[Schö08] Schönberger, J.; Kopfer, H.: On the Value of Objective Function Adaptation in Online Optimization. In: Operations Research 2007. Springer, 2008

On the Value of Objective Function Adaptation in Online Optimisation

Jörn Schönberger and Herbert Kopfer

University of Bremen, Chair of Logistics, Wilhelm-Herbst-Straße 5, 28359 Bremen, Germany {jsb,kopfer}@uni-bremen.de

Summary. We analyse a dynamic variant of the vehicle routing problem with soft time windows in which an average punctuality must be guaranteed (e.g. lateness is allowed at some customer sites). The existing objective function does not support both the aspiration for punctuality and least cost so that additional efforts are necessary to achieve an acceptable punctuality level at least possible costs. Within numerical experiments it is shown that static penalties are not adequate in such a situation but that an adaptation of the objective function before its application to the next problem instance supports the search for high quality solutions of the problem.

1 Introduction

We consider a vehicle routing problem in which the adequate fulfilment mode of consecutively arriving customer requests is to be selected. Arriving requests are fulfilled using the cheap but not necessarily reliable self-fulfilment mode (SF) or the more expensive but reliable subcontracting mode (SC). We propose to adaptively adjust the weights of the costs of the two fulfilment modes in a mono-criterion objective function. Doing so, we refine the idea of Gutenschwager et al. [1] who initially propose to adjust higher-ranked objective functions to the recent decision situations while solving the next instance in a sequence of optimisation models in online-fashion. This adaptive approach is compared in numerical experiments with the typically used penalty approach in which a static unchangeable value is used to depreciate late visits at customer sites.

Section 2 introduces the investigated decision problem. Section 3 outlines the decision algorithms. The computational experiments are reported in Section 4.

2 Dynamic Decision Problem

We investigate the following generalisation of the vehicle routing problem with time windows.

Soft time windows. Lateness at a customer site is possible but causes penalty costs. The portion p_t of the requests completed or scheduled for completion in the interval $[t-t^-;t+t^+]$ is observed. At least p_t of all requests of the interval around the current time t must be started within the agreed time windows.

Subcontraction. Each request can be served by a vehicle from the own fleet in the SF-mode or it can be subcontracted (SC-mode) to a logistic service provider (LSP). In the former case, late arrivals at customer sites cannot be prevented but in the latter case, an in-time service is assured. A once subcontracted request cannot be re-integrated into the routes of the own vehicles.

Uncertain demand. Only a subset of all requests is known to the planning authority at the time when the decision concerning subcontracting is made and the routes for the own vehicles are generated. The planning authority decides about the fulfilment mode of a request as soon as it becomes known.

Additional requests arriving at time t_i trigger the update of the so far followed transportation plan TP_{i-1} which contains the decision how the waiting requests will be served. For the update of TP_{i-1} to TP_i we have introduced an optimisation model [2] whose solving identifies a least costs refresh of the transportation plan. Each possible update is evaluated by the resulting costs using the objective function (1).

$$\underbrace{C_1(\mathcal{RP}(\mathcal{R}_i^{int})) + C_2(\mathcal{RP}(\mathcal{R}_i^{int}))}_{\text{self-fulfilment costs}} + \underbrace{C_3(\mathcal{SC}(\mathcal{R}_i^{ext}))}_{\text{SC usage costs}} \rightarrow \min. \qquad (1)$$

The set $\mathcal{R}_i^{int}$ contains all requests for which the self-fulfilment mode has been selected and the set $\mathcal{R}_i^{ext}$ comprises all subcontracted requests at time t_i. With $\mathcal{RP}$, we denote the least cost collection of paths for the own vehicles and $\mathcal{SC}$ refers to the minimal-charge bundling of subcontracted requests. Then, $C_1(\mathcal{RP}(\mathcal{R}_i^{int}))$ denotes the travel costs of the own vehicles, $C_2(\mathcal{RP}(\mathcal{R}_i^{int}))$ gives the penalty costs to be paid for late customer site visits of own vehicles. Finally, $C_3(\mathcal{SC}(\mathcal{R}_i^{ext}))$ gives the costs of the subcontracted requests.

If both fulfilment modes would lead to the same costs for a given request r, e.g. if

$$\alpha := \frac{C_3(TP3(r))}{C_1(TP1(\mathcal{R}_i^{int})) + C_2(TP1(r))} \approx 1, \tag{2}$$

then both fulfilment modes SF and SC will be used to the same extend as long as the limited capacity of the own fleet is exhausted. As soon as the capacity of the own fleet is exhausted then some requests are shifted into the SC-mode. However, if $\alpha \gg 1$ then the aspiration for cost minimal modes prevent the usage of the SC mode. If C_2 is not stringent and severe enough, then the number of late severed requests increases, so that p_t falls down.

In the remainder of this article we investigate the dependencies between the severeness of the penalisation of late requests and the selection of the fulfilment mode. Thereby, we assume that $\alpha \gg 1$, so that the aspiration for least cost transportation plans does not support the selection of the mode that leads to the highest percentage of punctually served requests.

We use artificial test cases [2] constructed from the 100-customer Solomon [3] instances $\{R103, R104, R107, R108\}$ for an experimental analysis of the aforementioned situation. In these scenarios, a demand peak leads to a temporal exhaustion of the cheaper SF mode. We propose and test ideas to overrule the cost-based mode decision in order to consider punctuality issues to a larger extend.

3 Algorithm Details

We use the Memetic Algorithm described e.g. in [2] to derive a new transportation plan after additional requests have arrived. In such a case, the execution of TP_{i-1} is interrupted and TP_{i-1} is replaced by TP_i.

A piece-wise linear penalty function h is deployed, which is 0 for delays up to T_{max} time units and which increases proportionally up to a maximal penalty value P_{max} (money units) for delays longer than the threshold delay of 100 time units. Using this penalty calculation the sum of penalty payments is $C_2(\mathcal{RP}(\mathcal{R}_i^{int})) := \sum_{r \in \mathcal{R}^{int}} h(delay(r))$, where $delay(r)$ gives the distance to the latest allowed visiting time at the customer site corresponding to request r.

The previously introduced penalty function h is deployed with different parameter settings. We perform simulations with the maximal penalty values $P_{max} \in \{50, 75, 100, 125\}$ and the tolerance ranges $T_{max} \in \{0, 25, 50, 75\}$. A parameterisation of the penalty function is denoted by $P(P_{max}, T_{max})$.

Alternatively, we deploy an adaptation mechanism that re-weights the costs of the two fulfilment modes in the objective function in dependence from the currently observed punctuality p_t. The idea of this approach is to artificially lower the costs of the SC mode (compared to the SF mode) if p_t is low in order to make the usage of the SC mode more attractive.

$$f(t_i) \cdot \Big[C_1(\mathcal{RP}(\mathcal{R}_i^{int})) + C_2(\mathcal{RP}(\mathcal{R}_i^{int}))\Big] + C_3(\mathcal{SC}(\mathcal{R}_i^{ext})) \rightarrow \min \quad (3)$$

The coefficient $f(t_i)$ is adjusted before the update of TP_{i-1} to TP_i starts. It is $f(t_0) = 1$ and $f(t_i) = 1 + \alpha \cdot \Omega(t_i, p_{t_i})$ for $i \geq 1$. We use the piece-wise linear function Ω which is 0 if the current punctuality p_{t_i} is larger than $p^{target} + 0.05$ and which equals 1 if $p_{t_i} \leq p^{target} - 0.05$. In the latter case, it is $f(t_i) = 1 + \alpha$ and subcontracting a request is identified by the solver via the objective function to be cheaper than the self-fulfilment with respect to the currently used objective function (3). For p_t-values between $p^{target} - 0.05$ and $p^{target} + 0.05$ the function Ω decreases proportionally from 1 down to 0. Since the re-definition of the coefficient affects the search trajectory heading of the solving algorithm, we call this approach Search Direction Adaptation (SDAD).

4 Numerical Experiments

Experimental Setup. We analyse two scenarios. In scenario I, SF and SC have the same prices ($\alpha = 1$) but in scenario II SC is quite more expensive than SF ($\alpha = 3$).

A single simulation run $(P, \omega, \epsilon, \alpha)$ is determined by the request set $P \in \{R103, R104, R107, R108\}$, the algorithm seeding $\omega \in \{1, 2, 3\}$, the applied strategy $\epsilon \in \{SDAD\} \cup \{PEN(a, b) \mid a \in \{50, 75, 100, 125\}, b \in \{0, 25, 50, 75\}\}$ and $\alpha \in \{1, 3\}$. Thus, $4 \cdot 3 \cdot 17 \cdot 2 = 406$ simulation runs have been executed.

Throughout the simulations we observed the maximal punctuality decrease (in percent) $\delta(\epsilon, \alpha)$ after the demand peak and the cumulated overall costs $C(\epsilon, \alpha)$.

The results observed for the PEN$(\cdot, \cdot)$-experiments in scenario I are presented in Fig. 1. The left isoline-plot shows the observed maximal punctuality decreases $\delta(\alpha, \epsilon)$. In $A_1^{\delta}(-0.6)$ maximal punctuality variations between -0.6% and 0 (light grey shaded area) are observed. Punctuality variations between -1% and -0.6% appear in $A_1^{\delta}(-1)$. Decreases of p_t between 1% and 1.4% take place in $A_1^{\delta}(-1.4)$

The right isoline plot compiles the average of the cumulated costs $C(\cdot,\cdot)$ occurred during the simulation runs within the PEN$(\cdot,\cdot)$-experiments. Additional costs of less than 5% ($A_1^C(0)$) are observed for small penalties and high tolerance values (light grey shaded) areas. A cost increase of more than 15% is realized if $T_{max} \leq 25$ and $P_{max} \geq 75$.

The application of SDAD leads to a maximal punctuality decrease of 1% at nearly the same costs (dark grey shaded areas in the two plots). We conclude, that if $\alpha = 1$ (same costs for SF and SC) then the static parameter setting $(50, 75)$ of h performs sufficient with respect to a sufficiently high service quality as well as service cost minimisation.

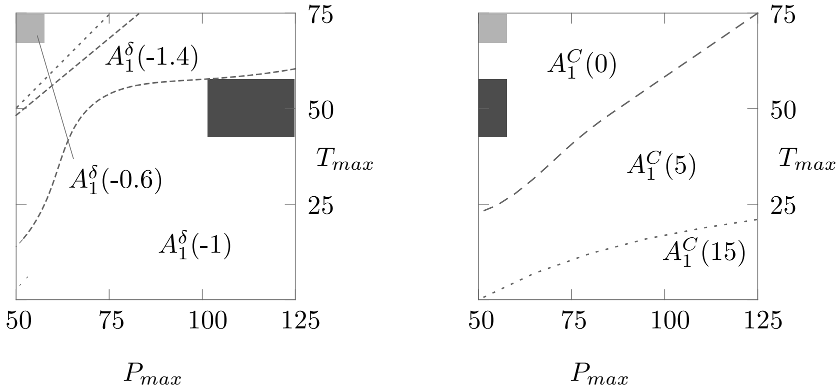

Fig. 1. Scenario I ($\alpha = 1$): punctuality decrease (left) and costs (right)

Quite different results are observed in scenario II (Fig. 2). The service quality optimal parameter setting is $P_{max} = 125$ and $T_{max} = 25$ with a maximal punctuality reduction of 2.9% at costs of 72284,10 (light grey shaded area in the left plot in Fig. 2). This setting causes additional costs of more than 15% (cf. right plot in Fig. 2). On the other hand, the cost optimal parameter setting ($P_{max} = 50, T_{max} = 75$, (light grey shaded area in the right plot in Fig. 2) results in a punctuality collapse of around 20%. It is therefore not possible to find a parameter setting for h that satisfies both goals costs minimisation and punctuality preservation to the maximal extend at the same time.

For both reasonable tradeoff parameter settings (100,50) and (75,50) we observe a significantly higher punctuality decrease (compared to the punctuality preserving setting) or quite enlarged costs (compared to the cost optimal setting).

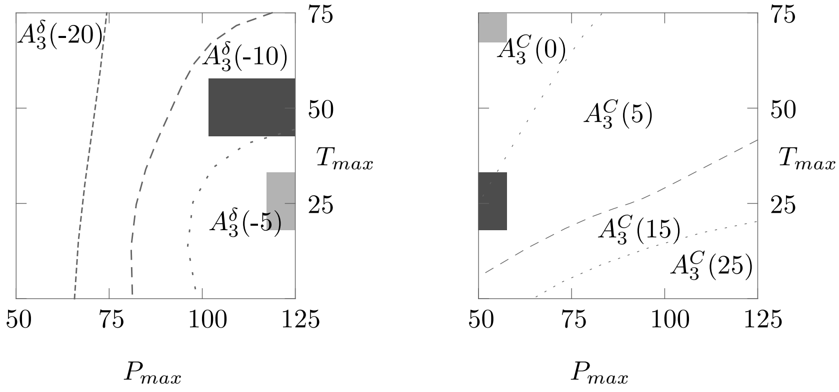

Fig. 2. Scenario II ($\alpha = 3$): Punctuality decrease (left) and costs (right)

In contrast, SDAD performs very well. It comes along with an acceptable punctuality decrease of only 5.7% which is better than the performance of the two trade-off proposals (dark gray shaded areas in Fig. 2). The costs resulting from the application of SDAD are only 63888,7 which is a significant reduction of the costs compared to the two proposed trade-off parameter settings.

5 Conclusions and Outlook

We have shown that the adaptive definition of an objective function supports the achievement of a good trade-off between service quality and service costs in a volatile environment. Further research activities are dedicated the identification of the right key indicators to be used to derive the right adaptation decisions.

References

1. K. Gutenschwager, F. Böse, and S. Voß. Effiziente Prozesse im kombinierten Verkehr – ein neuer Lösungsansatz zur Disposition von Portalkränen. *Logistik Management*, 5(1):62–73, 2003.
2. J. Schönberger and H. Kopfer. Flexible transport process planning in volatile environments and the adaptation of a cost-based objective function. In *Proceedings of Logistik Management 2007*, pages 263–283. Gabler-Verlag, 2007.
3. M. Solomon. Algorithms for the vehicle routing and scheduling problem with time window constraints. *Operations Research*, 35(2):254–265, 1987.