively itself, i.e. the fulfillment of this order was delayed.
- Delete the assignments needed to fulfill the withdrawn order. And evaluate if other assignments in the existing plan can be moved to improve the plan's quality. Thereby trying to use the freed resources of the withdrawn order to improve the fulfillment of the remaining orders.

Actually it is an open question to what events an online planning system has to react. One would expect that online planning reacts when a plan becomes invalid. But a literature review shows that this is not common ground, so far. For example in [13] only the event of new customers is regarded, which is not the only event that can invalidate a plan. The set of regarded events in transportation planning was extended e.g. in [28] to at least five different events. The number of events that can occur is even higher in the job shop scheduling domain, for example Reheja et al. [22] list 17 different events that can occur in their model. It is arguable only to react to negative events, as the plan quality is typically decreased by those events. And the potentials to improve a plan, as a result of other events is not used.

Reacting to events means typically change the current plan, which can decrease the stability of a plan. It is only *can* here, because a plan adoption might save future adjustments of the plan, as a consequence of upcoming events in the future, which depends on the characteristics of future events and the event handling strategy.
Current discussions on plan stability can be found in [10, 19].

3.4 Robustness: way out or digging in?

Closely related to plan stability is the term of robustness or robust plans. Robustness is commonly seen as the means to achieve plan stability. A plan is called robust if it "is likely to remain valid under a wide variety of different types of disturbance" [7]. Robustness is typically achieved by allowing the potentials to incorporate upcoming events without changing the already planned activities. A typical technique therefore is to use buffers or slacks in the plan, reduce resource capacity or increase operation durations artificially in the original plan. Of course this means all decrease typically plan quality. Therefore a tradeoff

[4] Expect you have some contribution margin elements in your objective function, which is rather infrequent.

between plan's quality and robustness has to be made. Finding a tradeoff is complicated by the fact that robustness cannot easily be measured.

Whether a plan is robust or not depends heavily on the kinds and sequence of the occurring events. There can exist a sequence of events that can all be incorporate without changing the plan even one time, while another sequence of events might cause plan repair after each event, starting with exactly the same plan. Thus measuring robustness depends essentially on the events that are assumed to occur. As this is typically unknown different distributions of event sequences have to be explored. Thereby the sequence can be changed according to the

- different kinds of events that can occur,
- the percentage rate for each kind of event,
- the ordering of the events and
- the time when the events occur.

Moreover robustness and plan stability do not only depend on the sequence of events but on the current situation within the execution system, e.g. the shop floor, as well [24].
So the degrees of freedom while measuring robustness and consequently plan stability are comparable high and it becomes hard to find reasonable metrics. And even if there exist such a metric that indicates that a given plan is robust. There is always the potential that while executing this plan a lot of adoption becomes necessary caused by an event sequence that was assumed rather unlikely during the evaluation of the plan.
Speaking about robustness without indicating towards which events is misleading, too. As there most often exist a sequence of events that will result in an instable plan. What can be achieved is robustness against a certain type of event or defined event mixes.

3.5 Evaluation: Or why my system is always best

As it becomes clear so far, evaluating online scheduling systems is a hard task to accomplish. Thereby nearly every presented approach is evaluated. But most often this is done using a specialized scenario that mainly focus on specific characteristics of the presented method. With the result that most approaches claim to have some unique advantages in comparison to other existing methods.

Surveying existing papers about topics like online planning or reactive planning shows that only a few people deal with a more systematic approach to compare and evaluate existing approaches e.g. [19, 20, 24, 32]. Technically all those approaches base on simulation as primary methodology for evaluation of online planning systems (see also [1, 6, 19, 20, 32]). In doing so the planning system is used within a simulation environment that simulates the dynamic environment and rises the events. Thereby it has to be mentioned that even if the systematic is commonly agreed on, there exist no tool or system that is dominate in use. Even if there exist suggestions of testbeds [20, 32] those are not widespread either.

At least four objective reasons can be identified for the lack of work that characterize and evaluate online planning systems.

- Online planning has been identified as important. A consequence of the terminological mess, the different fields of research that are interested in this topic and the different application domains it is hard to survey relevant work, and maybe find appropriate existing methodologies.
- There exist no common ground on what online planning typically comprises. So which set of events should be handled, which functionalities an online planning tool should offer. This is at least one requirement to compare different approaches. As mentioned before the number of events that can be handled by different approaches vary from one to at least 17.
- A consequence of the aforementioned aspects is, that there exist no commonly accepted benchmark scenarios. And as outlined before a lot of approaches are not strive to be comparable or applicable to scenarios already existing in the literature, as they highlight to handle situation that are special in their scenario. Another challenging aspect is the high degree of freedoms that already have been mentioned by modifying a sequence of events. Expect using random functions there exist rarely systematic testing approaches for generating event sequences. Thereby those sequences may have some random elements but on the other hand allow exact replications of simulation runs that are needed to compare different approaches in identical situations.
- We have a lack of metrics. We have currently no adequate metrics that are commonly accepted either for
  - dynamic itself [13, 28, 11],
  - stability [19] or
  - robustness

  Metrics based on a competitive analysis [3] require a comparison to the optimal solution. But in most cases of job shop scheduling or transportation planning those optimal solutions are unknown and cannot be computed efficiently. Thus those analysis are only of theoretical interest.

4 Discussion

In this paper we have addressed challenges that currently arise in online transportation planning. We identified similarities between these planning problems and problems arising in dynamic scheduling domains. It has to be pointed out that planning in dynamic environment has a common ground, that have to be explored in more depth. So far research is in an early stage. More work seems to be investigated in finding of new methods than in comparing and evaluating existing techniques. As a reason for this we state that important fundamentals have to be investigated before a reasonable evaluation or comparison of online planning can be achieved.

Our hypothesis is that a common understanding of what online planning should comprise, common benchmarks and widely accepted metrics are needed.

This could lead to a more engineering oriented catalog and characterization of existing and upcoming online planning approaches. Thereby the evaluation of methods will be regarding different aspects. It is not enough, and in dynamic environment typically not feasible, to find that method a is superior to method b. Which method performs better typically depends on characteristics by the application domain and on characteristics of the changes within the environment. Therefore for a more engineering oriented approach it would be necessary to identify those characteristic and evaluate the abilities of existing methods according to those characteristics. A step towards that approach is for example [24].

Aspects discussed in this paper would lead to a more coherent, sound and although more engineering view of the field of online planning. Thereby different fields of research are involved investigating general challenges in online planning.

References

1. Christoph Aigner, Jürgen Dorn, and Mario Girsch. Cooperation between a simulation environment and a reactive scheduling system,. In *Proceedings of the Eighth International Conference on Manufacturing Engineering (CME 2000)*, Sydney Australia, 2000.
2. Marlene Arangú, Antonio Garrido, and Eva Onaindia. A general technique for plan repair. In *Proceedings of the 20th IEEE International Conference on Tools with Artificial Intelligence*, volume 1, pages 515–518, Dayton, Ohio,USA, 2008. IEEE Computer Society.
3. Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis.* Cambridge Univ. Press,, 1998.
4. Sergio Cavalieri, Marco Garetti, Marco Macchi, and Marco Taisch. An experimental benchmarking of two multi-agent architectures for production scheduling and control. *Computers in Industry*, 43:139 – 152, 2000.
5. William Cushing and Subbarao Kambhampati. Replanning: a new perspective. In *Poster Program. In Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling (ICAPS-05)*, Montery, 2005.
6. Jürgen Dorn. Evaluating reactive scheduling systems. In *IAT '04: Proceedings of the Intelligent Agent Technology, IEEE/WIC/ACM International Conference on (IAT'04)*, pages 458–461. IEEE Computer Society, 2004.
7. Jürgen Dorn, Roger Kerr, and Gabi Thalhammer. Reactive scheduling: Improving the robustness of schedules and restricting the effects of shop floor disturbances by fuzzy reasoning. *International Journal on Human-Computer Studies*, 42(6):687 – 704, 1995.
8. Klaus Fischer, Jorg P. Muller, Markus Pischel, and Darius Schier. A model for cooperative transportation scheduling. In Victor R. Lesser and Les Gasser, editors, *Proceedings of the First International Conference on Multiagent Systems (ICMAS)*, pages 109–116, San Francisco, California, USA, 1995. The MIT Press.
9. Truls Flatberg, Geir Hasle, Oddvar Kloster, Eivind J. Nilssen, and Atle Riise. Dynamic and stochastic vehicle routing in practice. In Vasileios Zeimpekis, Christos D. Tarantilis, George M. Giaglis, and Ioannis Minis, editors, *Dynamic Fleet Management: Concepts, Systems, Algorithms & Case Studies*, Operations Research / Computer Science Interfaces Series, pages 41 – 63. Springer, 2007.

10. Maria Fox, Alfonso Gerevini, Derek Long, and Ivan Serina. Plan stability: Replanning versus plan repair. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proceedings of the Sixteenth International Conference on Automated Planning and Scheduling (ICAPS 2006)*, pages 212–221, Cumbria, UK, 2006. AAAI Press.
11. G.D.M. Frizelle, D. McFarlane, and L. Bongaerts. Disturbance measurement in manufacturing production systems. In *Advanced Summer Institute (ASI 98)*, Bremen, Germany, 1998.
12. Herwig Henseler. *Aktive Ablaufplanung mit Multi-Agenten*, volume 180 of *DISKI*. infix Verlag, 1998.
13. Allan Larsen. *The Dynamic Vehicle Routing Problem*. PhD thesis, Technical University of Denmark (DTU), 2001.
14. David Luckham. *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*. Addison-Wesley Longman, Amsterdam, 2002.
15. Bernhard Nebel and Jana Koehler. Plan reuses versus plan generation: A theoretical and empirical analysis. *Artificial Intelligence*, 75(1-2):427–454, 1995.
16. Djamila Ouelhadj, Sanja Petrovic, Peter I. Cowling, and Amnon Meisels. Inter-agent cooperation and communication for agent-based robust dynamic scheduling in steel production. *Advanced Engineering Informatics*, 18(3):161–172, 2004.
17. Giselher Pankratz. Dynamic vehicle routing by means of a genetic algorithm. *International Journal of Physical Distribution & Logistics Management*, 35(5):362–383, 2005.
18. Y. Peng, Tim Finin, Yannis Labrou, R. Cost, B. Chu, J. Long, W. Tolone, and A. Boughannam. An agent-based approach for manufacturing integration - the ciimplex experience. *International Journal of Applied Artificial Intelligence*, 13(1-2):39 – 64, 1999.
19. András Pfeiffer, Botond Kádár, and László Monostri. Stability-oriented evaluation of rescheduling strategies, by using simulation. *Computers in Industry*, 58:630–643, 2007.
20. Martha E. Pollack. Planning in Dynamic Environments: The DIPART System. In A. Tate, editor, *Advanced Planning Technology: Technology Achievements of the ARPA/Rome Laboratory Planning Initiative*, pages 218 – 225. AAAI Press, 1996.
21. Warren B. Powell. Dynamic models of transportation operations. In A.G. de Kok and S.C. Graves, editors, *Handbooks in Operations Research and Management Science: Supply Chain Management: Design, Coordination and Operation*, volume 11 of *Handbools in Operations Research and Management Science*. Elsevier Publishing Company, 2003.
22. Amritpal Singh Raheja, K. Rama Bhupal Reddy, and Velusamy Subramaniam. A generic mechanism for repairing job shop schedules. *Innovation in Manufacturing Systems and Technology (IMST)*, 2003.
23. Andreas Richter. *Dynamische Tourenplanung: Modifikation von klassischen Heuristiken für das Dynamische Rundreiseproblem (DTSP) und das Dynamische Tourenplanungsproblem (DVRP) mit der Möglichkeit der Änderung des aktuellen Fahrzeugzuges*. PhD thesis, Technischen Universität Dresden, 2005.
24. I. Sabuncuoglu and M. Bayõz. Analysis of reactive scheduling problems in a job shop environment. *EJORS European Journal of Operational Research*, 126:567 – 586, 2000.
25. Tuomas Sandholm and Victor R. Lesser. On automated contracting in multi-enterprise manufacturing. In *Proceedings of Improving Manufacturing Performance*

in a Distributed Enterprise: Advanced Systems and Tools, pages 33–42, Edinburgh, Scotland, 1995.

26. Jürgen Sauer. Vertical data integration for reactive scheduling. *Künstliche Intelligenz*, 2009, to appear.
27. John Sauter and H. van Dyke Parunak. Ants in the supply chain. In *Workshop on Agent based Decision Support for Managing the Internet-Enabled Supply Chain, Agents 99,*, Seattle, 1999.
28. René Schumann, Thomas Timmermann, and Ingo J. Timm. Transportation planning in dynamic environments. In Bernhard Fleischmann, Karl-Heinz Borgwardt, Robert Klein, and Axel Tuma, editors, *Operations Research Proceedings 2008*, Operations Research Proceedings, pages 319 – 324, Augsburg, 2009. Springer.
29. Stephen F. Smith. Opis: A methodology and architecture for reactive scheduling. In Monte Zweben and Mark S. Fox, editors, *Intelligent Scheduling*, pages 29 – 66. Morgan Kaufmann Publishers, San Francisco, 1994.
30. Paul Valckenaers and Hendrik Van Brussel. Holonic manufacturing execution systems. *CIRP annals International Institution for Production Engineering Research annals*, 2006.
31. Paul Verstraete and Paul Valckenaers. Towards cooperating planning and manufacturing execution systems. In *INCOM*, 2006.
32. Paul Verstraete, Paul Valckenaers, Hendrik Van Brussel, Karuna Hadeli, and Bart Saint Germain. Multi-agent coordination and control testbed for planning and scheduling strategies. In *Fifth international joint conference on Autonomous agents and multiagent systems*, pages 1451 – 1452, Hakodate, Japan, 2006.
33. Guilherme E. Vieira, Jeffrey W. Herrmann, and Edward Lin. Rescheduling manufacturing systems: A framework of strategies, policies, and methods. *Journal of Scheduling*, 6(1):39–62, 2003.
34. Jan Wörner and Heinz Wörn. Benchmarking of multiagent systems in a production planning and control environment. In Stefan Kirn, Otthein Herzog, Peter Lockemann, and Otto Spaniol, editors, *Multiagent Engineering: Theory and Applications in Enterprises*, International Handbook on Information Systems, pages 115 –134. Springer, Berlin et al, 2006.