The final chapter is in essence a review of the past, an assessment of the present, and a forecast of possible futures of BOR. Covering a wide range of topics *"ranging from the modeling of thought and behavior and the incorporation of behavioral factors in models to how people understand models and how thought and behavior is influenced by modeling work"*, the book makes a clear case for incorporating behavior into OR projects. The objective of illustrating "behavior with models", "behavior in models" and "behavior beyond models" has clearly been achieved in this publication. This is a great reference for those who want to get familiarized with BOR.

As early as September 1964 during the first international conference held by the Operational Research Society with the theme "Operational Research and the Social Science", Sir Charles Goodeve, stated: "*operational research people are very much concerned with change and can deal with the logic, including the economics, of it. But attitudes of people - managers, technicians, workpeople, salesmen, customers etc. - can throw the best of predictions into confusion*". These are words that ring true, even today. The book Behavioral Operational Research is therefore an important and essential addition to further enhance and broaden the impact of OR as a discipline.

TUTORIALS

Combinatorial Auctions In Collaborative Vehicle Routing

Margaretha Gansterer
Production and Operations Management, University of Vienna
margaretha.gansterer@univie.ac.at

Richard F. Hartl
Production and Operations Management, University of Vienna
richard.hartl@univie.ac.at

The highly competitive transportation industry has put pressure on carriers to increase efficiency. Collaboration has been encouraged by public authorities since it serves such ecological goals as reducing road congestion, noise pollution, and harmful emissions.

In horizontal collaboration, carriers form coalitions in order to perform parts of their logistics operations jointly. By exchanging transportation requests, they can operate more efficiently. In the full truckload market, different types of auctions have been successfully applied to meet these goals (e.g. Sheffi 2014; Kuyzu et al. 2015). However, in less than truckload (LTL) settings, the application of these methods becomes more challenging. The exchange of requests can be organized through combinatorial auctions, where collaborators submit requests for exchange to a common pool. In combinatorial auctions, requests are not traded individually but are combined in bundles. This is of particular importance in vehicle routing, where a request might be worthless unless combined with others. Figure 1 shows such an example for 3 carriers with their pickup and delivery requests.

If a bidding carrier's price is accepted, the carrier receives the full bundle. The carrier with a rejected bid does not get any item in the bundle. This eliminates the risk of obtaining only a subset of requests that does not fit into the current request portfolio.

Alternatively, generation of bundles can be moved to the carriers themselves. Thus, the auctioneer could offer the set of requests without grouping them into bundles. The carriers then give their bids on self-created packages of requests. Since the auctioneer cannot guarantee that all requests will be assigned to carriers, the outsourcing option is included. In this case, a carrier can get a set of requests that exceeds its capacity.

Without outsourcing, an LTL combinatorial transportation auction typically follows a 5-phase procedure (Berger and Bierwirth, 2010):

1. Carriers decide which requests to put into the auction pool.
2. Auctioneer generates bundles of requests and offers them to the carriers.
3. Carriers place their bids for the offered bundles.
4. Auctioneer allocates bundles to carriers based on their bids (winner determination problem).
5. Collected profits are distributed among the carriers.

Combinatorial auctions can be effective mechanisms to allocate transportation requests (e.g., Ledyard et al., 2002; Song and Regan, 2003; Krajewska and Kopfer 2006; Berger and Bierwirth, 2010; Ackermann et al., 2011; Dai et al., 2014). Nevertheless, each of the 5 auction phases bears a complex and at least a partly unsolved decision problem. In the first phase, participating collaborators can decide either on self-fulfilment, i.e. they plan and execute their transportation requests with their own capacities, or to offer some of them to other carriers. Aiming at network profit maximization, carriers should try to offer requests that are valuable for other network participants. Otherwise, the auction mechanism will not yield improved solutions. However, the identification of requests that are valuable for collaborators is not trivial since the actors do not want to reveal sensitive information. This auction phase is illustrated in Figure 2.

An intuitive solution would be to let the carriers solve a team orienteering problem and put requests that do not appear in the optimal tour into the pool (Archetti et al., 2014). However, Gansterer and Hartl (2016a) show that the best request evaluation criteria take geographical aspects into account. They clearly dominate pure profit-based strategies. Schopka and Kopfer (2016) analyse several other heuristic pre-selection strategies.

In the second phase, the requests in the pool are grouped into bundles. These are then offered to participating carriers. From a practical point of view, offering all possible bundles is not manageable, since the number of bundles grows exponentially with the number of requests that are in the pool. Gansterer

Combinatorial Auctions In Collaborative Vehicle Routing

and Hartl (2016b) show that the complete set of bundles can efficiently be reduced to a subset of attractive ones. They develop a proxy function for assessing the attractiveness of bundles under incomplete information. This proxy is then used in a Genetic Algorithms-based framework that aims to produce attractive and feasible bundles. With only a little loss in solution quality, instances can be solved in a fraction of the computational time compared to the situation where all possible bundles are evaluated.

Bundles are then offered to the carriers in Phase 3. By giving a bid on a bundle, a carrier reflects its willingness to acquire the bundle. Typically, a bid is based on the carrier's marginal profit, which is the difference of the profits including and excluding the bundle in the tour. Hence, for each bid, a routing problem has to be solved. Buer (2014) proposes heuristics for the identification of promising bundles in order to decrease the bidding effort.

Phase 4 addresses the winner determination problem. Bundles are assigned to carriers in such a way that the total coalition gain is maximized. Models for the winner determination problem need to ensure that requests (that can be part of more than one bundle) are only assigned once. Carriers should not receive more requests than they can handle, unless the outsourcing option is included. The framework proposed by Gansterer and Hartl (2016b) guarantees that a feasible assignment of bundles to carriers is found without making it necessary for carriers to reveal sensitive information and without the need of outsourcing. The winner determination problem is an NP-hard optimization problem (Rothkopf et al., 1998). Figure 3 shows a possible re-assignment of requests after the winner determination problem.

The main advantage of horizontal collaboration in logistics is that the total profit of the coalition will be higher than the sum of the carriers' stand-alone profits. However, the collaboration gain should be distributed in such a way that each partner benefits from participating in the coalition (Phase 5). If this is not the case, carriers may leave the coalition. The profit sharing method should additionally force companies to avoid strategic behavior that negatively influences the coalition gain, and reward behavior that benefits the coalition (Vanovermeire et al., 2014). Strategic behavior can occur in each of the phases. Behavior that negatively impacts the coalition might involve carriers not using their real marginal costs for generating bids or not following the rules for selecting requests. For example if in Figure 2, Carrier A selects A1 and A3 which for it are not attractive, the opportunity for greater profits for the coalition is affected. Nonetheless, the impact of strategic behavior on the outcome of these collaborations is still unknown.

Admittedly, many challenging questions still have to be answered to make combinatorial auctions efficiently applicable to real-world LTL settings. For instance, the strong relationship between the phases has not been investigated. In addition, the realistic assumption that carriers might behave strategically opens many interesting research questions. Therefore, collaborative vehicle routing is an active research field with a high practical importance in the LTL area. Combinatorial auctions have the huge potential of becoming powerful mechanisms for increasing collaboration profits.

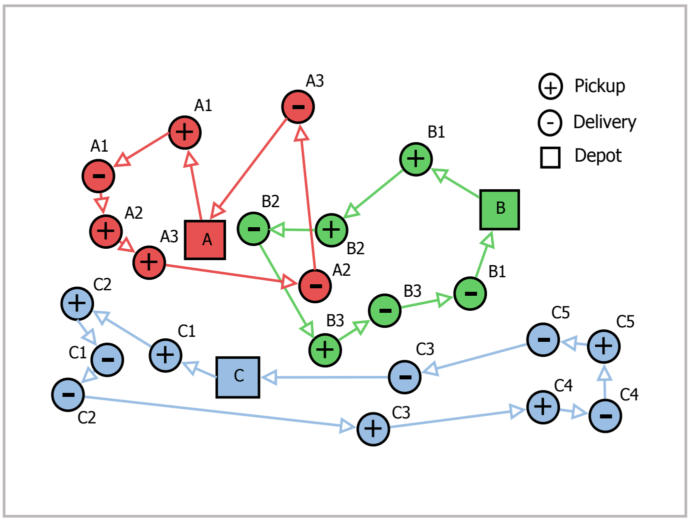

Figure 1: Illustration of the non-collaborative tours of 3 carriers. For carrier B, buying request C4 alone will probably not be profitable. However, the extra travel cost will probably be compensated if C5 is acquired additionally.

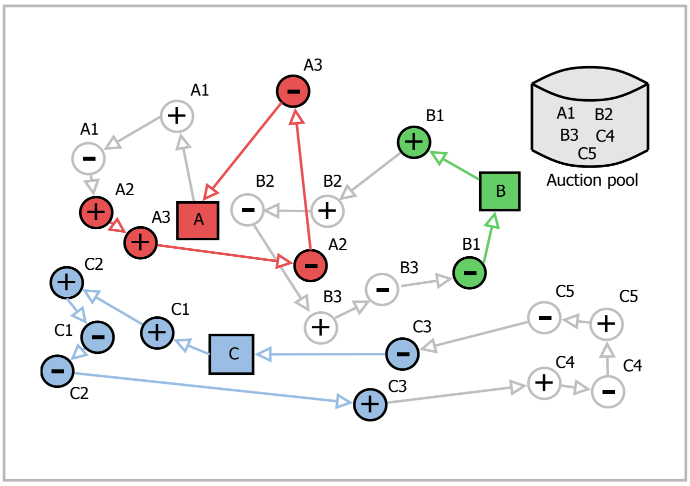

Figure 2: Carriers submit request to the pool. Carrier A selects requests based on marginal profits, while carriers B and C take geographical information into account.

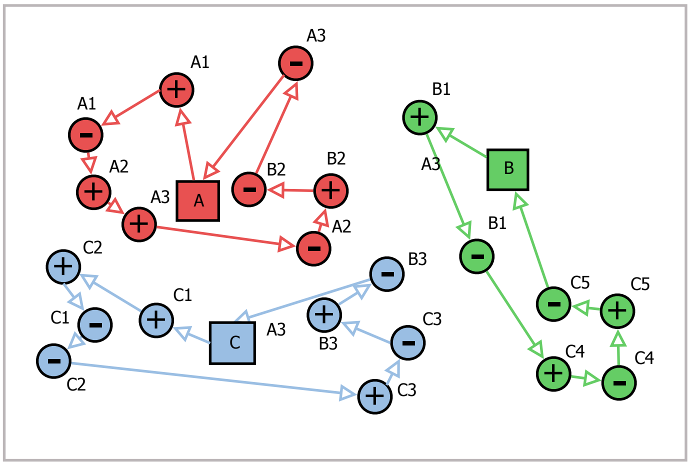

Figure 3: Possible re-assignment of requests to carriers.

References
Ackermann, H., Ewe, H., Kopfer, H., Küfer, K., 2011. Combinatorial auctions in freight logistics. In: Böse, J., Hu, H., Carlos, C., Shi, X., Stahlbock, R., Voss, S. (Eds.), Computational Logistics. Vol. 6971 of Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 1–17.
Archetti C, Speranza MG, Vigo D (2014) Vehicle routing problems with profits. In: Toth P, Vigo D (eds) Vehicle routing: problems, methods, and applications. MOS-SIAM series on optimization, Philadelphia, pp 273–297
Berger, S., Bierwirth, C., 2010. Solutions to the request reassignment problem in collaborative carrier networks. Transportation Research Part E: Logistics and Transportation Review 46, 627–638.
Buer T (2014). An exact and two heuristic strategies for truthful bidding in combinatorial transport auctions. CoRR. arxiv.org/abs/1406.1928.
Dai, B., Chen, H., Yang, G., 2014. Price-setting based combinatorial auction approach for carrier collaboration with pickup and delivery requests. Operational Research 14 (3), 361–386.
Gansterer, M., Hartl, R. F., 2016a. Request evaluation strategies for carriers in auction-based collaborations. OR Spectrum 38(1), 3–23.
Gansterer, M., Hartl, R. F., 2016b. Bundle generation in combinatorial transportation auctions, working paper.
Krajewska, M., Kopfer, H., 2006. Collaborating freight forwarding enterprises. OR Spectrum 28(3), 301–317.
Kuyzu, G., Akyol, C.G., Ergun, Ö., Savelsbergh, M., 2015. Bid price optimization for truckload carriers in simultaneous transportation procurement auctions, Transportation Research Part B: Methodological 73: 34-58.
Ledyard, J., Olson, M., Porter, D., Swanson, J., Torma, D., 2002. The first use of a combined-value auction for transportation services. Interfaces 32 (5), 4–12.
Rothkopf, M.H., Pekec, A., Harstad, R.M., 1998. Computationally manageable combinatorial auctions. Management Science, 44:1137–1147.
Schopka, K.; Kopfer, H.: Pre-selection Strategies for the Dynamic Collaborative Transportation Planning Problems. In: Lübbecke, M., Koster, A., Letmathe, P., Madlener, R., Peis, B., Walther, G. (Eds.) Operations Research Proceedings 2014: Selected Papers of the Annual International Conference of the German Operations Research Society (GOR), RWTH Aachen University, Germany, September 2-5, 2014, Springer, 523-529.
Sheffi, Y., 2004. Combinatorial auctions in the procurement of transportation services. Interfaces 34(4), 245–252.
Song, J., Regan, A.C., 2003. Combinatorial auctions for transportation service procurement: the carrier perspective. Transportation Research Record 1883: 40–46.
Vanovermeire, C., Sörensen, K., Van Breedam, A., Vannieuwenhuyse, B., Verstrepen, S., 2014. Horizontal logistics collaboration: decreasing costs through flexibility and an adequate cost allocation strategy. International Journal of Logistics Research and Applications, 17(4):339–355.

Lifelong Learning in Optimisation

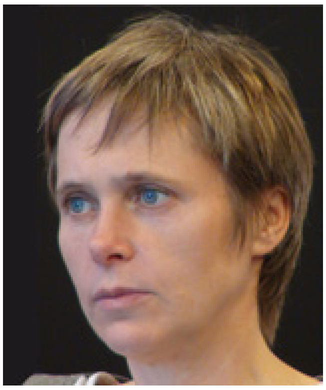

Emma Hart

School of Computing
Edinburgh Napier University
E.Hart@napier.ac.uk

Kevin Sim

School of Computing
Edinburgh Napier University
K.Sim@napier.ac.uk

A food-processing factory manufactures different products based on customer orders from a stock of raw ingredients. They want to create weekly schedules that ensure that all orders are manufactured on time. They commission an optimisation algorithm from the OR department at the local University, supplying historical data to allow the algorithm to be tuned, and at first are very pleased with the results. Over time, however, the company begins to notice that the quality of the schedules produced seems to be deteriorating, with several orders delivered late, causing them to lose confidence in the algorithm.

Unfortunately, the experience of the factory is typical – the market in which the factory operates is dynamic: their customer base slowly changes over time, the type of products demanded by consumers changes, and the availability of the raw ingredients fluctuates due to changes in suppliers. The optimisation algorithm that was tuned to work well on the original problem instances no longer works on the new problem instances that now have very different characteristics. Of course, the manager can return to the OR department to ask for the algorithm to be re-tuned – or perhaps a completely new algorithm is required – but either way it's a time consuming process that requires considerable expert input.

As human problem-solvers, we cope much better than machines with adapting our problem-solving processes to change. We modify existing processes, drawing on prior experience if relevant, and generate completely new strategies if the magnitude of the change demands it. In contrast, optimisation algorithms rarely work in this way - an algorithm is usually designed and tuned to work well across a class of instance using a range of examples drawn from the class and then 'fixed'. At best, some adaptive methods can alter parameter values online as a single instance is solved, based on feedback from the solver, but any knowledge gained during this process is lost as soon as the algorithm terminates.

In response to this, we propose a new model of optimisation system – in which systems not only learn how to solve a problem but learn continuously over a lifetime. Such systems improve their problem solving abilities over time: retaining knowledge, using it to improve future learning, and generating new knowledge when required.