[Schö07g] Schönberger, J.; Kopfer, H.: Autonomously Controlled Adaptation of Formal Decision Models - Comparison of Generic Approaches.
In: Proceedings of the International Conference on Dynamics in Logistics. 2007

Autonomously Controlled Adaptation of Formal Decision Models – Comparison of Generic Approaches

Jörn Schönberger and Herbert Kopfer

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

1 Introduction

This article is about different strategies to adapt automatically a decision model to a changed problem situation. By means of an example from transportation logistics, we assess two strategies for implementing recent knowledge about a volatile decision problem into an optimisation model representing the problem at hand: Adaptation of the constraint set and adjustment of the objective function.

Section 2 introduces the considered decision problem. Section 3 describes the proposed model adaptation strategies. Section 4 describes the experimental setup and reports the achieved numerical results.

2 Vehicle Scheduling Problem

The problem we are investigating in this contribution generalises the common vehicle routing problem with time windows (Solomon, 1987) in three aspects.

a) Soft Time Windows. Lateness at a customer site is possible but causes penalty costs. Although a particular request is allowed to be late, the portion p^{target} of the f_t requests completed in $[t^-;t]$ and expected to be completed in $[t;t^+]$ must be in time. Let f_t^{comp} be the number of the requests completed timely within the last t^- time units and let f_t^{expec} be the number of punctually scheduled requests within the next t^+ time units, then $p_t := (f_t^{comp} + f_t^{expec})/(f_t) \geq p^{target}$ has to be achieved.

b) Subcontracting. Logistics service providers (LSP) are paid for the reliable fulfilment of selected requests. An LSP receives a certain amount of payment for this service and ensures that the request is fulfilled within the specified time window. A subcontracted request remain unconsidered while constructing the routes for the own vehicles. If a request has been subcontracted then this decision cannot be revised later on.

c) 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 subcontracting or self-fulfilment of a request as soon as it becomes known.

Decision model. A transportation plan describes how the known requests are fulfilled. Subsequently arriving requests are accepted and handled by updating the existing transportation plan. A sequence of transportation plans TP_0, TP_1, TP_2, … is generated reactively at the ex ante unknown update times t_0, t_1, t_2, … and each single transportation plan is executed until it is updated. It is aimed at keeping the costs for the execution of the additional requests as low as possible but, on the other hand, to provide a sufficiently high reliability within the request fulfilment.

A single request r attains consecutively different states. Initially, r is known but not yet scheduled (F). Then, r is assigned to an own vehicle (I) or subcontracted (E). If the operation at the corresponding customer site has already been started but not yet been finished the state (S) is assigned to r. The set $R^+(t_i)$ is composed of additional requests released at time t_i. Requests completed after the last transportation plan update at time t_{i-1} are stored in the set $R^C(t_{i-1},t_i)$. The new request stock $R(t_i)$ is determined by $R(t_i):=R(t_{i-1}) \cup R^+(t_i) \setminus R^C(t_{i-1},t_i)$ and the set $R^X(t_i)$ contains all currently available requests belonging to the state $X\in\{F,E,I,S\}$.

The transportation plan update problem at time t_i is as follows. Let V denote set of all own vehicles, $P_v(t_i)$ the set of all routes executable for vehicle v and let $P(t_i)$ denote the union of the sets $P_v(t_i)$ $(v\in V)$. If the request r is served in path p then the binary parameter a_{rp} is set to 1, otherwise it is set to 0. A request r, already known at time t_{i-1} that is not subcontracted in TP_{i-1} is served by vehicle v(r) according to TP_{i-1}. The travel costs associated with route p are denoted as $C^1(p)$ and $C^2(p)$ refers to the penalties associated with p. Finally, $C^3(r)$ gives the subcontracting costs of request r.

We deploy two families of binary decision variables. Let $x_{pv}=1$ if and only if path $p\in P(t_i)$ is selected for vehicle $v\in V$ and let $y_r=1$ if and only if request r is subcontracted.

$$\sum_{p\in P(t_i)}\sum_{v\in V}(C^1(p)+C^2(p))x_{pv}+\sum_{r\in R(t_i)}C^3(r)y_r \to \min \tag{1}$$

$$\sum_{p\in P_v(t_i)} x_{pv}=1 \quad \forall v\in V \tag{2}$$

$$x_{pv}=0 \quad \forall p\notin P_v(t_i), v\in V \tag{3}$$

$$y_r + \sum_{p \in P(t_i)} \sum_{v \in V} a_{rp} x_{pv} = 1 \quad \forall r \in R(t_i) \tag{4}$$

$$y_r = 1 \quad r \in R^E(t_i) \tag{5}$$

$$\sum_{p \in P_{v(r)}} a_{rp} x_{pv(r)} = 1 \quad \forall r \in R^S(t_i) \tag{6}$$

The overall costs for TP_i are minimised (1). One route is selected for each vehicle (2) and the selected path p is realisable by vehicle v (3). Each single request known at time t_i is served by a selected vehicle or given away to an LSP (4) but a once subcontracted request cannot be re-inserted into the paths of the own vehicles (5). An (S)-labelled request cannot be re-assigned to another vehicle or LSP (6).

Test cases. The construction of artificial test cases from the Solomon instances (Solomon, 1987) is described by Schönberger and Kopfer (2007). In these scenarios, demand peaks that represent significant changes in the decision situations interrupt balanced streams of incoming requests. Two fulfilment modes are available for each request: self-fulfilment and subcontracting. The costs for the first mode are normalized to 1 monetary unit for each travelled distance unit. For each subcontracted distance unit α monetary units have to be paid to the LSP. Each single request r causes overall costs of F_r:=C^3(r) monetary units calculated by multiplying the distance between the LSP depot and the customer site location of r with α.

3 Online Decision Strategies

We use the Memetic Algorithm (MA) scheduling framework introduced in Schönberger (2005) for the repeated update of the transportation plans. Whenever new requests are released then the MA is re-called in order to produce an updated transportation plan. This approach is a realisation of the generic online optimisation concept (Krumke, 2001). Is referred to as NONE because neither the search space nor the objective function are adapted to current situations.

Search Space Adaptation. The severeness of the consecutively solved optimization models changes due to the variation of the system load. In order to consider load modifications, we enforce the subcontraction if the currently observed punctuality p_t falls below the desired least punctuality p^{target}.

A continuous piece-wise linear control function h that is equal to 1 if $p_t \leq p^{target}$-0.05, 0 if $p_t \geq p^{target}$+0.05 and that falls proportionally from 1 down to 0 if p_t increases from p^{target}-0.05 up to p^{target}+0.05 is used in order to determine the intensity of the enforcement of the fulfillment mode "subcontraction". Therefore, at first we select randomly the portion h(p_t) of R^+(t) and the fulfillment mode of the selected requests (collected in R^{pre}(t)) is pre-select as "subcontraction". It is not allowed to change the selected mode anymore so that we have to replace the constraint (5) by the new constraint (7).

$$y_r = 1 \quad r \in R^E(t_i) \cup R^{pre}(t_i) \tag{7}$$

Since the constraint (7) is adapted with respect to the current punctuality, this online decision strategy is referred to as Constraint Set Adaptation (CSAD).

Search Direction Adaptation. An adaptation of the cost coefficient in the objective function (1) leads to a re-valuation of the costs. This adjustment corresponds to the adaptation of the search direction of the used solver. In order to enable such a search direction adaptation (SDAD), the objective function (1) is replaced by the function (8) and the coefficient f_i is set to $f_0=1$ and $f_i=1+\alpha h(p_{t_i})$.

$$\sum_{p \in P(t_i)} \sum_{v \in V} f_i (C^1(p) + C^2(p)) x_{pv} + \sum_{r \in R(t_i)} C^3(r) y_r \rightarrow \min \tag{8}$$

As long as the punctuality is sufficiently large ($p_{t_i} \geq p^{target} + 0.05$) both fulfilment modes are weighted only by the original costs. As soon as the punctuality decreases f_i increases and if p_{t_i} has reached an unsatisfactory level ($p_{t_i} \geq p^{target} - 0.05$), then $f_i = 1+\alpha$, which means that in (8) the costs for subcontraction are weighted less than the self-entry costs. After the coefficient has been adjusted, the updated model is solved and is used for the generation of a transportation plan update.

4 Numerical Experiments

Experimental setup. We deploy a piece-wise linear penalty function, which is 0 for delays shorter than 10 time units and which increases proportionally up to a maximal value of 25 money units for delays longer than 100 time units. The target punctuality is set to $p^{target}=0.8$.

We perform experiments for different tariff levels $\alpha \in \{1, 1.25, 1.5, 1.75, 2, \text{and } 3\}$ ($\alpha=1$ represents a fair and comparable LSP tariff which leads to the same costs as in the case of self-fulfillment). Requests are taken from the Solomon instances R103, R104, R107, R108. Each of the |{R103, R104, R107, R108}| · |{1, 1.25, 1.5, 1.75, 2,3 }| · |{NONE, CSAD, SDAD}|=72 scenarios is simulated three times leading to overall 216 performed simulation experiments. Here, we report the average results observed for each scenario. For analysing the impacts of the objective function adaptation and the constraint set adaptation with respect to the tariff level α and the decision policy $\varepsilon \in$ {NONE,CSAD,SDAD}, we first calculate p_t's maximal decrease $\delta(\varepsilon,\alpha) := \min_{t \geq 1500} (p_t(\varepsilon,\alpha)/p_{1000}(\varepsilon,\alpha))$ after the demand peak's start.

Let $T_{\varepsilon,\alpha}^{below}$ denote the first time in which ptarget is not achieved and $T_{\varepsilon,\alpha}^{heal}$ refers to the time in which p^{target} is finally re-achieved. We define $\pi(\varepsilon,\alpha) := (T_{\varepsilon,\alpha}^{heal} - T_{\varepsilon,\alpha}^{below})/(4000)$ as the percentage of low quality situations ($p_t<p^{target}$) within the observation interval [1000,5000].

Beside the effects on the process reliability, we have recorded the resulting process costs. Let $C_{\varepsilon,\alpha}(t)$ denote the cumulated overall costs realized up to time t. In order to quantify the impacts of tariff level rising, we calculate the relative growth $c(\varepsilon,\alpha):=(C_{\varepsilon,\alpha}(5000))/(C_{\varepsilon,1}(5000))-1$ caused by increasing freight tariffs. To compare the impacts of the cost criteria in the mode decision, we calculate the cost increase $r(\varepsilon,\alpha):=c(\varepsilon,\alpha)/c(\text{NONE}, \alpha)-1$ caused by switching from NONE to CSAD and SDAD. Similarly, we calculate the relative growth of the travel costs $c^{int}(\varepsilon,\alpha)$, of the subcontracting costs $c^{ext}(\varepsilon,\alpha)$ as well as of the penalty costs $c^{pen}(\varepsilon,\alpha)$. The contribution of the travel costs to the overall costs is defined as $m^{int}(\varepsilon,\alpha):=C_{\varepsilon,\alpha}^{int}(5000)/C_{\varepsilon,\alpha}(5000)$ (the portion $m^{ext}(\varepsilon,\alpha)$ of the subcontracting costs as well as the portion $m^{pen}(\varepsilon,\alpha)$ of the penalties are determined in the same way).

Simulation Results. Independently from the used strategy ε, the reliability decreases if the freight level α is increased. A severe decrease of δ(ε,α) is observed for increasing α. The application of CSAD as well as SDAD leads to values of δ(ε,α) above 91%. However, SDAD seems to be able to keep δ(ε,α) on a slightly higher level than CSAD (Table 1).

Table 1. Minimal punctuality δ(ε,α)

ε	α					
	1	1,25	1,5	1,75	2	3
NONE	98,0%	97,0%	93,6%	80,7%	81,5%	72,0%
CSAD	92,9%	96,4%	93,6%	93,5%	93,3%	91,1%
SDAD	99,0%	97,0%	95,6%	96,1%	93,1%	92,2%

The percentage of situations with a punctuality below p^{target} increases if the tariff level is lifted by the tariff level increase but the application of a model adaptation defers the occurrence to higher tariff levels (Table 2).

Table 2. Percentage π(ε,α) of replanning situations with a punctuality below p^{target}

ε	α					
	1	1,25	1,5	1,75	2	3
NONE	--	5,0%	15,0%	42,5%	42,5%	47,5%
CSAD	--	--	5,6%	8,3%	13,9%	19,4%
SDAD	--	--	--	--	5,0%	7,5%

The increase in the service level is mainly based by an extension of the LSP incorporation. The results in Table 3 show that the application of CSAD as well as SDAD leads to an increase of σ(ε,α).

Table 3. Maximal quote of subcontracted requests σ(ε,α)

ε	α					
	1	1,25	1,5	1,75	2	3
NONE	18,8%	13,5%	9,6%	10,9%	5,7%	9,3%
CSAD	21,2%	17,1%	20,0%	20,2%	20,3%	20,9%
SDAD	21,9%	17,1%	16,0%	15,2%	15,5%	16,5%

CSAD and SDAD overrule the cost criteria in the mode selection for $\alpha>1$ so that additional costs occur. Table 4 shows that SDAD leads to significantly less additional costs (10,9%) than the application of CSAD (49,8%). However, the severeness of the cost increase after an increase of the freight tariff level is quite different in the three strategies (Table 5).

Table 4. Cost increase r(ε,α) after switching from NONE to CSAD or SDAD

ε	α					
	1	1,25	1,5	1,75	2	3
CSAD	2,8%	11,4%	17,3%	23,1%	25,5%	49,8%
SDAD	4,6%	8,2%	7,5%	6,8%	5,9%	10,9%

Table 5. Increase c(ε,α) of the cumulated overall costs

ε	α					
	1	1,25	1,5	1,75	2	3
NONE	0,0%	4,9%	13,6%	20,4%	27,6%	43,0%
CSAD	0,0%	13,8%	29,6%	44,3%	55,7%	108,5%
SDAD	0,0%	8,6%	16,7%	23,0%	29,2%	51,7%

Table 6. Cost increase

ε		α					
		1	1,25	1,5	1,75	2	3
NONE	$c^{int}(\varepsilon,\alpha)$	0,0%	70,0%	119,5%	148,7%	167,1%	206,4%
	$c^{ext}(\varepsilon,\alpha)$	0,0%	-45,1%	-67,0%	-77,6%	-80,8%	-91,4%
	$c^{pen}(\varepsilon,\alpha)$	0,0%	93,3%	141,4%	182,6%	237,3%	430,0%
CSAD	$c^{int}(\varepsilon,\alpha)$	0,0%	50,3%	87,8%	115,9%	132,6%	154,1%
	$c^{ext}(\varepsilon,\alpha)$	0,0%	-10,7%	-8,5%	-1,7%	6,7%	78,3%
	$c^{pen}(\varepsilon,\alpha)$	0,0%	78,1%	113,7%	129,2%	140,4%	182,9%
SDAD	$c^{int}(\varepsilon,\alpha)$	0,0%	39,0%	66,1%	90,5%	106,8%	137,1%
	$c^{ext}(\varepsilon,\alpha)$	0,0%	-7,8%	-9,9%	-13,1%	-12,3%	5,8%
	$c^{pen}(\varepsilon,\alpha)$	0,0%	56,7%	92,4%	119,2%	137,9%	182,1%

The overall costs are split into three cost drivers: travel costs (int), subcontracting costs (ext) and penalties (pen) as shown in Table 6. Independently from the applied adaptation strategy, the highest cost increase is caused by additional penalty costs. The second highest cost driver is the travel costs. However, the highest discrepancies are observed for the subcontracting costs. In the NONE experiment, the pure cost-based decision strategy disqualifies the subcontracting mode as more as the tariff level is increased. At the end (α=3), the external costs have been reduced by 91.4%. If the search direction is adapted in the SDAD experiment then the subcontracting cost increase is rather small but if the constraint set is adapted then the promotion of this fulfilment mode leads to additional subcontracting costs of 78.3%. As long as the freight tariffs are comparable (α=1, 1.25), the LSP costs contribute mostly to the overall costs independently from the applied strategy ε. For higher freight tariffs, travel costs of the own vehicles become the most important cost driver in the NONE-experiment but if the decision model is adapted then LSP charges as well as travel costs contribute similarly to the overall costs although CSAD leads to a higher contribution of LSP charges (Table 7).

Table 7. Split of costs

ε		α					
		1	1,25	1,5	1,75	2	3
NONE	$m^{int}(\varepsilon,\alpha)$	39,8%	64,5%	76,9%	82,2%	83,3%	85,2%
	$m^{ext}(\varepsilon,\alpha)$	57,2%	29,9%	16,6%	10,6%	8,6%	3,4%
	$m^{pen}(\varepsilon,\alpha)$	3,1%	5,6%	6,5%	7,2%	8,1%	11,4%
CSAD	$m^{int}(\varepsilon,\alpha)$	36,0%	47,6%	52,1%	53,9%	53,8%	43,9%
	$m^{ext}(\varepsilon,\alpha)$	61,2%	48,1%	43,2%	41,7%	41,9%	52,3%
	$m^{pen}(\varepsilon,\alpha)$	2,8%	4,4%	4,6%	4,4%	4,3%	3,8%
SDAD	$m^{int}(\varepsilon,\alpha)$	32,3%	41,3%	45,9%	50,0%	51,7%	50,5%
	$m^{ext}(\varepsilon,\alpha)$	65,7%	55,8%	50,7%	46,4%	44,6%	45,8%
	$m^{pen}(\varepsilon,\alpha)$	2,0%	2,9%	3,3%	3,6%	3,7%	3,7%

In order to understand the different costs, we have analysed the recorded control function values h(t) in the CSAD as well as in the SDAD experiment. Fig. 1 shows the intervention intensities for CSAD (dashed bold line) and SDAD (continuous bold line). Although the CSAD control function produces a lower average value than the SDAD control function, the amplitude of the CSAD control function oscillation is quite larger than the amplitude of the SDAD control function oscillation. This might be a reason, why CSAD leads to more intensive subcontracting usages with a lower number of requests served by own vehicles.

With respect to the costs, SDAD dominates CSAD (cf. Table 4 and Table 5), however, CSAD is able to enforce the subcontracting of requests to a larger extend (cf. Table 3). This might be an idea to combine SDAD and CSAD into a common strategy that uses SDAD mechanisms to heal small punctuality deficiencies and that employs CSAD capabilities to correct severe punctuality decreases.

5 Conclusions

We have analysed an evolving decision problem from transportation logistics. In order to implement ex ante unknown problem knowledge into the formal model decision, we have proposed extensions to the general online optimization framework that allow the contextual formulation of the instances in an online optimization problem. Within comprehensive numerical experiments, we have proven the general applicability of the adaptation strategies and a comparison has been carried out. Future research will address the combination of the so far separately used adaptation strategies.

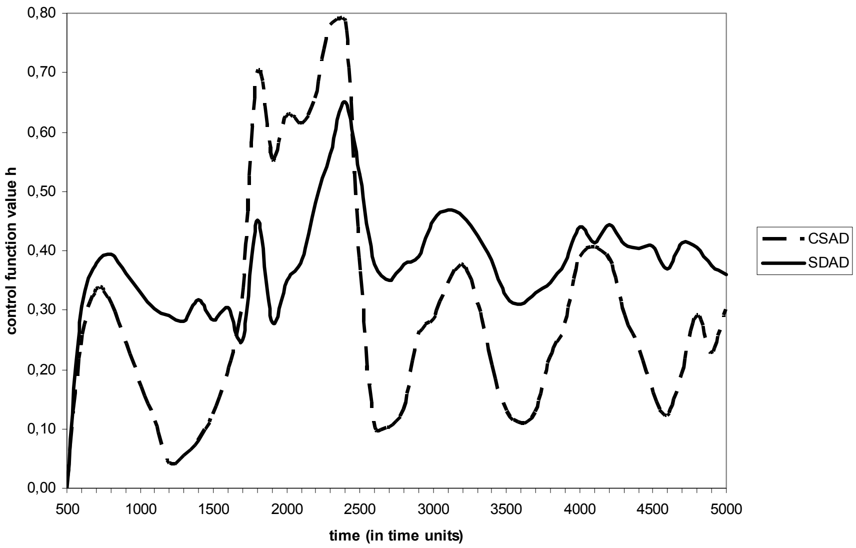

Fig. 1: Control function values and self entry quotes observed in the α=3 experiment

Acknowledgement. This research was supported by the German Research Foundation (DFG) as part of the Collaborative Research Center 637 "Autonomous Cooperating Logistic Processes" (Subproject B7)

References

Krumke S (2001) Online-Optimization – Competitive Analysis and Beyond. Habilitation Thesis, Technical University Berlin, Germany.

Schönberger J; Kopfer H (2007) On Decision Model Adaptation in Online Optimization of a Transport System, in: Günther H-O; Mattfeld DC; Suhl L: Management logistischer Netzwerke, Springer 2007, Berlin, Heidelberg, S. 361-381

Schönberger, J (2005) Operational Freight Carrier Planning. Springer.

Solomon, M (1987) Algorithms for the vehicle routing and scheduling problem with time window constraints. Operations Research 35 (2), pp. 254-265.