Pick Again: Self-Adaptive Warehouse Commissioning

Oliver Meisch, Gerben Peet
Kramp BV
{oliver.meisch, gerben.peet}
@kramp.com

Stefan Rudolph, Jörg Hähner
Universität Ausgburg
{stefan.rudolph, joerg.haehner}
@informatik.uni-augsburg.de

Sebastian von Mammen
Universität Würzburg
sebastian.von.mammen
@uni-wuerzburg.de

***Abstract*—Picking in warehouses represents a central process for mail order businesses. It has a major impact on the required investments by the business and the speed of delivery to the clients. In fact, due to the steady rise in expectations regarding delivery times - even private households receive goods on a same-day-delivery basis, nowadays - ongoing optimisation efforts are rather important. They are all the more necessary, as the economic environment changes rapidly, which includes emerging or receding trading channels as well as the generation of vast amounts of sales and usage data. These dynamics, in turn, equally challenge traditional optimisation approaches and monolithic IT systems. Therefore, in this collaborative project between academia and industry, we have spatially re-modelled an existing small parts store, modelled the workers, dollies, and goods, and optimised various aspects such as the chosen routes and storage locations of the goods. Altogether, we achieved a significant increase in efficiency: The picking process yielded 14% more picks per time and the walking distance was reduced by 37%. Our agent-based, and Q-learning-based approach lends itself well for adapting to changes in the environment as well as changes in the clients' shopping behaviours.**

I. Introduction

Traders in the mail order business can have a major impact on on the timespan between the reception of an order and delivery of the goods by strategically designing the commissioning process. If the shipment procedure is outsourced, the only remaining steps for the trader are preparing the goods for transport and passing them to the parcel service. In this rather common case, the commissioning process becomes the most important factor for reducing delivery times.

Nowadays, general "in-night" deliveries, i.e. deliveries on the next day (sometimes als referred to as "over-night" deliveries), are considered a high-quality asset in the business to customer (B2C) segment. In the business to business (B2B) segment, i.e. the trade between companies, "in-night" deliveries are crucial to a company's success—delays in delivery can result in delays of whole supply chains.

At the same time, companies need to adapt to increasingly dynamic environments. Especially mail order businesses are exposed to strong seasonal variations of commissioned goods. As goods and suppliers keep changing, the determining factors for the commissioning work change as well. Even the purchase order behaviour by the clients has changed dramatically towards high frequencies of small orders as opposed to rare orders of large numbers of items. In addition, the customers' ties to individual companies have waned, which further the importance of customer satisfaction with respect to once business' success.

In this paper, we present an approach to optimise the picking process in a small parts store. In particular, we created a spatial, agent-based model of the warehouse, implemented various heuristic and optimal schemes for determining picking paths, and first learned which scheme to apply in which situation. Next, we tested different storage locations for different goods, and measured and optimised their impact on the picking process; again, we resorted to the optimised path strategy, in order to achieve a net benefit. In addition, we learn how to optimally utilise the concept of selecting certain commissioning items before the picking process starts. Our optimisation efforts are concluded with a combination of all of the mentioned approaches which yielded the highes efficiency gains in our simulation.

The remainder of this paper is structured as follows. In the next section, we briefly touch upon aspects of agent-based modelling and organic computing that served as a basis for our work. We also reference works that link these fields to supply chain management. In Section III, we present the warehouse and the picking process. In Section IV and systematically analyse the optimisation potential and unfold the approach taken. Concrete experimental results are presented and discussed in Section V. We conclude with a summary and an outlook on potential future work.

II. Related Work

Multi-agent-systems (MAS) are an established field of research [1]. They enable us to quickly develop and compute multi-factorial and multi-scale models that are otherwise hard to solve. They allow us to model, simulate and optimise complex system architectures [2] or dynamic networked systems [3]. MAS also play a central role in the field of Organic Computing which translated features of natural systems, such as robustness, flexibility or scalability, to technical systems [4]. In the context of technical systems, these features are referred to as self-x properties and they included, for instance, self-calibration, self-optimisation, self-explanation, self-healing, and, most foundational, self-organisation [5], [6]. In [7], an Organic Computing system for the control of traffic flow is presented. The link from self-adapting traffic control and traffic management systems is closely aligned with optimising transport and logistic processes, as for instance emphasised by [8]. More specific works that link agent-based models to the domain of supply chain management stress how artificial intelligence approaches can improve globally defined system metrics [9]. It has been demonstrated on numerous occasions how the bridge to practical implementations can be realised. For instance, in [10], a self-controlling system is presented in which software agents are integrated with a company's infrastructure and conduct transactions across several businesses. Their ability to self-control allows the system to adapt to changing software requirements. Another example for the rigorous integration of agent-based optimisation models can be found in the Agent.Enterprise project [11]. Here, a set of multi-agent systems is interlinked hierarchically to modularly implement non-functional as well as functional aspects of a supply chain management system.

III. Model

The warehouse that we modelled spans across three floors. The floors are connected by stairs and conveyor belts that carry reusable transport boxes for the stored goods. Each floor is 87m wide and 44m long and divided into six sections, which we will detail in the next section. The warehouse is used as a small parts store and workers need to pick up certain goods/parts at fixed storage locations. During the first work shift, from 4am to 12pm, new goods arrive at the receiving area and are stored away. The storage logistics follow the so-called chaotic inventory system, which does not foresee a fixed location assignment for specific goods [12]. The reason for this dynamic allocation of locations lies in the seasonality of the goods that are stored in the specific warehouse that we investigated—the commissioned parts are technological components or systems of the agricultural industry. A certain rule set ensures that certain parts with sufficient availability are stored in several stations distributed across several floors. These rules aim at balancing the workload across all workers as well as all the materials-handling technology. During the second shift, from 1pm to 9pm, the incoming orders are processed by 25 to 54 workers. The orders are processed in the order of their scheduled time of departure, which depends on the distance to the goal region of the client.

A. Sections

All 18 sections (6 per floor) follow the same design. At the lower end, hoisting and conveying technology is setup. There are two stations for each conveyor belt: One for regular orders and another one for express orders. The warehouse management channels transport boxes to the respective stations based on their respective priorities. There is a aisle measuring 1.2m in width that runs along the material-handling strip at the bottom of the floor (a schematic figure is shown in Figure 1). There are 18 shelves expanding orthogonally from the aisle across the whole warehouse. As pairs of shelves lean on each other, there are 9 lanes (width 1.04m) protruding from the aisle. The lanes are interrupted at the centre of the warehouse by a second aisle that is 3.3m wide, and by another aisle close to the upper perimeter that is 1.8m wide. The shelves are divided into up to 31 modules, each 1.2m-wide and 30cm deep. At the end wall, there is a module which is assigned to the next shelf of even number.

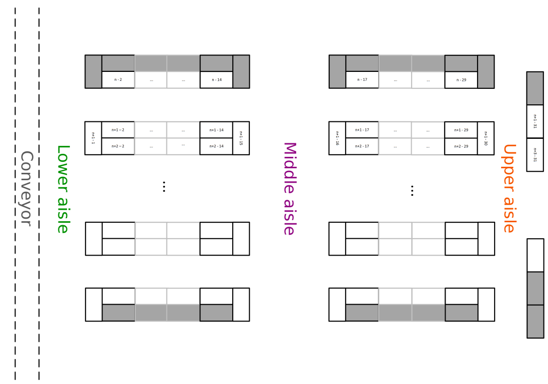

Fig. 1. Schematic representation of a station.

2.5 workers can work at each station: Two permanent workers and one who moves between stations in accordance with the workload. Empirical evidence has shown that greater numbers of workers at the stations tend to slow down the picking processes due to congested aisles and lanes.

B. Commissioning Performance

At the company we are looking at, the commissioning performance p is measured in number of picks n per time t (Eqn. 1). p in combination with the amount of required human labour determines the amount of available operating resources. This approach to measuring commissioning performance is common in practice and well-researched in academia [13]. p can be used as a measure for individual workers, stations, floors or for capturing the performance of the whole warehouse.

$$p(n,t) = \frac{n}{t} \tag{1}$$

We consider an hour the time measure of reference. Picks are considered only by the number. There is no distinction, for instance, between the volume or mass of the picked goods, their storage height, or additionally required activities including unwrapping or boxing of goods. These factors could be considered to achieve a more precise calculation of the actual performance. However, in the warehouse we investigated, the simple performance measure from Eqn. 1 represents the established way for determining the system behaviour, which we set out to improve.

IV. Analyses, Representation & Optimisation

Walking from one module to the next takes the longest time during the commissioning process. At the same time, we expected the greatest optimisation potential in this activity. Therefore, we analysed the established routes chosen by the workers and presented an according optimisation concept.

In general, there are two levers for optimisation: One can optimise (1) the chosen route for a given set of picking locations that need to be traversed, or (2) the arrangement of goods to result in shorter paths (on average) to begin with. Close proximity between all the modules that need to be visited during one picking tour reduces the overall distance, on average. If a worker can select a set of picking tasks, he can minimise the distance accordingly.

A. *Established Routing*

For each module, a value o is determined which reflects the distance required to take a good out of the warehouse. In order to calculate o, one needs to consider the lane l of shelf s, the module number m, and its face f. The face is encoded as 1/2 for modules to the right/left in a lane, starting from the lower aisle. The numbers of aisles, shelves, modules are incremented from the lower aisle, which automatically yields a greater value of o with greater indices. An according equation of o, considering the numbering scheme for the investigated warehouse, is provided in Eqn. 2.

$$o(l, m, s, f) = 10^5 * l + 10^4 * m + 10 * s + f \qquad (2)$$

In the established approach all the picks of a tour would be communicated to the worker as a list ordered with increasing o. Equation 2 was designed to reflect the relative distances of a high-rise storage setup, which is why it does not seamlessly translate to the investigated small parts store. As one can see in Figure 2, it does not, for instance, account for aisles at the centre or back of the warehouse, which make it possible to switch from one lane to another without returning to the starting aisle.

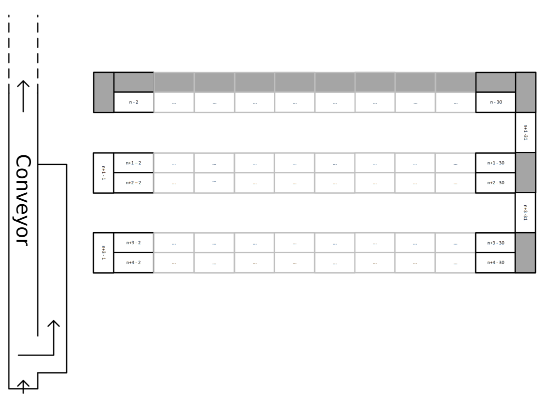

Fig. 2. Originally assumed layout of the warehouse.

This originally assumed layout does not leave much leeway for optimisation: One has to enter all the lanes with prospective picks and follow them until the last good of the respective lane has been picked. However, this approach and its associated o calculation are not applicable to the actual warehouse layout, which results in unnecessarily longer paths with increasing numbers of picks. Therefore, a first optimisation step in the scope of our work, was to identify better or optimal routes.

B. *Established Storage Locations*

As we mentioned above, the investigated warehouse implements a chaotic inventory system, which allows to assign goods to changing locations in order to compensate for seasonal fluctuations. Items that might be in high demand during the summer time might be stored in up to 40 storage units and occupy only one during winter. An important reason to spread articles across several storage units lies in the width of the lanes which limits the number of workers at one station to 2.5. In order to compensate for this constraint, it is a goal to maintain approximately the same workload at all the stations and to make a specific part commissionable from at least two stations. This is currently achieved by evenly distributing the goods across all stations and across the whole area. This distribution, in turn, results in picking tours covering the whole area throughout a single shift. Figure 3 visualises the access frequency (higher frequency ~darker hue) across one floor over the course of one shift. One clearly sees that the most frequently picked modules are scattered across the whole floor.

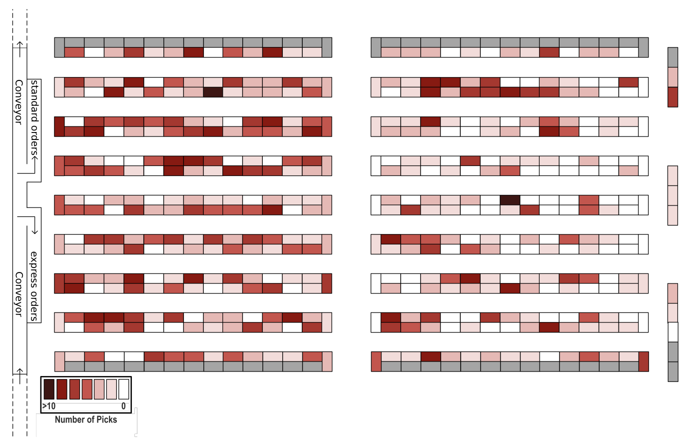

Fig. 3. Distribution of picks over the course of one shift (darker red hue means higher frequency).

In order to visualise the scattering of picks that result from this distribution for a single tour, Figure 4 exemplarily highlights all the visited modules of one tour in blue, emphasising the most frequented module in red. Each shelve was visited at least once.

C. *Warehouse Representation*

As the workers are always assigned to one station or floor, we only require a two-dimensional representation of the warehouse. Next, we assumed 90^o angles between construction elements and furniture to facilitate mapping the warehouse to Cartesian coordinates and to simplify the path calculations. Based on the width of the smallest infrastructural element, the module, one unit in our coordinate system translates to 1.2m.

D. *Optimisation of Routes*

In order to determine optimal routes, we rely on two algorithms. First, we calculate the shortest distances between

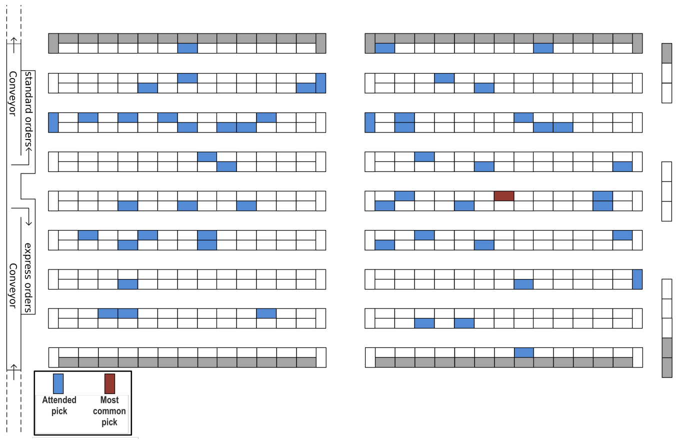

Fig. 4. Modules picked from over the course of one tour (in blue). In red: The module with the highest total access frequency.

all pairs of locations that we need to visit by means of the A* algorithm [14]. Second, we determine the best possible route by means of a branch-and-bound approach [15]. It starts with a nearest neighbour heuristic and determines the lower bounds a 1-trees [16]in accordance with Prim's algorithm [17].

E. Machine Learning Approach

We apply Q-learning [18] to optimise our simulated agents' behaviours. The high degree of similarity between the situations faced by the agents allows this algorithm to converge to optimal decisions in static environments. At the same time, its online capability, i.e. its realisation of iterative improvement, prepares the stage for a self-adaptive system that can handle variations of client habits and market fluctuations at different time scales.

However, for our early results, we distinguished between a learning and an application phase in order to directly quantify the learning effect over time. Learning rate α and discount factor γ have a major impact on the learning speed and accuracy. During the learning phases, we started out with a learning rate $\alpha = 1$, which means that only new information is considered, which is continuously reduced to $\alpha = 0.1$, which means new information only marginally impacts the agents' behaviour. We implemented a constant discount factor of $\gamma = 0.9$, which accounts for preceding actions' impact on a result. Instead of selecting the action with the highest Q-value, we apply roulette wheel selection and select an action with a probability proportional to its contribution to the fitness value in a given state [19]. We calculate this probability p for a specific action a by means of Equation 3, whereas f_a is the Q-value or fitness of a and N is the numbers of all possible actions in a given state.

$$p_a = f_a / \sum_{j=1}^{N} f_j \tag{3}$$

V. Experiments & Results

The calibration of the agent-based model is an important first step. Our goal is the simulation of a regular work day that matches empirical data, without any optimisations applied. Such parameterised and validated model ensures that our simulation results be translated to real-world scenarios and it provides a basis for comparisons of our optimisation efforts.

We detail our empirical methodology and its application for calibrating our simulation model in the next section. Afterwards, we step through the individual optimisation experiments, explain their rationale, and present and discuss their results.

A. Simulating a Regular Work Day

Our empirical data is based on sampling picking processes over a period of one month (June 2016). We logged the raw picking information, providing information on the good, the time, and location of a pick. In addition, we calculated and logged inferred values such as the walked distance, the total work time, and the performance measure p at a granularity of one hour. In order to validate our trained model, we considered the first three weeks of the month for learning and the remaining week for validation only.

Simulating a regular work day, our trained model yielded a performance of 57.45 picks per hour. In total, the workers walked 480.5km and worked for 248 hours and 14 minutes. These values closely correspond with our empirical data.

B. Optimal Routes

As explained above, the original routing method in place was misguided. Based on our model that also considers the centre and back aisles as well as rather accurate spatial dimensions, we ran A* in combination with a branch-and-bound approach to calculate shortest paths. With this first change to the original model, we achieved the performance increases captured in Table I. The total distance walked per day dropped by 29%, which is a significant improvement. As a result, the commissioning performance p rose by 11%. The marginal drop in the average number of picks per tour may have resulted from the limited amount of transport boxes available at the station—in analogy with the current situation in the warehouse.

TABLE I
Simulation Results: Introducing Optimal Routes

	Regular Day	Optimised Routes	Δ
Performance p	57,18	63,03	+ 10,83%
Total distance (km)	480,451	368,353	- 29,34%
Scattering of picks	12,21	12,47	+ 2,13%
Avg. #pickstour	4,08	4,02	- 1,47%
Total work time (h)	248:14	221:07	- 10,92%

C. Learning the Routing Selection

Following a simple heuristic for choosing the next pick location provides the quickest start of the tour, whereas any thoughts into planning the tour systematically requires additional computational burden and, thereby, introduces an additional delay. While optimal, computing the shortest paths for each pair of picks is costly, and so is the systematic search for their concatenation into a minimal route. Therefore, in this second optimisation step, we train our agent-based model to select one of several routing options: (1) Follow the established heuristic, (2) greedily perform a nearest neighbour search, (3) rely on A* and find a good route using branch-and-bound, (4) rely on A* and perform an exhaustive search to find the optimal route. If a route needs to consider only very few stops, an exhaustive search might still be more efficient overall, whereas finding the optimal route for many stops might take longer than the gained advantage from the reduction in the walking distance.

The behaviour that was learned is summarised in Table II: Depending on the number of picks, a certain routing was selected. The established, naive model is only good for routes with one pick - as it provides direct access and does not require any calculations, not even a nearest neighbour comparison. This comparision yields an optimal result when only two picks need to be ordered, which renders it the preference in this case. Between three to ten picks, branch-and-bound on pairs of picks and their shortest paths is the general method of choice. But in one case, when there are exactly four picks, an exhaustive search turns out to be more efficient overall than the branch-and-bound approach. Yet, we want to point out that these preferences are the result of a probabilistic learning algorithm and should not be considered in absolute terms. Any number of picks greater than ten renders nearest neighbour search the preferred method.

TABLE II
LEARNED ROUTING SELECTION

# Picks	Route Planning Approach
1	Established Model
2	Nearest Neighbour
3	Branch-and-Bound
4	Exhaustive Search
5 to 10	Branch-and-Bound
> 11	Nearest Neighbour

Learning the above route planning preferences yielded mixed results, as can be seen in Table III: Most importantly, the commissioning performance p is further increased, if only by 1.03%. The total work time is reduced by a negligibly small amount, and the picks are slightly less scattered across the section/floor. However, at the same time, the total distance increased as much as the performance did and the average number of picks of a tour slightly decreased as well.

D. Introducing an ABC Inventory Control System

The next optimisation step foresees the allotment of storage locations with respect to the accessing frequency of different goods. To recall, the established approach was to evenly distribute the goods. While the even distribution is maintained, we introduce the ABC inventory control system. This means

TABLE III
SIMULATION RESULTS: LEARNING ROUTING SELECTION

	Optimised Routes	Routing Selection	Δ
Performance p	63,06	63,73	+ 1,03%
Total distance (km)	339,493	342,996	+ 1,03%
Scattering of picks	12,47	12,09	- 3,05%
Avg. #pickstour	4,02	4,00	- 3,05%
Total work time (h)	221:07	220:38	- 0,22%

that at each station, goods and storage modules are assigned to either of three classes based on their access frequency, whereas A-class items are closest to the station and C-class items furthest. When comparing the results to a regular work day (Table IV), we do not see any improvements. On the contrary, the total distance increases by 6% and the commissioning performance drops by 2%. In addition, scattering of picks as well as the work time are increased. Clearly, one reason for the degradation in performance lies in the fact that the established, naively guided routes are closely intertwined with the bias-free, even distribution of goods.

TABLE IV
SIMULATION RESULTS: INTRODUCTION OF AN ABC SYSTEM

	Regular Day	ABC System	Δ
Performance p	57,45	56,09	- 2,37%
Total distance (km)	480,451	509,354	+ 6,02%
Scattering of picks	12,21	13,11	+ 7,37%
Avg. #pickstour	4,13	4,09	- 0,97%
Total work time (h)	248:14	253:50	+ 0,07%

Analysing the distribution of picks across the station, one recognises that modules close to the lower aisle are accessed much more frequently than those further away. This is exactly the ABC system's expected effect. Figure 5 illustrates this effect considering the same period of time and the same station as the one analysed in Section IV-B. Considering that the expected qualitative result has been achieved, we decided to also test this approach from another angle.

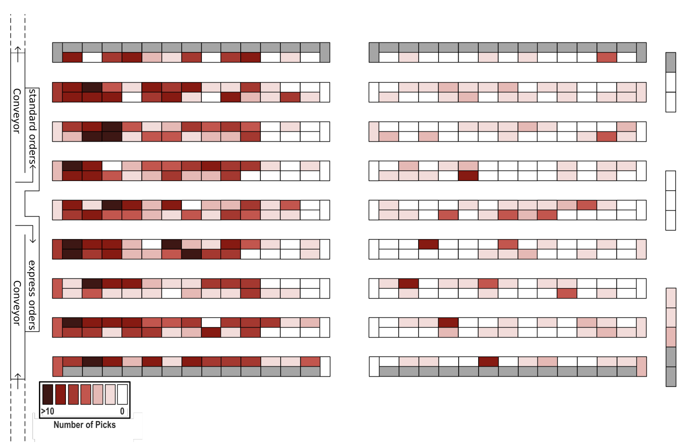

Fig. 5. Picks with the introduction of the ABC system.

E. ABC Control System and Optimal Routes

As the introduction of ABC classes in combination with the established, naive routing mechanism did not yield desirable results, we investigated next, whether an ABC system can be advantageous, if optimal routing is adopted as well. We compare the results with the optimal routing approach, as we have previously witnessed that the deployment of optimal routing yields a general improvement, even without adaptation of article and module locations. As one can see in Table V, introducing the ABC system does indeed yield another improvement, if the routing scheme is adapted. The total distance is reduced by 3%, and the performance p is increased by 0.4%. But it also results in a slightly more scattered distribution of goods.

TABLE V
SIMULATION RESULTS: ABC AND OPTIMISED ROUTES

	Optimised Routes	ABC & Optimised Routes	Δ
Performance p	63,67	63,89	+ 0,35%
Total distance (km)	339,493	330,782	- 2,57%
Scattering of picks	12,47	13,05	+ 0,35%
Avg. #pickstour	4,02	4,02	+ 0,00%
Total work time (h)	221:07	220:14	- 0,40%

F. Learning to Leave Transport Boxes

As non of the reported approaches could reduce the scattering of picks, we decided on testing another distinct approach. Instead of trying to achieve improvements while touring, we now consider, whether to accept a commission with only one or with several picks, or whether individual picking tasks should be left behind at the station. The concrete approach foresees that a worker picks up the first transport box and decides which of its successors he wants to skip or add to his tour. We first simulated this new approach extending the regular work day model. As one can see in Table VI, this new approach achieves the desired results. Scattering of picks drops by 55%. At the same time, the total distance is reduced and the commissioning performance p is increased by 2%.

TABLE VI
SIMULATION RESULTS: LEARNING TO LEAVE TRANSPORT BOXES

	Regular Day	Leaving Boxes	Δ
Performance p	57,45	58,63	+ 2,05%
Total distance (km)	480,451	466,601	- 2,88%
Scattering of picks	12,21	5,55	- 54,55%
Avg. #pickstour	4,13	3,92	- 5,08%
Total work time (h)	248:14	246:19	+ 0,77%

In order to visualise the newly achieved distribution of picks, we take up the illustration scheme shown in Section IV-B. In Figure 6, the module with the highest access frequencies is depicted in red, whereas all the other modules that have been visited are shown in blue. One can see that the picks are not scattered across the whole area any longer. Rather, mostly the modules around the red one in the back are visited, whereas the front area is only visited as an exception.

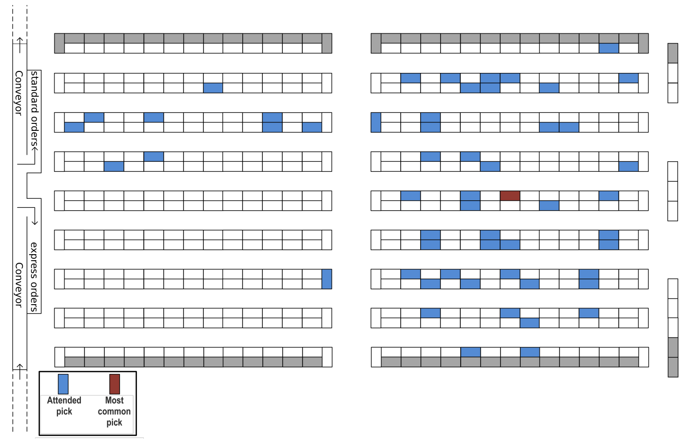

Fig. 6. Scattering of picks of a tour after learning to leave transport boxes.

G. Learning to Leave Transport Boxes and Optimal Routes

As we have shown that leaving transport boxes at the station can increase the efficiency and decrease the distribution of picks in the previous section, we now extend it by and compare it to the dynamic optimisation of routes. Table VII reveals that learning which commissioning tasks should be processed in combination further improves on the optimal routing selection scheme. The total distance could be reduced by another 8%. The commissioning performance p increased by 3%, or two picks per hour. At the same time, the scattering of picks was reduced by 45%, which led to an overall decrease of work time by 5 hours. The average number of picks per tour also decreased, which might be a result from the increased efficiency.

TABLE VII
SIMULATION RESULTS: LEARNING TO LEAVE TRANSPORT BOXES AND OPTIMAL ROUTES

	Optimised Routes	Leaving Boxes	Δ
Performance p	63,89	65,72	+ 2,86%
Total distance (km)	330,782	305,027	- 7,79%
Scattering of picks	13,05	7,16	- 45,13%
Avg. #pickstour	4,02	3,89	- 3,23%
Total work time (h)	220:14	215:33	- 2,13%

H. Combined Benefit of all Advantageous Approaches

For the conclusion of our empirical studies, we combined all the aforementioned, advantageous approaches and analysed how well they played together. In particular, this last model includes the application of the ABC inventory system, learning the routing scheme, and learning which transport boxes should be combined. Again, the results of this combined approach to self-adaptive warehouse optimisation are compared to the basic regular work day model. Table VIII summarises the results. In combination, our optimisation steps yield a significant increase in efficiency. The commissioning performance p is increased by 14%, or eight picks per hour. The total distance is reduced by 37%, or 177km per day and the scattering of picks is diminished by 57%. These improvements lead to a reduction of the total work time by 13%, or 32 man-hours per day. The average number of picks per tour are reduced by 21% as well, which may, again, be owed to lack of storage boxes at the station. These numbers emphasise that especially the interplay of the proposed optimisation approaches leads to significant efficiency improvements.

TABLE VIII
SIMULATION RESULTS: BENEFIT OF COMBINED APPROACH

	Regular Day	**Combined Approach**	Δ
Performance p	57,45	65,54	+ 14,08%
Total distance (km)	480,541	303,043	- 36,94%
Scattering of picks	12,21	5,29	- 56,67%
Avg. #pickstour	4,13	3,28	- 20,58%
Total work time (h)	248:14	216:11	- 12,91%

VI. SUMMARY & FUTURE WORK

For traders in the mail order business the commissioning process is of seminal importance. Often, it provides them the greatest lever for processing an order with maximal costumer satisfaction in mind. Therefore, it is decisive to analyse, fully understand and continuously evolve this process. In this work, we have conducted research to deploy a combination of optimisation and learning techniques to result in dynamic, self-adaptive warehouse commissioning processes. We based our representation, simulation, and optimisation on empirically captured model data. Despite its very concrete application, our approach can be adapted to suit other warehouse situations. Most importantly, the dynamics that are inherent in the trading business, e.g. strong variations in customer behaviour, seasonal changes, or changes in product availability, subcontractor selection, etc. render it necessary that any optimisation approaches can entertain self-adaptation to changing environments. Aiming at high commissioning performance, our approach can self-adapt accordingly by means of an agent-based, reinforcement learning scheme in place.

In particular, based on the empirical model we increased the degree of freedom of worker agents step by step. We subjected their behavioural freedom to reinforcement learning and compared the outcome to the model results of lesser complexity, starting with the un-optimised status quo. Each extension brought about some improvements, with the final combination of all aspects resulting in performance gains of 14%. The individual measures towards this success included (1) the introduction of shortest path calculations for pairs of picks, (2) planning picking routes based on heuristics, heuristic search or exhaustive search, (3) learning to choose the right planning approach automatically, (4) introduction of ABC classes for goods and modules, and (5) strategic skipping of individual picking jobs to arrive at lean picking routes.

The presented model has the ability to self-adapt. Yet, it would be interesting to investigate the extent and the speed of environmental dynamics it can cope with. A rough estimate for according boundaries could potentially be inferred just by theoretical analysis of the model. Next, comprehensive experiments for a sensitivity analysis could be conducted. In addition to this theoretical work, conceptual work would have to follow suit that sheds light on the medium and long-term integration of such a self-adaptive system in the warehouse ecosystem. We have already introduced some of the outlined aspects into the investigated warehouse. Yet, the tandem of self-adaptation and feedback, and respectively, goal specification, would have to be concretely designed. The interface of the commissioning system generally requires some additional efforts. Currently, the workers are supported by a hand-held device. Possibly, this could be further supported by smart, wireless in-ear headphones, which could help navigation without loosing visual focus. In this context, precise indoor location could become a major asset, also, for instance, for detecting jammed lanes.

REFERENCES

[1] M. Wooldridge, *An introduction to multiagent systems*. John Wiley & Sons, 2009.

[2] ——, “Agent-based software engineering,” *IEE Proceedings-software*, vol. 144, no. 1, pp. 26–37, 1997.

[3] H. Qi, Y. Xu, and X. Wang, “Mobile-agent-based collaborative signal and information processing in sensor networks,” *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1172–1183, 2003.

[4] C. Müller-Schloer, H. Schmeck, and T. Ungerer, Eds., *Organic Computing - A Paradigm Shift for Complex Systems*, ser. Autonomic Systems. Birkhäuser Verlag, 2011.

[5] S. Fischer, “Naturinspirierte verfahren in der informatik—anregungen für die logistik,” in *Das Beste der Logistik*. Springer, 2008, pp. 137–144.

[6] S. Tomforde, H. Prothmann, J. Branke, J. Hähner, M. Mnif, C. Müller-Schloer, U. Richter, and H. Schmeck, “Observation and control of organic systems,” in *Organic Computing—A Paradigm Shift for Complex Systems*. Springer, 2011, pp. 325–338.

[7] H. Prothmann, S. Tomforde, J. Branke, J. Hähner, C. Müller-Schloer, and H. Schmeck, “Organic traffic control,” in *Organic Computing—A Paradigm Shift for Complex Systems*. Springer, 2011, pp. 431–446.

[8] T. Becker, C. Illigen, B. McKelvey, M. Hülsmann, and K. Windt, “Using an agent-based neural-network computational model to improve product routing in a logistics facility,” *International Journal of Production Economics*, vol. 174, pp. 156–167, 2016.

[9] H. Beckmann, *Supply Chain Management: Strategien und Entwicklungstendenzen in Spitzenunternehmen*. Springer-Verlag, 2012.

[10] B. Scholz-Reiter and H. Höhns, “Selbststeuerung logistischer prozesse mit agentensystemen,” in *Produktionsplanung und-steuerung*. Springer, 2006, pp. 745–780.

[11] D. Frey, L. Mönch, T. Stockheim, P.-O. Woelk, and R. Zimmermann, “Agent. enterprise-integriertes supply chain management mit hierarchisch vernetzten multiagenten-systemen.” in *GI Jahrestagung (1)*, 2003, pp. 47–63.

[12] K. Neumann, *Produktions-und Operations-Management*. Springer-Verlag, 2013.

[13] T. Gudehus, *Logistik: Grundlagen-Strategien-Anwendungen*. Springer-Verlag, 2013.

[14] N. J. Nilsson, *Principles of artificial intelligence*. Morgan Kaufmann, 2014.

[15] F. Jaehn and E. Pesch, *Ablaufplanung: Einführung in Scheduling*. Springer-Verlag, 2014.

[16] G. Reinelt, *The traveling salesman: computational solutions for TSP applications*. Springer-Verlag, 1994.

[17] E. Horowitz and S. Sahni, *Fundamentals of computer algorithms*. Computer Science Press, 1978.

[18] P. Dayan and T. J. Sejnowski, “Td (λ) converges with probability 1,” *Machine Learning*, vol. 14, no. 3, pp. 295–301, 1994.

[19] T. M. Mitchell *et al.*, *Machine learning. WCB*. Boston, MA: McGraw-Hill, 1997.