



GRASP with Hybrid Path Relinking for Bi-Objective Winner Determination in Combinatorial Transportation Auctions

Tobias Buer, Department of Information Systems, University of Hagen, Germany, E-Mail: tobias.buer@fernuni-hagen.de
Giselher Pankratz, Department of Information Systems, University of Hagen, Germany, E-Mail: giselher.pankratz@fernuni-hagen.de

Abstract

The procurement of transportation services via large-scale combinatorial auctions involves a couple of complex decisions whose outcome highly influences the performance of the tender process. This paper examines the shipper's task of selecting a subset of the submitted bids which efficiently trades off total procurement cost against expected carrier performance. To solve this bi-objective winner determination problem, we propose a Pareto-based greedy randomized adaptive search procedure (GRASP). As a post-optimizer we use a path relinking procedure which is hybridized with branch-and-bound. Several variants of this algorithm are evaluated by means of artificial test instances which comply with important real-world characteristics. The two best variants prove superior to a previously published Pareto-based evolutionary algorithm.

Keywords: transportation procurement, combinatorial auction, set covering problem, carrier performance, multi-criteria decision making, Pareto optimization, GRASP

Manuscript received August 28, 2009, accepted by Karl Inderfurth (Operations and Information Systems) July 8, 2010.

1 Introduction

Shippers, such as industrial and trading companies, regularly use framework agreements in order to contract out their transportation tasks to motor carriers. In a framework agreement the shipper arranges with the freight carrier which transportation services he or she is to take over on what level of service and at what cost they are to be carried out. In this case, a framework agreement (denoted as *contract* in the remainder) comprises the transportation (repeated, if necessary) of a volume of goods from a pickup point to a delivery point; for example, the weekly delivery to a chain store from a central depot for a period of 12 months.

Many shippers tender their transportation contracts via transportation procurement (reverse) auctions. Transportation procurement auctions are of high economic relevance. Caplice and Sheffi (2006) reported on the size of real-world transportation auctions in which they were involved over a period of five years. According to

their report, in a single transportation auction up to 470 (median 100) carriers participated, up to 5,000 (median 800) lanes were tendered, and the annual cost of transportation amounted to US-\$ 700 million (median US-\$ 75 million). Elmaghraby and Keskinocak (2004) presented a case study of a procurement auction event in which a do-it-yourself chain operating mainly in North America procured transportation services for about a quarter of the in-bound moves to their chain stores, which corresponds to over 600 lanes (in the study at hand, the terms *lane* and transportation *contract* are used interchangeably). Similarly, shippers in Europe strive to consolidate their transportation procurement activities by Europe-wide tendering. In recent years, specialized Internet portals have emerged which offer European contractors a neutral environment for issuing their logistics contracts. Sizes of tenderings processed via such platforms reportedly scale up to several hundreds of contracts (Cargoclix Dr. Meier & Schmidt

GmbH 2010).

The process of transportation contract tendering is carried out, as a rule, in three steps (Caplice and Sheffi 2006): In the *qualification stage* (pre-auction stage) the shipper chooses the carriers that can provide transportation services at a minimum specified level. Regularly applied quality criteria are, for example, solid financial ratios, suitable IT systems for a smooth exchange of data, a vehicle fleet that is suitable for transporting the goods, or the reliable adherence to delivery deadlines (Caplice and Sheffi 2006, 2003). Freight carriers that were able to pass the qualification stage are allowed to take part in the *bidding stage* (auction stage). In this stage the qualified carriers submit concrete bids for the tendered transportation contracts. After that, the *allocation stage* (post-auction stage) follows. In this stage, the shipper assigns the contracts to carriers based on the bids received in accordance with previously established criteria.

In addition to cost aspects, the shipper has to consider several other criteria in the allocation stage. In particular, qualitative aspects – despite preselection of carriers in the qualification phase – can be of importance also in the allocation phase. The remaining carriers may fulfill the minimum level of service required, but beyond that, they vary with regard to their suitability for performing particular services. Varying service-related demands can arise, for example, because of the different nature of the goods (fragile, high-value or dangerous goods versus goods that are none of these things) or the varying strategic importance of the respective customer. For these reasons it seems wise to explicitly match carriers to contracts, in order to strive for the highest possible service quality from the carriers that are performing the tasks.

Electronic support for the tendering process, e.g., via Internet-based information and communications systems, usually extends to the publication of the auction event and the bidding requirements, the structured compilation of the interested carriers and their characteristics relevant to the tender in the qualification stage, as well as the gathering of the bids in the bidding phase. For supporting the selection decisions that must be made in both the qualification and the allocation stages, the respective systems offer a (graphic or tabular) format for the carrier-related bidding information that has been compiled. Further support for these

decisions, for example in the form of automatically generated decision-suggestions, is not yet part of the standard function-scope of the respective systems. This, for the most part, can be attributed to the complexity of the shipper's decision situation in both phases.

This article addresses this gap. In doing so it concentrates on how to support the shipper's contract assignment decision in the allocation stage. In contrast to the bulk of the literature (cf. Section 3), we use an enhanced model of the contract assignment problem. This model which was first introduced by Buer and Pankratz (2010) allows better consideration of the important features of the shipper's decision situation as stated above.

For solving this problem we present a novel approach based on the GRASP metaheuristic (greedy randomized adaptive search procedure). As an extension we apply a path relinking technique which is hybridized by an exact branch-and-bound procedure. Our results show that the GRASP approach clearly outperforms an evolutionary algorithm previously published in Buer and Pankratz (2010).

The article is structured as follows: In Section 2 important features of the shipper's decision problem during the allocation stage are explained in detail. Section 3 then discusses combinatorial auctions as an appropriate auction format for transportation procurement. Against the background of the previously described requirements, one-sided focusing on cost minimization is identified as a weak point of conventional approaches. In Section 4 the investigated winner determination problem is defined as a bi-objective set covering problem. Section 5 reviews related work and motivates the choice of a Pareto-based GRASP approach to solve the problem described. Details of the chosen approach are given in Section 6. The results of a test with artificial problem instances are presented and discussed in Section 7. Finally, Section 8 summarizes the main findings and outlines further research issues.

2 Features of decision making during the allocation stage of transportation procurement auctions

A crucial basis for the further considerations are the features of the decision situation in the allocation phase, namely

- the existence of valuation interdependencies between transport contracts and
- the multi-criteria character of carrier assignment.

Both aspects are explained in more detail in the following.

From the bidder's point of view, transport contracts feature *valuation interdependencies*. Thus, the costs of a contract for a bidder can crucially depend on which other orders he or she is awarded in the allocation stage. Complementary and substitutional interdependencies can be distinguished: If two contracts are complementary, e.g., because they can be combined to make a busy route, then the costs of the combined execution of both orders are lower than the sum of the costs that result for each of the orders when executed separately (cost subadditivity). On the other hand, if the costs of a combined execution exceed the sum of the costs of isolated execution (cost superadditivity), e.g., if the bidder can take on each individual contract, but has to invest in additional transport capacity for executing both of them, then a substitutional relationship between the contracts is existent (Kopfer and Pankratz 1999; Pankratz 2000).

It has already been stated that during the allocation stage in practice it is not only the costs but also a series of other criteria that influence the decision of the shipper; thus it is a question of *multi-criteria decision making*. Some of these criteria can be formulated in a tolerable approximation of practical custom as side constraints, for example, a limitation of the number of assigned carriers in order to avoid overdependence on individual suppliers or rather to avoid coordination overhead. Other criteria, however, are to be considered explicitly as minimization or maximization objectives, in order to adequately reflect practical requirements. Along with the common objective of minimizing total costs, this contains the goal of achieving the highest possible overall performance quality in carrying out the contracts. Since in practice there can regularly be observed a positive correlation between the service level of a freight carrier and the level of the cargo rate, there is obviously a conflict of objectives between cost minimization and overall performance quality maximization. Therefore, it is sensible for a shipper to have important orders conveyed by particularly reliable freight carriers even on acceptance of higher cargo rates, while,

on the other hand, it appears appropriate to award orders with lower requirements to such companies that offer a reduced but still reasonable level of performance.

3 Combinatorial auctions for transportation procurement

In conventional auction formats which allow the carriers only to bid separately on each of the contracts, the bidders are unable to adequately express valuation interdependencies between the contracts in their bids. To address the former issue the use of so-called combinatorial tendering (also known as combinatorial reverse auctions) is therefore recommended in the literature (Sheffi 2004), which allows the submission of bids on any desired combination of the tendered contracts (bundle-bids). Because of the complex time/space interdependencies between transport contracts, the transport domain is considered to be a classic application area of combinatorial auctions (Cantillon and Pesendorfer 2006; Fischer, Müller, Pischel, and Schier 1995; Gomber, Schmidt, and Weinhardt 1997; Kopfer and Pankratz 1999; Pankratz 2000; Sandholm 1993). In comparison to simple auctions, the assignment of contracts to carriers in combinatorial auctions is, however, clearly more difficult. The task of the auctioneer is to determine from the bundle-bids of the participants the most favorable allocation of all the contracts. This problem, described in the literature as the Winner Determination Problem (WDP), is an NP-hard, combinatorial optimization problem (Sandholm, Suri, Gilpin, and Levine 2002). In the literature there is as yet only a fairly limited number of papers that address a WDP in the context of the procurement of transport services (Cantillon and Pesendorfer 2006; Caplice and Sheffi 2006, 2003; Elmaghraby and Keskinocak 2004; Ledyard, Olson, Porter, Swanson, and Torma 2002; Song and Regan 2003; Buer and Pankratz 2010).

Nevertheless, in the recent past combinatorial auctions have already been applied in shipping with success. For example, Caplice and Sheffi (2006) reported that between 1997 and 2001 well-known US companies assigned transport services amounting to a total of 175 million US dollars on average using combinatorial auctions.

However, except for the recent paper of Buer and Pankratz (2010) none of the studies published so

far consider the multi-criteria decision situation in the allocation stage. Instead, winning bids are selected taking into account only the objective of minimizing total transportation costs. A consideration of further qualitative criteria is not supported. In this respect all existing models rely on the assumption of a uniformly high level of service of all carriers after getting through the qualification phase, which puts them in the position of being able to carry out each contract at the standard required in each case. This assumption is sensible as long as the requirement level of the orders tendered is homogeneous; however, if the requirements of the transport orders are heterogeneous, then it appears necessary to incorporate the level of performance of the execution of the contract explicitly into the problem as a further objective. Therefore, we have chosen the model introduced by [Buer and Pankratz \(2010\)](#) as the basis for developing a novel GRASP-based solution approach. We describe this model in the next section.

4 The bi-objective winner determination problem for combinatorial transportation procurement auctions (2WDP-SC)

This section formally describes a winner determination problem (WDP) for combinatorial transportation procurement auctions which explicitly takes into account the considerations of the previous section. Since the problem at hand involves two objective functions and is formulated as a generalization of the well-known set covering problem (SC), it is abbreviated as 2WDP-SC. Note that the model presented here is equivalent to the 2WDP-SC model presented in [Buer and Pankratz \(2010\)](#). Given is a set of transport contracts T , a set of carriers C and a set of bundle-bids B . A bundle-bid $b \in B$ is defined as 3-tuple $b := (c, \tau, p)$. This means a carrier $c \in C$ is willing to execute the subset of transport contracts $\tau \subseteq T$ at a price of p . Let a_{tb} be a binary parameter which obtains the value 1, if transport contract $t \in T$ is part of bundle-bid $b \in B$, i.e., $t \in \tau_b$. In that case bundle-bid b covers contract t . Otherwise, if $t \notin \tau_b$ then $a_{tb} = 0$. Given is furthermore a matrix $Q = (q_{tc})_{t \in T, c \in C}$ where q_{tc} indicates the quality level at which carrier c is able to fulfill the transport contract t .

The task is to find a set of winning bids $X \subseteq B$, such that every transport contract t is covered by at least one bid $b \in X$. Furthermore the total procurement costs, expressed in objective function f_1 , are to be minimized and the total service quality, expressed in objective function f_2 , is to be maximized. The 2WDP-SC is modelled as follows:

$$\begin{aligned}
 (1) \quad & \min f_1(x) = \sum_{b \in B} p_b \cdot x_b \\
 (2) \quad & \max f_2(x) = \sum_{t \in T} \max_{b \in B} \{q_{tc} \cdot a_{tb} \cdot x_b\} \\
 (3) \quad & \text{s.t.} \quad \sum_{b \in B} a_{tb} \cdot x_b \geq 1 \quad \forall t \in T \\
 (4) \quad & x_b \in \{0, 1\} \quad \forall b \in B
 \end{aligned}$$

The first objective function (1) minimizes the total cost of the winning bids. The second objective function matches transport contracts to carriers to maximize the service quality. As the $\max\{\cdot\}$ -formulation of (2) is affected by side constraint (3), this side constraint is explained first. Each transport contract t has to be executed exactly once but the \geq operator in (3) allows the procurement of a contract more than once. To avoid this, a strict equal sign in (3) could be used which would result in a set-partitioning-based model.

However, the set partitioning formulation appears inappropriate from the shipper's point of view. Suppose the same data, the costs of an optimal set partitioning solution are at least as high as the costs of an optimal set covering solution. In this sense, the \geq operator in (3) is less restrictive than the $=$ operator. Therefore the set covering model should be preferred from the shipper's point of view.

From the carrier point of view the issue seems indifferent. The carrier suffers no disadvantages under the set covering formulation. This is justified by assuming *free disposal* ([Sandholm, Suri, Gilpin, and Levine 2002](#)). In the transportation procurement context, free disposal means that a carrier does not charge extra costs if he or she is asked by the shipper to execute fewer contracts than agreed upon. That is, the carrier suffers no disadvantages if he or she executes fewer contracts and is paid the previously agreed price. This definitely seems the case in the scenario at hand. Due to the stated reasons, a set covering formulation is preferred to a set partitioning formulation.

This choice affects the second objective function

(2) which maximizes the total service quality level of all transport contracts. Since contracts need to be executed only once, but may be part of more than one winning bid, it is not appropriate to simply add up the respective qualification values of all $b \in X$. Instead, it appears reasonable to assume that the shipper will break ties in favor of the bidder who offers the highest service level for a given contract. Hence, by assumption, for each transport contract t only the maximum qualification values of contracts covered by all winning bids are added up. Note that this rule might introduce an incentive for the carriers towards undesired strategic-bidding behavior (a similar issue would arise with a set partitioning formulation). As this paper does not focus on auction-mechanism design (compare, e.g., [Pekeć and Rothkopf 2003](#) or [Chu 2009](#)), we leave this issue to forthcoming research.

5 Literature review and selection of a solution approach

Up to now, the approaches described in Buer and Pankratz (2010) are the only ones for the 2WDP-SC as stated in Section 4. In their article, the authors proposed a Pareto-based evolutionary algorithm which is based on SPEA2, a problem-independent evolutionary framework for multi-objective optimization introduced by [Zitzler, Laumanns, and Thiele \(2002\)](#). Additionally, they developed an exact branch-and-bound procedure based on the so-called ε -constraint approach which is able to solve small 2WDP-SC instances to optimality.

This contribution aims at setting an alternative approach against the previously described approaches. To lay the foundations for choosing a solution method, this section first briefly reviews the literature on multi-objective optimization approaches. After that, an overview of solution approaches for problems closely related to the 2WDP-SC is given.

5.1 Multi-objective optimization approaches

In the literature, methods for solving combinatorial multi-objective optimization problems have greatly gained in interest in recent years ([Ehrgott and Gandibleux 2000](#); [Gandibleux and Ehrgott 2005](#)).

Depending on whether subjective preference information is exploited during search, the various

methods can be subdivided into preference-based and preference-independent approaches (see, e.g., [Miettinen 2008](#), for a more detailed classification). Preference-based methods take advantage of preference information explicitly provided by the decision-maker. On the one hand, these preferences may be directly integrated into the problem model (a priori approach). On the other hand, preference knowledge can be elicited and considered in an interactive fashion during search (interactive approach). Classical preference-based methods use the preference information to transform the multi-objective problem in a way that it can be treated by single-objective problem-solving methods. This is done, for example, by lexicographic ordering of objectives, by formulating objectives as so-called ε -constraints ([Haimes, Lasdon, and Wismer 1971](#)), or by using the weighted sum method ([Gass and Saaty 1955](#)). Unlike the preference-based methods, preference-independent approaches restrain from making use of any subjective preference information, both in the model and during search. Instead, they generate a number of (objectively) non-dominated solutions, each representing a particular trade-off between the various objectives. Among these solutions, the decision-maker chooses the one which best matches his or her subjective preferences. Methods which support this choice are discussed, e.g., by [Vincke \(1992\)](#). In the remainder, such approaches are referred to as Pareto-based approaches.

For solving the 2WDP-SC, a Pareto-based approach appears appropriate. As reported by [Sheffi \(2004\)](#), identifying the decision-maker's preferences is one of the most time-consuming steps in the course of processing a transportation procurement auction. Furthermore, since transportation contracts are tendered at intervals of one up to three years, subjective preferences may significantly change from auction to auction. Thus, it is even uncertain to which extent preference information from a prior auction can be re-used in a current auction. These circumstances speak in favor of a Pareto-based approach. Moreover, a Pareto-based approach could also support the identification of preferences by providing the decision-maker with a set of feasible, objectively equally-valued, high-quality solutions as an excellent starting point for selecting the (subjectively) best solution.

In the literature on Pareto-based approaches, there are few reports on exact algorithms for

solving multi-objective problems (e.g., [Dhaenens, Lemesre, and Talbi 2010](#), [Bérubé, Gendreau, and Potvin 2009](#), [Mezmaz, Melab, and Talbi 2007](#)). However, the applicability of exact algorithms still is often limited to rather small instances. For example, the only exact approach for solving the 2WDP-SC which has been published so far proved unsuitable when applied to solving problem instances of practical dimensions ([Buer and Pankratz 2010](#)).

Therefore, by far the most Pareto-based approaches in the literature employ metaheuristics. Among these, evolutionary algorithms enjoy the greatest popularity (e.g., [Deb 2001](#); [Obayashi, Deb, Poloni, Hiroyasu, and Murata 2007](#); [Coello, Aguirre, and Zitzler 2005](#)). Due to their distinctive ability to manage a huge number of solutions in a population, evolutionary algorithms are regarded as particularly suitable for treating a set of non-dominated solutions. Meanwhile, evolutionary multi-objective optimization frameworks like NSGA-II ([Deb, Pratap, Agarwal, and Meyarivan 2002](#)) and SPEA2 ([Zitzler, Laumanns, and Thiele 2002](#)) have emerged which facilitate the development of evolutionary Pareto-based methods. Applications based on these frameworks have proven successful in a number of different domains ([Coello, Aguirre, and Zitzler 2005](#); [Obayashi, Deb, Poloni, Hiroyasu, and Murata 2007](#); [Ehrgott, Fonseca, Gandibleux, Hao, and Sevaux 2009](#); [Doerner, Gutjahr, and Nolz 2009](#)).

The already mentioned Pareto-based evolutionary approach for solving the 2WDP-SC ([Buer and Pankratz 2010](#)) also was developed using one of these frameworks (SPEA2).

5.2 Approaches for related problems

In search of a promising alternative approach for the 2WDP-SC, previous work on the general winner determination problem in combinatorial auctions should be examined first. There exists a multitude of WDP models for different scenarios (for an overview, see, e.g., [De Vries and Vohra 2003](#) or [Abrache, Crainic, and Rezik 2007](#)). However, multi-objective formulations of the winner determination problem are still missing. Furthermore, since most of the models are based on set packing or set partitioning formulations of the WDP, the proposed approaches cannot be easily transferred

to the 2WDP-SC which is based on a set covering formulation.

Therefore, related work in the set covering literature may offer more inspiration for finding a suitable approach.

An overview of successful methods for solving (single-objective) set covering problems is given by [Caprara, Toth, and Fischetti \(2000\)](#) and [Umetani and Yagiura \(2007\)](#). In particular, the 3-flip neighborhood search heuristic proposed by [Yagiura, Kishida, and Ibaraki \(2006\)](#) and the multi-start heuristic Meta-RaPS of [Lan, DePuy, and Whitehouse \(2007\)](#) have proven superior to previous approaches. Furthermore, the Bi-objective Set Covering Problem (BOSC) investigated by [Jaszkievicz \(2004\)](#) and [Prins, Prodhon, and Calvo \(2006\)](#) shows similarities to the 2WDP-SC; however, the BOSC differs in the objective function f_2 . Jaszkievicz applied ten different preference-independent methods to the BOSC, including three simulated annealing algorithms, three genetic algorithms and three hybrid genetic algorithms. Furthermore, a multi-start method is proposed which randomly changes the weights of a weighted Tchebycheff scalarizing function at each restart. The multi-start method is also used as construction heuristic in the hybrid genetic algorithms which performed best in the experiments by Jaszkievicz. The preference-independent method proposed by [Prins, Prodhon, and Calvo \(2006\)](#) transforms the BOSC into a single-objective set covering problem by calculating weights and applying Lagrangian relaxation.

From a problem-related point of view, multi-start heuristics have recently recommended themselves as successful methods for solving single-objective and bi-objective set covering problems. Multi-start methods require the outcome of different solutions with each restart. The mentioned multi-start methods for the BOSC apply different scalarization functions in each iteration to find different solutions. In contrast, the well-known multi-start method GRASP ([Feo and Resende 1995](#)) randomizes the greediness in the construction phase to find different solutions. GRASP methods have proven successful in solving a large number of (single-objective) combinatorial optimization problems for many years. It should be noted that the Meta-RaPS heuristic introduced by [Lan, DePuy, and Whitehouse \(2007\)](#) also exhibits strong similari-

ties to GRASP (DePuy, Moraga, and Whitehouse 2005). There exists a small, but increasing number of successful Pareto-based restart-algorithms for other multi-objective problems. In particular, the Pareto-based GRASP approaches of AdensoDíaz, García-Carbajal, and Gupta (2008), Higgins, Hajkowicz, and Bui (2008), Ribeiro and Urrutia (2007), Boudia and Prins (2009), Arroyo, Vieira, and Vianna (2008), Ishida, de Carvalho, Pozo, Goldberg, and Goldberg (2008), Reynolds and de la Iglesia (2009), and Parragh, Doerner, Hartl, and Gandibleux (2009) have to be mentioned. None of these deals with multi-objective set covering problems. Consequently, this article is to propose a Pareto-based GRASP heuristic which does not require the calculation of different weighted scalarizing functions for the 2WDP-SC.

6 A Pareto-based GRASP for the 2WDP-SC

As stated in the previous section, the presented Pareto optimizer is based on the multi-start heuristic GRASP (Feo and Resende 1995). Instead of a single solution, the method searches for a set of non-dominated solutions. For that purpose, it does not require any preference information. The choice of a subjective best solution by applying preference information of the decision-maker is not the focus of this contribution. The reader interested in this topic is referred to, e.g., Vincke (1992).

Before the algorithm is presented, the applied concept of dominance shall be introduced. For notational convenience, the 2WDP-SC is treated as a pure minimization problem in the following. Hence, $\max f_2(X)$ is replaced by $\min f_2(X)$ with $f_2(X) := -f_2(X)$ from now on. Then the required terminology can be introduced as follows (compare, e.g., Deb 2001 or Zitzler, Thiele, Laumanns, Fonseca, and da Fonseca 2003):

Let \mathcal{S} denote the set of feasible solutions and let $\mathcal{Z} \subseteq \mathbb{R}^2$ be the objective function space. It is assumed that every solution $X^i \in \mathcal{S}$ is mapped to a unique objective vector $z^i \in \mathcal{Z}$ with $z^i := (f_1(X^i), f_2(X^i))$.

Given two solutions $X^1, X^2 \in \mathcal{S}$, solution X^1 *weakly dominates* solution X^2 , if every component of z^1 is less or equal than the corresponding component of z^2 (written $X^1 \preceq X^2$). If every component of z^1 is less or equal than the corresponding component of z^2 and at least one

component of z^1 is less than the corresponding component of z^2 , then X^1 *dominates* X^2 (written $X^1 \prec X^2$). If X^1 does not dominate X^2 (written $X^1 \not\prec X^2$) and $X^2 \not\prec X^1$, then X^1 and X^2 are incomparable, i.e., both are not dominated with respect to each other (written $X^1 \sim X^2$).

If a set $A \subseteq \mathcal{S}$ contains only non-dominated solutions with respect to each other, i.e., $X^i \sim X^j$ for all $X^i, X^j \in A, X^i \neq X^j$, it is called (Pareto) *approximation set*. The set which holds all non-dominated solutions with respect to \mathcal{S} is called the Pareto set A^* . The images z^i of the $X^i \in A$ are called *front*. The objective function vectors z^i of the Pareto set A^* are called *Pareto front*.

To shorten notation the *non-dominated union operator* \uplus is introduced. Let A^1, A^2 be two approximation sets. The operator \uplus merges both sets and removes all solutions which are dominated by another solution in the merged set. The outcome of the operator \uplus again is an approximation set:

$$(5) \quad A^1 \uplus A^2 := A^1 \cup A^2 \setminus \{X \in A^1 \cup A^2 \mid \exists X' \in A^1 \cup A^2 \wedge X' \prec X\}$$

Having introduced some notational conventions, an overview of the solution approach is given in Algorithm 1. The basic idea of the algorithm is the repeated application of local search to different initial solutions. The algorithm restarts until a termination criterion is met. In order that the local search procedure attains different regions of the solution space, GRASP constructs initial solutions not deterministically but randomly. As in the present case two objective functions have to be optimized simultaneously; the result of the GRASP is not a single best found solution, but the approximation set attained over all iterations. The algorithm does not require any preference information.

The building blocks of the procedure are described in greater detail below: Section 6.1 presents the construction heuristic (constructSolution). Section 6.2 introduces a clean-up procedure (cleanUp) which removes redundant bids from the constructed solution. Section 6.3 describes the local search procedure (localSearch). Both, the construction heuristic as well as the local search heuristic, apply a rating operator to decide which bid to add to the solution. Section 6.4 explains three alternative rating operators. Finally, section 6.5 gives some details on the hybrid path relinking

procedure.

Algorithm 1: Overview Pareto-GRASP

Input: parameters $rclSize$,
 $nChangeInitialSolution$,
 $nImprovementTrials$, $destroyRate$

Output: approximation set A

```

 $A \leftarrow \{\}$ 
while termination criterion not met do
   $X^{initial} \leftarrow \text{constructSolution}(rclSize)$ 
   $A^{local} \leftarrow \text{cleanUp}(\{X^{initial}\})$ 
   $A^{local} \leftarrow \text{localSearch}(A^{local}, \text{parameters})$ 
   $A \leftarrow A \uplus A^{local}$ 
end
 $A \leftarrow \text{cleanUp}(A)$ 
 $A \leftarrow \text{hybridPathRelinking}(A)$ 

```

6.1 Construction heuristic

The construction heuristic (cf. Algorithm 2) initializes a feasible solution. As long as the constructed solution is infeasible, it is augmented in every construction step by an up-to-now unselected bid. To determine the bid to append, every unselected bid is rated by one of three rating operators RO . Details of these rating operators are given in Section 6.4. After the bids are rated, the $rclSize$ best bids are saved in the so-called restricted candidate list (RCL). Finally, a bid is selected randomly from the RCL and added to the solution. By this stochastic selection step, the important GRASP principle of randomly constructing an initial greedy solution is realized.

Algorithm 2: constructSolution

Input: parameter $rclSize$

Output: feasible solution X

```

 $X \leftarrow \{\}$ 
while  $X$  infeasible do
   $RCL \leftarrow \{\}$ 
  forall  $b \in B \setminus X$  do
    rate  $b$  with  $RO(b, X)$ 
  end
  save the  $rclSize$  best-rated  $b$  in  $RCL$ 
  randomly select a  $b' \in RCL$ 
   $X \leftarrow X \cup \{b'\}$ 
end

```

6.2 Clean-up procedure

Immediately after initializing a solution a simple clean-up procedure tries to improve solution quality by heuristically removing redundant bundle-bids (cf. Algorithm 3).

For that purpose, every bid of a solution is checked in random order if the solution is still feasible after removing the bid. In that case, the bundle-bid is removed and the procedure tries to add the solution X to the approximation set. Afterwards the same principle is applied again to check whether another bundle-bid can be removed from the updated solution.

Algorithm 3: cleanUp

Input: approximation set A

Output: approximation set A'

```

 $A' \leftarrow \{\}$ 
forall  $X \in A$  do
  forall  $b \in X$  do
    if  $X \setminus \{b\}$  is feasible then
       $X \leftarrow X \setminus \{b\}$ 
       $A' \leftarrow A' \uplus \{X\}$ 
    end
  end
end

```

6.3 Local search heuristic

Local search (Algorithm 4) works on a local approximation set A^{local} . First, from this approximation set an unexplored solution X is selected randomly (initial solution). A percentage of $destroyRate$ of all bundle-bids in this initial solution X is removed randomly. In all likelihood, the resulting solution X' is infeasible ($X' \notin S$). Based on this infeasible solution the construction heuristic (Algorithm 2) is applied again (potentially with a different RO) to attain a better (and feasible) solution. In contrast to the construction of an initial solution the parameter $rclSize$ is set to the value one, i.e., a new feasible solution X'' is searched for using a pure greedy strategy.

If ($X'' \not\prec X$), i.e., the new solution X'' does not dominate the initial solution X , then the counter of unsuccessful improvement tries j is incremented by one. If X and X'' are incomparable ($X \sim X''$), then X'' is added to the local approximation set, while possibly dominated solutions are removed

from it.

In the case that the new solution dominates the initial solution ($X'' \prec X$), the counter of unsuccessful improvement tries is reset to zero, and local search recommences. However, in this attempt the new initial solution is set to X'' .

Local search goes in for $nImprovementTrials$ trials to find a solution dominating the initial solution. If this fails, the search process switches the initial solution, i.e., an up-to-now unexplored solution from A^{local} is chosen as the new initial solution.

Fig. 6.3 illustrates the search for a dominant solution. The initial solution of local search (solution X^1) is unchanged until, either within at most $nImprovementTrials$ trials a solution which dominates X^1 is found (e.g., solution X^2 or X^3) or – if this fails – another local non-dominated solution is chosen as the new initial solution (alternation from solution X^3 to solution X^4).

Algorithm 4: localSearch

Input: X , parameters $rclSize$,
 $nChangeInitialSolution$,
 $nImprovementTrials$, $destroyRate$

Output: approximation set A^{local}

```

 $A^{local} \leftarrow \{X\}$ 
for  $i = 1$  to  $nChangeStartSolution$  do
  select unprocessed  $X \in A^{local}$ 
   $j \leftarrow 0$ 
  while  $j \leq nImprovementTrials$  do
     $X' \leftarrow destroySolutionRandomly(X,$ 
     $destroyRate)$ 
     $X'' \leftarrow constructSolution(X',$ 
     $rclSize = 1)$ 
    if  $X'' \prec X$  then
       $X \leftarrow X''$ 
       $j \leftarrow 0$ 
    else
       $j \leftarrow j + 1$ 
      if  $X'' \sim X$  then  $A^{local} \leftarrow A^{local} \uplus \{X''\}$ 
    end
  end
end

```

The multi-start procedure terminates if the appropriate termination criterion is met. This may be, for example, a predefined number of iterations or the expiration of a specified computation time.

6.4 Three rating operators for

bundle-bids

This section presents three concrete variants of the abstract rating operator $RO(b, X)$ introduced in 6.1. The first variant $P(b, X)$ rates a bundle-bid from the cost point of view, whereas the second variant $Q(b, X)$ emphasizes the quality point of view. Finally, the third variant $S(b, X)$ applies both operators $P(b, X)$ and $Q(b, X)$ simultaneously in a vector-based manner. Let $b \in B$ a bundle-bid and X a solution (feasible $X \in S$ or infeasible $X \notin S$); for an operator $RO \in \{P, Q, S\}$ the respective formula applies:

$$(6) \quad P(b, X) = \begin{cases} \frac{p(b)}{|\tau(b) \setminus \tau(X)|} & \text{for } |\tau(b) \setminus \tau(X)| > 0, \\ \infty & \text{for } |\tau(b) \setminus \tau(X)| = 0 \end{cases}$$

$$(7) \quad Q(b, X) = \frac{f_2(X \cup \{b\}) - f_2(X)}{\sum_{b' \in X \cup \{b\}} |\tau(b')|}$$

$$(8) \quad S(b, X) = \begin{pmatrix} P(b, X) \\ Q(b, X) \end{pmatrix}$$


The first rating operator $P(b, X)$, P for short, examines the ratio of the cost $p(b)$ of a bundle-bid b to the number of *new* transport contracts in this bid (Chvátal 1979). A transport contract is denoted as new, if it is not covered by the current (infeasible) solution. The smaller the value of P , the better a bundle-bid b is rated.

The second rating operator $Q(b, X)$, Q for short, rates a bundle-bid with respect to its ability to improve the service quality level f_2 . A decline of f_2 by adding a further bundle-bid is not possible, see (2). The increment of f_2 is divided by the total number of covered transport contracts. As the case may be, contracts covered by multiple bids are counted repeatedly. In order that Q assigns lower values to better bundle-bids, just like P , the fraction is multiplied by -1.

The third rating operator $S(b, X)$, S for short, is composed of the operators P and Q in a vector-valued manner.


For that purpose the dominance concept introduced at the beginning of Section 6 for solutions is transferred correspondingly, without further formalization, to the dominance of bundle-bids rated by S . The RCL contains consequently a set of bundle-bids for which a dominance relation exists. To select the $rclSize$ best bundle-bids a procedure is applied as follows: All bundle-bids are rated by S . The rated bundle-bids form one or more fronts (cf. Fig. 2). From the first front $rclSize$ bundle-bids

Figure 1: Core principles of the local search procedure



are selected randomly. If the first front contains fewer than $rclSize$ bundle-bids, some bids from the second front are selected randomly and so on.

Figure 2: Fronts of bundle-bids with rating operator S



Note the dimensions $P(b, X)$ and $Q(b, X)$ of the axes instead of f_1 and f_2 .

6.5 Hybrid Path Relinking heuristic

As a post-optimization procedure, a hybrid approach based on a combination of heuristic path relinking and exact branch-and-bound is applied. Although path relinking was introduced by Glover (1996) as an adaptive memory strategy for tabu search, it became a well-established intensification strategy within the GRASP framework (Resende and Ribeiro 2003). It is applied either as a mechanism integrated into GRASP after each local search phase to escape from local optima or as

a separate post-optimization step. In both cases, path relinking tries to reach better solutions by combining attributes of high-quality solutions. To find promising combinations of attributes an *initial* solution $X^{initial}$ and a *guiding* solution $X^{guiding}$ are *relinked* by a path through the neighborhood space outlined by the two solutions. A path consists of a sequence of moves where each move adds at least one attribute from the guiding solution to the initial solution; in other words: each move reduces the distance between the initial and the guiding solution.

Algorithm 5: hybridPathRelinking

Input: approximation set A , *criticalHamDist*

Output: approximation set A

while *termination criteria not met* **do**

select $X^{initial}, X^{guiding}$ randomly from A

$B^+ \leftarrow X^{guiding} \setminus X^{initial}$

$B^- \leftarrow X^{initial} \setminus X^{guiding}$

$X^{current} \leftarrow X^{initial}$

while $|B^+| + |B^-| > \text{criticalDistance}$ **do**

reductionOperator($X^{current}, B^+, B^-$)

end

$X^{base} \leftarrow X^{current} \setminus B^-$

$A' \leftarrow \epsilon\text{LookaheadBB}(X^{base}, B^+ \cup B^-)$

$A \leftarrow A \uplus A'$

end

Here, the attributes of a solution are given by its x_b values (cf. section 4). While B^- contains only those bids of the initial solution which are not in the guiding solution (and thus have to be removed from the initial solution in order to reach the guiding solution), B^+ contains all bids of the

guiding solution which are missing in the initial solution (and thus have to be added). The distance between two solutions is equal to the number of attributes by which the two solutions differ, i.e. $|B^-| + |B^+|$.

The approach at hand combines a truncated path relinking algorithm with an exact *subspace search* mechanism (Algorithm 5). The goal of exact subspace search is to find *all* non-dominated solutions that can be obtained by recombining the attributes of the initial and the guiding solution. Basically, exact subspace search is equivalent to exploring all possible paths between the initial and the guiding solution.

We implemented exact subspace search using the ϵ Lookahead-branch-and-bound procedure (ϵ LBB) presented in Buer and Pankratz (2010). This procedure is based on the ϵ -constraint method (Haimes, Lasdon, and Wismer 1971; Chankong and Haimes 1983). The idea of the ϵ -constraint method is to optimize a single objective function, treating the other objective function as a side constraint whose value is bounded by a certain ϵ . To obtain the complete Pareto set for the 2WDP-SC, a proper sequence of such ϵ -constrained single-objective optimization problems is solved for different values of ϵ . Here, the 2WDP-SC is scalarized by treating f_2 as side constraint. The derived single-objective minimization problem consists of the objective function (1) with the covering constraint (3) and the epsilon constraint $f_2(x) < \epsilon$. For further details the reader is referred to the original article.

However, if the distance between the initial and the guiding solution is above a critical level, the exact procedure requires excessive computing time. Therefore, after an initial and a guiding solution have been chosen randomly from the current approximation set, the distance between the guiding and the initial solution is reduced via a truncated path relinking mechanism until a predefined critical distance is reached.

Two reduction operators are available for truncated path relinking (not presented as pseudo code). Starting from $X^{initial}$, the *RandomlyAddAttributes* operator (R for short) in each step randomly selects a bundle-bid from either B^- or B^+ and removes it from or adds it to $X^{current}$, respectively. B^- or B^+ are updated accordingly. This is repeated until the critical distance is reached. Note that intermediate solutions are neither stored nor checked for feasibility. Thus, this procedure may

end up with an infeasible solution after termination.

The *DestroyAndInsertNeighborhood* operator (N for short) also starts with $X^{initial}$. In contrast to the first reduction operator, this operator explicitly aims at creating feasible solutions while exploring the neighborhood. In each step, a random number between 0 and 1 is drawn. If this number is below 0.5, a $b \in B^-$ is randomly selected and removed from $X^{current}$ (and from B^-). This step repeats until the first occurrence of a random number above 0.5. Thereafter, bundle-bids from B^+ are inserted into $X^{current}$ by applying the GRASP construction heuristic (Section 6.1) until the solution gets feasible (B^+ being updated accordingly). To rate candidate bundle-bids during construction, rating operator S is used (cf. section 6.4). If the solution meets the non-dominance property, it is added to the current approximation set. Then the procedure starts over again, until the critical distance is reached.

After termination of truncated path relinking, the solution $X^{current}$ attained at the endpoint of the incomplete path is handed over to the exact subspace search mechanism. The ϵ LBB starts with a base solution X^{base} containing only those bundle-bids which both the current and the guiding solution have in common. The corresponding decision variables x_b are fixed, whereas all x_b whose values are different in the initial and the guiding solution are set free. Note that in the narrower sense, the sequence of search steps within branch-and-bound does not necessarily follow a "path" through some neighborhood. However, the solutions found during branch-and-bound are those solutions, which had been discovered if all possible paths between the initial and the guiding solution had been constructed.

After subspace search via ϵ LBB has finished, the resulting approximation set is joined with the existing approximation set and path relinking starts over again by randomly selecting a new initial and a new guiding solution from the updated Pareto approximation set. The hybrid path relinking procedure terminates according to some reasonable criteria like a predefined amount of computation time.

7 Evaluation

To assess the performance of the developed GRASP variants, several tests with 37 benchmark instances are performed.

At first, Section 7.1 introduces two indicators to assess the solution quality of Pareto optimization procedures. Subsequently, Section 7.2 describes the test instances. Finally, Section 7.3 in detail presents and discusses the results of the computational study.

7.1 Criteria for evaluating the solution quality of Pareto-based search procedures

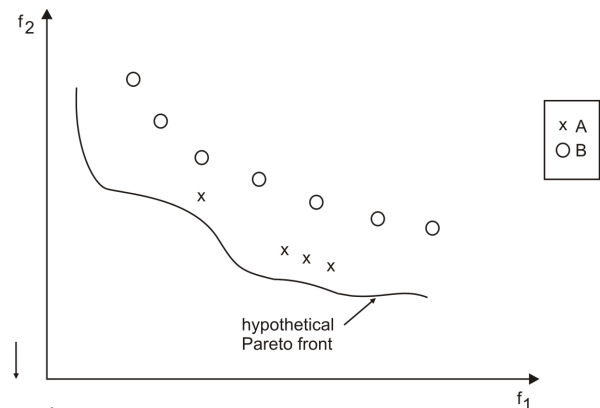
The assessment of an approximation set achieved by a Pareto optimization procedure is usually based on two criteria: The first criterion is *closeness* of the found approximation front to the Pareto front which characterizes the distance of the objective function values of the found solutions to the Pareto front. The smaller this distance, the better the solution quality. The second criterion for quality assessment is the *diversity* of the solutions found. A large diversity exists, if the front spreads wide and is evenly distributed in the direction of all dimensions of the objective function space.

For clarification of closeness and diversity, Fig. 3 shows a hypothetical Pareto front and two approximation fronts A, B . The front A has a small distance to the Pareto front, but features only little diversity. However, the diversity of B is superior, but the closeness of the solutions in B to the Pareto front is inferior with respect to the solutions in A . To measure closeness and diversity of approximation sets, we use two indicators established in the literature: the hypervolume indicator and the epsilon indicator. Both indicators consider the fact that the true Pareto front is unknown. Therefore, they define measures which allow the user to compare approximation sets regardless of the actual Pareto front. A detailed discussion of these concepts can be found in, e.g., [Deb \(2001\)](#) or [Zitzler, Thiele, Laumanns, Fonseca, and da Fonseca \(2003\)](#).

The hypervolume indicator $I_{HV}(A)$ introduced by [Zitzler and Thiele \(1999\)](#) measures the volume of the objective function space which is weakly dominated by all solutions of an approximation set A . To measure the respective volume, the space has to be bounded by a reference point. The reference point


must be weakly dominated by all solutions of the approximation set. The hypervolume is illustrated in Fig. 4 by a gray area. It is evident that this area is always enlarged, if the (Euclidean) distance of a solution to the (unknown) Pareto front is reduced. Since the area also enlarges with each additional non-dominated solution that is found, and grows the bigger, the wider the solutions are spread over the objective space, the hypervolume indicator is able to measure closeness and diversity simultaneously. Therefore, an approximation set A outperforms an approximation set B , if $I_{HV}(A) > I_{HV}(B)$. This is only true, if the chosen reference point is dominated by all solutions in $A \cup B$. To obtain reproducible results for the hypervolume indica-

Figure 3: Closeness and diversity as characteristics of the quality of an approximation set



Approximation front A is closer to the hypothetical Pareto front but less diverse than the front B .

Figure 4: Hypervolume indicator



The light gray area attained by the front B is overlapped by the dark gray area attained by front A . A achieves a higher hypervolume and therefore outperforms B .

tor, for all instances the reference point z^R was set to $z^R := (f_1(B), 0)$. For the test instances introduced in the following section, the choice of that reference point z^R guarantees that z^R is dominated by every feasible solution. Since the range of values of both dimensions is very different in the present instances, each objective function vector z is normalized according to formula (9).

$$(9) \quad z'_i = \frac{z_i - z_i^{\min}}{z_i^R - z_i^{\min}} \quad \text{with } i \in \{1, 2\}.$$

Here, let $z_1^{\min} := 0$ and $z_2^{\min} := f_2(B) - 1$. The reason for the choice of z_2^{\min} is to avoid that a solution which is optimal with respect to f_2 does take on the value zero after normalization. Otherwise, the calculation of the epsilon indicator (see below) would require a division by zero.

Unlike the hypervolume indicator I_{HV} , the epsilon indicator I_ϵ introduced by Zitzler, Thiele, Laumanns, Fonseca, and da Fonseca (2003) does not require a singular reference point, but a second approximation set.

For all solutions in A^2 , $I_\epsilon(A^1, A^2)$ is the minimum factor by which each of the corresponding objective vectors has to be scalar multiplied, such that any of the resulting objective vectors is weakly dominated by the objective vector of at least one solution in A^1 .

Given two approximation sets A^1, A^2 , then $I_\epsilon(A^1, A^2)$ is formally defined as:

$$(10) \quad I_\epsilon(A^1, A^2) = \max_{X^2 \in A^2} \min_{X^1 \in A^1} \max_{1 \leq i \leq n} \frac{f_i(X^1)}{f_i(X^2)}.$$

n denotes the number of objective functions.

Approximation sets of algorithms have to be compared pairwise by means of I_ϵ . Thus to evaluate n algorithms, one had to make $n \cdot (n - 1)$ comparisons. In order to facilitate evaluation, we introduce the reference approximation set A^R which is the non-dominated union set of the approximation sets of all algorithms to be evaluated. Let \mathcal{A} be the set of all algorithms to be evaluated. Then we introduce $I_{\epsilon R}(A)$ which compares the approximation set of a given algorithm $A \in \mathcal{A}$ to the reference approximation set, i.e.,

$$(11) \quad I_{\epsilon R}(A^1) := I_\epsilon(A^1, A^R) \quad \text{with } A^R := \biguplus_{A \in \mathcal{A}} A.$$

The best possible value of $I_{\epsilon R}$ is 1; inferior values are greater than 1.

7.2 Test instances

To test the GRASP variants the test instances of Buer and Pankratz (2010) are used. These satisfy some important domain-specific criteria, which are not guaranteed by the well-known instances for the classical set covering problem, e.g., from ORLIB, or other available test instance generators for winner determination problems, e.g., Leyton-Brown, Pearson, and Shoham (2000) and Leyton-Brown, Nudelman, and Shoham (2002). In particular, the instances comply with the following characteristics of transportation procurement situations (for details of the problem generation procedure, see Buer and Pankratz 2010):

- Bidders compose bundle-bids according to synergies between the contracts. These synergies are highly carrier-specific and depend on several factors like the geographical dispersion of the carrier's depots, other already existing transportation commitments, as well as the size and the mixture of the fleet. Thus, a carrier-specific synergy matrix with a preset density ρ indicating the occurrence of pairwise synergies between contracts is used when generating bids.
- In order to maintain the opportunity of being awarded, carriers predominantly bid on bundles of contracts which exhibit such synergies. Hence, the prices of all bundle-bids B^c submitted by bidder c are *strictly subadditive*. The price of a set of contracts τ is called strictly subadditive, if for each partition of set τ the price of τ is strictly lower than the sum of the prices of all parts of the respective set partition (recall that it is assumed that prices are always equal to the bidder's true valuations). Formally this is expressed in (12), in which $\mathbf{P}(\tau)$ denotes the powerset of τ :

$$(12) \quad \forall \mathcal{T} \subseteq \mathbf{P}(\tau) : \bigcup_{\tau' \in \mathcal{T}} \tau' = \tau \wedge \bigcap_{\tau' \in \mathcal{T}} \tau' = \emptyset \\ \wedge p^c(\tau) < \sum_{\tau' \in \mathcal{T}} p^c(\tau').$$

- Furthermore, the prices of a bundle-bid comply with the *free disposal* assumption, i.e., the price charged by a carrier c for a set of contracts $\tau_b, b \in B^c$ is always as least as high as every

subset of τ_b :

$$(13) p(\tau_{b'}) \leq p(\tau_b) \mid \forall \tau_{b'} \subseteq \tau_b \wedge b \in B^c.$$

As explained in Section 4, free disposal can be assumed in the procurement scenario at hand. Therefore the used instances should also feature free disposal.

- Finally, a positive correlation between the service quality of a contract specific to a bidder and the charged price for its execution is assumed, i.e., the better the service quality of a contract, the higher the demanded price tends to be.

The test instances divide up into two subsets of different problem size.

The first subset contains 30 problem instances of practical size as reported in Section 1. These instances are referred to as *large instances*. For the number $|B|$ of bundle-bids, the values 500, 1000 and 2000 were assumed. The number $|T|$ of transportation contracts tendered varied between 125, 250 and 500; the number $|C|$ of participating carriers in the procurement auction were 25, 50 and 100, respectively. Finally, the density ρ of the bidder-specific synergy matrices was set to 25, 50 and 75 percent, respectively. In practice, the number of bids submitted depends largely on the number of tendered contracts and these in turn affect the number of participating carriers in the auction. Therefore, only the parameter combinations shown in Tab. 3 (lines 8 through 37) are considered. The second subset consists of the seven *small instances* for which the optimal results have been reported in [Buer and Pankratz \(2010\)](#). The respective parameter combinations are displayed in lines 1 through 7 of Tab. 3.

7.3 Computational study

The GRASP algorithm and the hybrid path re-linking heuristic were implemented in Java (SDK 1.6). All tests were executed on a standard PC (Intel Pentium 4, 2 GHz, 768 MB RAM). Computational results and the used benchmark instances are provided in the electronic appendix.¹ During development and initial testing, the parameter values shown in Tab. 1 were identified as reasonable

¹ These files can be downloaded from www.business-research.org.

Table 1: Parameter values for all GRASP variants

Parameter	Value
target runtime	300s
relSize	10
destroyRate	10%
nChangeStartSolution	10
nImprovementTrials	5

and held constant throughout all tests of the computational study.

7.3.1 Test set up

The GRASP algorithm proposed incorporates three stages (construction, local search, path relinking) for each of which one of the operators introduced in Sections 6.4 and 6.5 has to be chosen. Let $G_{ijk} \in \mathcal{G}$ denote the GRASP instance which uses the rating operator $i \in \{P, Q, S\}$ in its construction stage, the rating operator $j \in \{_, P, Q, S\}$ in its local search stage, and the relinking operator $k \in \{_, R, N\}$ in its path relinking stage. The underscore $_$ in the domains of j and k means, that the respective stage is deactivated. For example, the GRASP variant $G_{SP_}$ uses the rating operator $S(b, X)$ in its construction stage, the rating operator $P(b, X)$ during local search, and the path relinking phase is not used at all.

The computational study is divided up into a preliminary study and a benchmark study. The goal of the *preliminary study* is to identify the best GRASP variants, i.e. the best configurations of the operators denoted by i, j, k . Solution quality is measured in terms of I_{HV} and I_{eR} as introduced in Section 7.1. Note that a performance ranking of GRASP variants according to I_{HV} will generally not be consistent with a performance ranking according to I_{eR} . This is indebted to the multi-objective optimization situation and will generally result in a set of non-dominated GRASP configurations. In the final *benchmark study*, the performance of the non-dominated GRASP variants is compared to an exact solution approach (for the small instances) and to an evolutionary algorithm (for the small and the large instances), both of which are from [Buer and Pankratz \(2010\)](#).

All experimental results are based on a single run of each considered GRASP variant per instance. For each run, a target computing time of 5 minutes (300s) was allowed. As termination criterion

the first excess of this target time was used. The termination criterion was checked at four break points: Firstly, before every restart of a method; secondly and thirdly, within local search (Algorithm 4) immediately before completion of any iteration of the for- or the while-loop, respectively. Finally, the termination criterion is checked within the ϵ LBB of the path relinking phase before a new sub-problem with an incremented ϵ -value in the added side constraint is solved. In case of termination, all solutions found by then were added to the global approximation set. All in all, the variation of computing times caused by this termination policy is small and seems not to affect the validity of the computational study.

Against the background of the business scenario, computing times of five minutes may seem short, at first sight. As the considered tenderings occur in intervals of several months or years, in principle longer computing times seem justifiable. On the other hand, one has to take into account that in practice, some side constraints are not fully defined beforehand, but are worked out in the course of the allocation stage. Such implicit side constraints are, e.g., an ad-hoc restriction of the number of carriers assigned, or tacit business guarantees. Therefore, often a series of slightly modified models have to be solved until an appropriate solution is found (Caplice and Sheffi 2006: 567), which makes shorter computing times preferable.

7.3.2 Preliminary Tests

Tab. 2 shows the results of five preliminary tests. Each test examines different GRASP variants which differ with regard to the operator chosen for one of the three stages construction, local search, or path relinking. For each test, the respective rank orders according to I_{HV} and I_{eR} are given for every algorithm. All five tests are based on the set of the 30 large 2WDP-SC instances.


Note that both rank orders are obtained by systematic pairwise comparison of the results on a per-instance basis by means of the Wilcoxon signed rank test (level of significance five percent). If for two algorithms the null hypothesis cannot be rejected, they are assigned the same average rank. The average indicator values \bar{I}_{HV} , \bar{I}_{eR} over the thirty instances are provided for illustrative purposes only. Detailed results for each instance are provided in the electronic appendix.

Choosing a rating operator for the construction stage. The first test answers the question of which of the three rating operators performs best during the construction stage (Tab. 2, Test No. 1). Both indicators show that there is a consistent ranking order at which procedure $G_{P_}$ performs worst and procedure $G_{S_}$ performs best. Obviously, considering cost and quality simultaneously not only equilibrates both objectives, but also boosts the effect of both operators in the construction phase.

Fig. 5 reveals some interesting insights into these performance differences using the example of a representative instance. Although $G_{P_}$ and $G_{Q_}$ specialize on minimizing cost and quality, respectively, they are clearly outperformed in terms of both cost and quality by $G_{S_}$.

Choosing a rating operator for the local search stage. The second test tries to identify the best rating operator for the local search stage. Due to the results of the previous test, the construction operator was uniformly set to S . This time, the rank orders determined by the two indicators are not consistent (Tab. 2, Test No. 2). If the preferences of the decision-maker are best represented by I_{HV} , algorithm $G_{SP_}$ should be chosen. Else, the decision-maker should chose $G_{SQ_}$. By comparing the results with those of Test No. 1, two further conclusions can be deduced: (i) local search considerably improves the mean indicator values. (ii)

Figure 5: Approximation fronts of the construction-only GRASP variants



Sample instance | B | = 2000; | T | = 125; | C | = 25; $\rho=0.5$.

Table 2: Results of preliminary tests

Test	Algorithm	\bar{I}_{HV}	\bar{I}_{eR}	Rank	
				I_{HV}	I_{eR}
1	$G_{P_}$	0.8489	56.47	3	3
	$G_{Q_}$	0.8838	8.15	2	2
	$G_{S_}$	0.9182	2.92	1	1
2	$G_{SP_}$	0.9225	3.09	1	3
	$G_{SQ_}$	0.9205	1.14	2.5	1
	$G_{SS_}$	0.9194	1.63	2.5	2
3a	$G_{SPR, 0}$	0.9224	3.19	3	2.5
	$G_{SPR, 20}$	0.9232	3.01	1	1
	$G_{SPR, 40}$	0.9224	3.18	2	2.5
3b	$G_{SPN, 0}$	0.9224	3.19	3	3
	$G_{SPN, 20}$	0.9231	2.94	1	1
	$G_{SPN, 40}$	0.9224	3.18	2	2
3c	$G_{SPR, 20}$	0.9232	3.01	1	1.5
	$G_{SPN, 20}$	0.9231	2.94	2	1.5
4a	$G_{SQR, 0}$	0.9202	1.14	2.5	2.5
	$G_{SQR, 20}$	0.9210	1.13	1	1
	$G_{SQR, 40}$	0.9201	1.14	2.5	2.5
4b	$G_{SQN, 0}$	0.9201	1.14	3	*
	$G_{SQN, 20}$	0.9206	1.13	1	1*
	$G_{SQN, 40}$	0.9203	1.14	2	*
4c	$G_{SQR, 20}$	0.9210	1.13	1	1
	$G_{SQN, 20}$	0.9206	1.13	2	2
5	$G_{SPR, 20}$	0.9232	3.01	1	2
	$G_{SQR, 20}$	0.9210	1.13	2	1

*) The asterisked entries of Test No. 4b indicate no consistent rank order. Indeed, $G_{SQN,20}$ outperformed $G_{SQN,0}$, but the other comparisons ($G_{SQN,0}$ vs $G_{SQN,40}$, $G_{SQN,20}$ vs $G_{SQN,40}$) tied regarding I_{eR} and the Wilcoxon signed rank test.

It is not the most promising search strategy to permanently seek to improve both objective criteria simultaneously throughout the whole search process, which would have been an obvious hypothesis in view of the Test No. 1 results.

Choosing a hybrid path relinking variant.

For hybrid path relinking, Tests No. 3 and 4 aim at identifying the best reduction operator and a suitable degree of hybridization in terms of the critical distance level. Distances of 0, 20 and 40 were tested (cf. corresponding entries in the second column of Tab. 2). A distance value of zero

represents the special case of a non-hybridized path relinking approach which does not use exact subspace search.

The total run time of five minutes was divided into four minutes for the GRASP procedure and one minute for the hybrid path relinking procedure.

With respect to the relatively short computation time assigned to the path relinking phase, the parameter $rclSize$ of the GRASP construction procedure (Algorithm 2) was decreased to 5 when applied in the context of the DestroyAndInsert-Neighborhood operator.

As Tests No. 3a/b and No. 4a/b show, the path relinking variants which use a critical distance of 20 are the most effective. Therefore, only these variants are discussed in the following. Test No. 3c discovers, that G_{SPR} outperforms G_{SPN} , and Test 4c shows that G_{SQR} outperforms G_{SQN} . Finally, those two best algorithms compete in Test No. 5 against each other: there is, however, no unique rank order as G_{SPR} is best regarding I_{HV} and G_{SQR} is superior regarding I_{eR} .

7.3.3 Final Benchmark Test

In this section, the best variants G_{SPR} and G_{SQR} are compared to an evolutionary algorithm (EA) and to an exact branch-and-bound algorithm from [Buer and Pankratz \(2010\)](#).

Solution Quality. The results of the final benchmark tests are presented in Tab. 3. For each of the algorithms, the I_{HV} and I_{eR} values for each instance are given. To further illustrate the results, $I_{\#}$ indicates how many solutions the respective algorithm contributed to the reference approximation set.

The last two columns of Tab. 3 report the I_{HV} values and the cardinalities of the corresponding reference approximation sets, respectively. Asterisked values indicate that they refer to the (optimal) Pareto set found by the exact branch-and-bound method of [Buer and Pankratz \(2010\)](#). Note that for each of the large instances (lines 8 through 37 in Tab. 3), the corresponding reference approximation set comprises all non-dominated solutions on the respective instance including those found during the preliminary tests.

All results, as in the preliminary tests, have been statistically evaluated using the Wilcoxon signed rank test. Statistical conclusions are stated at a

significance level of five percent.

As to the small instances (lines 1 through 7 in Tab. 3), the findings can be summarized as follows: G_{SPR} as well as G_{SQR} consistently outperform the EA with respect to both I_{HV} and I_{eR} . While the EA is able to find the complete Pareto set of the smallest instance only, G_{SPR} as well as G_{SQR} are able to find the optimal Pareto set of three further instances, respectively. Furthermore, G_{SPR} finds more solutions from the Pareto set than G_{SQR} in three instances, whereas G_{SQR} finds more optimal solutions than G_{SPR} only once. However, regarding the few small instances, neither the I_{HV} nor the I_{eR} values give statistical evidence that one of the two GRASP variants dominates the other.

Regarding the large instances ($|B| \geq 500$, starting from line 8 in Tab. 3), G_{SPR} and G_{SQR} again perform significantly better than the EA with respect to both I_{HV} and I_{eR} . In contrast to the small instances and as stated in Section 7.3.2 (Test No. 5), G_{SPR} dominates G_{SQR} according to I_{HV} , and G_{SQR} dominates G_{SPR} regarding I_{eR} . This statement is still true, if the results of all 37 instances are tested by the Wilcoxon rank test.

Robustness. In order to judge robustness of the competing approaches, the mean, the standard deviation and selected quantiles for each indicator and each algorithm are listed below the bottom rule of Tab. 3. Furthermore, the median values are shown in Fig. 6. This illustration also shows for each of the three algorithms the distance between the 25%-quantile and the 75%-quantile.

In accordance to the previous results, one can notice from Fig 6, that G_{SPR} and G_{SQR} dominate the EA with respect to the median values of I_{HV} and I_{eR} . The deviation from the I_{HV} -median is very similar with all three algorithms. However, the deviation from the median I_{eR} differs clearly. G_{SQR} achieves a significantly lesser deviation than both G_{SPR} and EA.


8 Summary and future research

This paper investigated the bi-objective winner determination problem in combinatorial transportation procurement auctions (2WDP-SC) as introduced in Buer and Pankratz (2010). In contrast to previous models for winner determination in transportation procurement auctions, the 2WDP-SC not only focuses on the objective of minimizing

total costs within the allocation phase but also considers the maximization of total service quality. In doing so, the model accommodates the important fact that the tendered contracts have different requirements which can be handled unequally well by different carriers.

To solve the 2WDP-SC, a Pareto-based GRASP was introduced with a post-optimization procedure which hybridizes truncated path relinking with exact branch-and-bound. Combinations of three different rating operators for bundle-bids and two reduction operators for the path relinking procedure were used to create a number of GRASP variants. The performance of these variants was evaluated using the seven small and thirty large instances introduced by Buer and Pankratz (2010). First, during a series of preliminary tests two best GRASP variants were identified: The first variant performs better according to the hypervolume indicator; the second variant reaches superior results as to the epsilon indicator. As a rule, those algorithms performed best, which consider both objectives during the construction stage and afterwards focus only on a single objective in local search. Furthermore, using a randomized reduction operator for the truncated path relinking phase proved best. In a final benchmark test the performance of the two best GRASP variants was tested against two previously published approaches. On the one hand, the results of the GRASP variants are compared to known optimal solutions by means of small

Figure 6: Median indicator values for G_{SPR} , G_{SQR} , and EA.



The lines indicate the distances between the 25%- and the 75%-quantile of the indicator values over all 37 test instances.



Table 3: Results of the final benchmark test for the 37 test instances (specified by columns 1 to 4). Asterisked values in the two leftmost columns are optimal.

Instance				G_{SPR}			G_{SQR}			EA			REF		
B	T	C	ρ	I_{HV}	I_{eR}	$I_{\#}$	I_{HV}	I_{eR}	$I_{\#}$	I_{HV}	I_{eR}	$I_{\#}$	I_{HV}	$I_{\#}$	
20	5	10	50	0.8576	1.00	7	0.8576	1.00	7	0.8576	1.00	7	0.8576*	7*	
				0.6095	1.00	11	0.6095	1.00	11	0.6029	1.03	6	0.6095*	11*	
	40	20	0.8168	1.01	12	0.8169	1.00	13	0.8126	1.47	6	0.8169*	13*		
60	40	20	50	0.5677	1.00	12	0.5677	1.00	12	0.5639	1.25	3	0.5677*	12*	
				0.8650	1.03	12	0.8626	1.03	11	0.8537	2.00	5	0.8652*	17*	
80	20	40	50	0.6988	1.00	10	0.6932	1.02	8	0.6913	1.09	0	0.6988*	10*	
				0.8911	1.11	10	0.8888	1.04	6	0.8872	1.66	2	0.8915*	17*	
500	125	25	25	0.8992	1.20	12	0.8923	1.08	22	0.8914	1.32	0	0.8997	51	
				0.9114	1.13	3	0.9053	1.08	7	0.9038	1.48	0	0.9115	28	
				0.9036	1.15	21	0.8958	1.09	3	0.8983	1.26	0	0.9038	35	
1000	125	25	25	0.9513	1.24	31	0.9498	1.04	30	0.9479	1.33	0	0.9516	69	
				0.9566	1.38	22	0.9531	1.11	14	0.9508	1.50	0	0.9571	63	
				0.9587	1.25	13	0.9544	1.13	11	0.9535	1.34	0	0.9589	41	
	250	25	25	0.9037	1.15	27	0.8993	1.05	17	0.9021	1.27	0	0.9038	44	
				0.9038	1.16	9	0.8972	1.08	19	0.9001	1.28	0	0.9042	47	
				0.9010	1.18	19	0.8970	1.04	17	0.8961	1.27	0	0.9011	37	
	500	50	25	25	0.9064	2.00	14	0.8995	2.00	25	0.8927	10.00	10	0.9081	105
					0.8984	1.38	5	0.8944	1.05	42	0.8943	1.27	0	0.8991	66
					0.9004	1.17	16	0.8962	1.05	35	0.8937	1.49	0	0.9008	54
2000	125	25	25	0.9694	6.00	4	0.9771	1.13	0	0.9720	4.00	0	0.9792	66	
				0.9807	1.24	15	0.9789	1.13	23	0.9786	1.56	0	0.9809	41	
				0.9756	3.00	23	0.9783	1.08	3	0.9746	2.00	0	0.9796	34	
	250	25	25	0.9552	1.11	19	0.9534	1.07	0	0.9531	1.20	0	0.9561	60	
				0.9546	1.26	17	0.9530	1.04	14	0.9532	1.26	0	0.9547	34	
				0.9531	1.11	1	0.9528	1.03	1	0.9510	1.28	0	0.9539	14	
	500	50	25	25	0.9461	10.00	29	0.9523	1.08	4	0.9337	21.00	0	0.9550	88
					0.9573	1.36	16	0.9534	1.11	18	0.9530	3.00	0	0.9576	59
					0.9532	1.32	2	0.9503	1.08	23	0.9508	2.00	0	0.9536	53
500	25	25	25	0.9067	1.02	0	0.9046	1.02	0	0.9022	1.19	0	0.9068	9	
				0.8997	1.01	0	0.8965	1.04	0	0.8942	1.16	0	0.9003	4	
				0.8962	1.00	4	0.8949	1.01	0	0.8944	1.13	0	0.8962	7	
	50	25	25	0.8928	1.15	23	0.8897	1.03	0	0.8902	2.00	0	0.8929	53	
				0.8979	1.16	13	0.8917	1.06	13	0.8947	1.24	3	0.8981	30	
				0.8925	1.18	4	0.8922	1.01	26	0.8887	1.22	0	0.8927	39	
	100	25	25	25	0.8893	5.00	8	0.8883	1.03	51	0.8708	41.00	0	0.8915	101
					0.8855	32.00	8	0.8959	2.00	61	0.8827	32.00	0	0.8988	108
					0.8942	5.00	26	0.8927	1.04	16	0.8939	1.27	4	0.8964	92
mean				0.8919	2.63	12.9	0.8899	1.10	15.2	0.8872	4.16	1.2	0.8933	43.8	
standard deviation				0.0893	5.29	8.4	0.0894	0.22	14.4	0.0900	8.65	2.5	0.0899	29.4	
10% quantile				0.8413	1.00	2.6	0.8413	1.01	0.0	0.8373	1.15	0.0	0.8413	9.6	
25% quantile				0.8925	1.11	7.0	0.8917	1.03	4.0	0.8887	1.25	0.0	0.8929	17.0	
median				0.9010	1.18	12.0	0.8965	1.05	13.0	0.8947	1.33	0.0	0.9011	41.0	
75% quantile				0.9531	1.36	19.0	0.9523	1.08	22.0	0.9508	2.00	0.0	0.9539	60.0	
90% quantile				0.9579	5.00	24.2	0.9538	1.13	32.0	0.9533	6.40	5.4	0.9581	89.6	



instances. On the other hand, they are evaluated against the results of an evolutionary algorithm using both the small and the large instances. For all instances, both GRASP variants consistently outperform the evolutionary algorithm. For all small instances, the GRASP algorithms find more often the whole Pareto set and all in all larger fractions of the optimal Pareto set.

From a managerial perspective, the presented GRASP approach offers the possibility to significantly improve the performance of transportation procurement auctions when both cost and quality have to be considered simultaneously. In following the Pareto principle, our approach does not force the shipper to subjectively weigh cost versus quality beforehand. Instead, it efficiently searches for solutions which are non-dominated in an objective sense. Thus, it gives the shipper valuable insights into the trade-off between transportation cost and transportation service quality. From the set of non-dominated solutions found by the GRASP, the shipper then can choose the solution which best trades off cost against quality according to his or her preferences. As we were able to show, the proposed GRASP approach clearly outperforms an evolutionary algorithm previously published for solving the same problem.

Ongoing and future research is focusing on additional ways of combining metaheuristic and exact solution approaches for solving the 2WDP-SC. In particular, we want to use the exact branch-and-bound approach already during the GRASP construction stage. Furthermore, we want to experiment with other neighborhood structures and we hope to improve solution quality by augmenting our GRASP with techniques of self-adaptive parameter control.

Acknowledgement

We thank the two anonymous referees for carefully examining our work and for their valuable suggestions which contributed to improving the paper.

References

- Abrache, Jawad, Teodor Crainic, Michel Gendreau, and Monia Rekik (2007): Combinatorial Auctions, *Annals of Operations Research*, 153 (34): 131–164.
- Adenso-Díaz, Belarmino, Santiago García-Carbajal, and Surendra M. Gupta (2008): A Path-Relinking Approach for a Bi-Criteria Disassembly Sequencing Problem, *Computers & Operations Research*, 35 (12): 3989–3997.
- Arroyo, José Elias Claudio, Pedro Sampaio Vieira, and Dalessandro Soares Vianna (2008): A GRASP Algorithm for the Multi-Criteria Minimum Spanning Tree Problem, *Annals of Operations Research*, 159 (1): 125–133.
- Bérubé, Jean-François, Michel Gendreau, and Jean-Yves Potvin (2009): An exact ε -Constraint Method for Bi-Objective Combinatorial Optimization Problems: Application to the Travelling Salesman Problem with Profits, *European Journal of Operational Research*, 194 (1): 39–50.
- Boudia, M. and C. Prins (2009): A Memetic Algorithm with Dynamic Population Management for an Integrated Production-Distribution Problem, *European Journal of Operational Research*, 195 (3): 703–715.
- Branke, Jürgen, Kalyanmoy Deb, Kaisa Miettinen, and Roman Slowinski (2008): *Multiobjective Optimization: Interactive and Evolutionary Approaches*, Springer: Berlin.
- Buer, Tobias and Giselher Pankratz (2010): Solving a Bi-Objective Winner Determination Problem in a Transportation Procurement Auction, *Logistics Research*, 2 (2): 65–78.
- Cantillon, Estelle and Martin Pesendorfer (2006): Auctioning Bus Routes: The London Experience, in: Peter C. Cramton, Yoav Shoham, and Richard Steinberg (eds.): *Combinatorial auctions*, MIT-Press: Cambridge (MA), 573–591.
- Caplice, Chris and Yossi Sheffi (2003): Optimization-Based Procurement for Transportation Services, *Journal of Business Logistics*, 24 (2): 109–128.
- Caplice, Chris and Yossi Sheffi (2006): Combinatorial Auctions for Truckload Transportation, in: Peter C. Cramton, Yoav Shoham, and Richard Steinberg (eds.): *Combinatorial Auctions*, MIT-Press: Cambridge (MA), 539–571.
- Caprara, Alberto, Paolo Toth, and Matteo Fischetti (2000): Algorithms for the Set Covering Problem, *Annals of Operations Research*, 98 (1–4): 353–371.
- Cargoclix Dr. Meier & Schmidt GmbH (2010): Success Stories, <http://cargoclix.com/info/de/success-stories/page.html> (Access date: 2010–03–22).
- Chankong, Vira and Yacov Y. Haimes (1983): *Multi-objective Decision Making: Theory and Methodology*, North-Holland: Amsterdam.
- Chu, Leon Yang (2009): Truthful Bundle/Multiunit Double Auctions, *Management Science*, 55 (7): 1184–1198.



- Chvátal, Vašek (1979): A Greedy Heuristic for the Set-Covering Problem, *Mathematics of Operations Research*, 4 (3): 233–235.
- Coello, Carlos A. Coello, Arturo Hernández Aguirre, and Eckart Zitzler (eds.) (2005): *Evolutionary Multi-Criterion Optimization: Third International Conference, EMO 2005, Guanajuato, Mexico, March 9–11, 2005: Proceedings*, Springer: Berlin.
- De Vries, Sven, and Rakesh V. Vohra (2003): Combinatorial Auctions: A Survey, *INFORMS Journal on Computing*, 15 (3): 284–309.
- Deb, Kalyanmoy (2001): *Multi-Objective Optimization Using Evolutionary Algorithms*, Wiley: Chichester.
- Deb, Kalyanmoy, Samir Agarwal, Amrit Partap, and T. Meyarivan (2002): A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II, *IEEE Transactions on Evolutionary Computation*, 6 (2): 182–197.
- DePuy, Gail W., Reinaldo J. Moraga, and Gary E. Whitehouse (2005): Meta-RaPS: A Simple and Effective Approach for Solving the Traveling Salesman Problem, *Transportation Research Part E: Logistics and Transportation Review*, 41 (2): 115–130.
- Dhaenens, C., J. Lemesre, and E. Talbi (2010): K-PPM: A New Exact Method to Solve Multi-Objective Combinatorial Optimization Problems, *European Journal of Operational Research*, 200 (1): 45–53.
- Doerner, Karl F., Walter J. Gutjahr, and Pamela C. Nolz (2009): Multi-Criteria Location Planning for Public Facilities in Tsunami-Prone Coastal Areas, *OR Spectrum*, 31 (3): 651–678.
- Ehrgott, Matthias, Carlos M. Fonseca, Xavier Gandibleux, Jin-Kao Hao, and Marc Sevaux (eds.) (2009): *Evolutionary Multi-Criterion Optimization: Fifth International Conference, EMO 2009, Nantes, France, April 2009: Proceedings*, Springer: Berlin.
- Ehrgott, Matthias and Xavier Gandibleux (2000): A Survey and Annotated Bibliography of Multiobjective Combinatorial Optimization, *OR Spectrum*, 22 (4): 425–460.
- Elmaghraby, Wedad and Pinar Keskinocak (2004): Combinatorial Auctions in Procurement, in: Terry P. Harrison, Hau L. Lee, and John J. Neale (eds.): *The Practice of Supply Chain Management: Where Theory and Application Converge*, Springer: New York, 245–258.
- Feo, Thomas A. and Mauricio G. Resende (1995): Greedy Randomized Adaptive Search Procedures, *Journal of Global Optimization*, 6 (2): 109–133.
- Fischer, Klaus, Jörg P. Müller, Markus Pischel, and Darius Schier (1995): A Model for Cooperative Transportation Scheduling, in: Victor Lesser (ed.): *ICMAS-95, First International Conference on Multi-Agent Systems, Proceedings, June 12-14, 1995, San Francisco, California*, AAAI Press: Menlo Park, 109–116.
- Gandibleux, Xavier and Matthias Ehrgott (2005): 1984–2004 – 20 Years of Multiobjective Metaheuristics. But What About the Solution of Combinatorial Problems with Multiple Objectives? in: Carlos A. Coello Coello, Arturo Hernández Aguirre, and Eckart Zitzler (eds.): *Evolutionary Multi-Criterion Optimization: Third International Conference, EMO 2005, Guanajuato, Mexico, March 9–11, 2005: Proceedings*, Springer: Berlin, 33–46.
- Gass, Saul and Thomas Saaty (1955): The Computational Algorithm for the Parametric Objective Function, *Naval Research Logistics Quarterly*, 2 (1-2): 39–45.
- Glover, Fred (1996): Tabu Search and Adaptive Memory Programming – Advances, Applications, and Challenges, in: Richard S. Barr, Richard V. Helgason, and Jeffery L. Kennington (eds.): *Interfaces in Computer Science and Operations Research: Advances in Metaheuristics, Optimization, and Stochastic Modeling Technologies*, Kluwer: Dordrecht, 1–75.
- Gomber, Peter, Claudia Schmidt, and Christof Weinhart (1997): Elektronische Märkte für die dezentrale Transportplanung, *Wirtschaftsinformatik*, 39 (2): 137–145.
- Haimes, Yacov Y., L. Lasdon, and D. Wismer (1971): On a Bicriterion Formulation of the Problems of Integrated System Identification and System Optimization, *IEEE Transactions on System, Man and Cybernetics*, 1 (3): 296–297.
- Higgins, Andrew J., Stefan Hajkovicz, and Elisabeth Bui (2008): A Multi-Objective Model for Environmental Investment Decision Making, *Computers & Operations Research*, 35 (1): 253–266.
- Ishida, Celso Y., Andre B. De Carvalho, Aurora T. R. Pozo, Elisabeth F. G. Goldberg, and Marco C. Goldberg (2008): Exploring Multi-Objective PSO and GRASP-PR for Rule Induction, in: Jano Van Hemert and Carlos Cotta (eds.): *Proceedings of the 8th European Conference on Evolutionary Computation in Combinatorial Optimization 2008, Naples, Italy, March 26–28, 2008*, Springer: Berlin, 73–84.
- Jaszkievicz, Andrzej (2004): A Comparative Study of Multiple-Objective Metaheuristics on the Bi-Objective Set Covering Problem and the Pareto Memetic Algorithm, *Annals of Operations Research*, 131 (1–4): 135–158.
- Kopfer, Herbert and Giselher Pankratz (1999): Das Groupage-Problem kooperierender Verkehrsträger, in:



- Peter Kall and Hans-Jakob Lüthi (eds.): *Operations Research Proceedings 1998*, Springer: Berlin, 453–462.
- Lan, Guanghui, Gail W. DePuy, and Gary E. Whitehouse (2007): An Effective and Simple Heuristic for the Set Covering Problem, *European Journal of Operational Research*, 176 (3): 1387–1403.
- Ledyard, John O., Mark Olson, David Porter, Joseph A. Swanson, and David P. Torma (2002): The First Use of a Combined-Value Auction for Transportation Services, *Interfaces*, 32 (5): 4–12.
- Leyton-Brown, Kevin, Eugene Nudelman, and Yoav Shoham (2002): Learning the Empirical Hardness of Optimization Problems: The Case of Combinatorial Auctions, in: *Proceedings of the International Conference on Principles and Practice of Constraint Programming CP 2002, Ithaca, 2002*, Springer: London, 556–572.
- Leyton-Brown, Kevin, Mark Pearson, and Yoav Shoham (2000): Towards a Universal Test Suite for Combinatorial Auction Algorithms, in: *Proceedings of the 2nd ACM Conference on Electronic (EC-00), Minneapolis, 2000*, ACM: New York, 66–76.
- Mezmaz, Mohand, Nouredine Melab, and El-Ghazali Talbi (2007): Combining Metaheuristics and Exact Methods for Solving Exactly Multi-Objective Problems on the Grid, *Journal of Mathematical Modelling and Algorithms*, 6 (3): 393–409.
- Miettinen, Kaisa (2008): Introduction to Multiobjective Optimization: Noninteractive Approaches, in: Jürgen Branke, Kalyanmoy Deb, Kaisa Miettinen, and Roman Slowinski (eds.): *Multiobjective Optimization: Interactive and Evolutionary Approaches*, Springer: Berlin, 1–26.
- Obayashi, Shigeru, Kalyanmoy Deb, Carlo Poloni, Tomoyuki Hiroyasu, and Tadahiko Murata (eds.) (2007): Evolutionary Multi-Criterion Optimization: *4th International Conference, EMO 2007, Matsushima, Japan, March 5–8, 2007, Proceedings*, Springer: Berlin.
- Pankratz, Giselher (2000): Analyse kombinatorischer Auktionen für ein Multi-Agentensystem zur Lösung des Groupage-Problems kooperierender Speditionen, in: Karl Inderfurth, Wolfgang Domschke, Friedrich Juhnke, Peter Kleinschmidt, and Gerhard Wäscher (eds.): *Operations Research Proceedings 1999*, Springer: Berlin, 443–448.
- Parragh, Sophie N., Karl F. Doerner, Richard F. Hartl, and Xavier Gandibleux (2009): A Heuristic Two-Phase Solution Approach for the Multi-Objective Dial-A-Ride Problem, *Networks*, 54 (4): 227–242.
- Pekeč, Aleksandar and Michael H. Rothkopf (2003): Combinatorial Auction Design, *Management Science*, 49 (11): 1485–1503.
- Prins, Christian, Caroline Prodhon, and Roberto Wolfler Calvo (2006): Two-Phase Method and Lagrangian Relaxation to Solve the Bi-Objective Set Covering Problem, *Annals of Operations Research*, 147 (1): 23–41.
- Resende, Mauricio and Celso C. Ribeiro (2003): Greedy Randomized Adaptive Search Problems, in: Fred Glover and Gary Kochenberger (eds.): *Handbook of Metaheuristics*, Kluwer: Dordrecht, 219–249.
- Reynolds, Alan P. and Beatriz de la Iglesia (2009): A Multi-Objective GRASP for Partial Classification, *Soft Computing*, 13 (3): 227–243.
- Ribeiro, Celso C. and Sebastián Urrutia (2007): Heuristics for the Mirrored Traveling Tournament Problem, *European Journal of Operational Research*, 179 (3): 775–787.
- Sandholm, Tuomas (1993): An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations, in: *Proceedings of the 11th National Conference on Artificial Intelligence, Washington D.C., July 11–15, 1993*, AAAI Press: Menlo Park, 256–263.
- Sandholm, Tuomas, Subhash Suri, Andrew Gilpin, and David Levine (2002): Winner Determination in Combinatorial Auction Generalizations, in: *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), Bologna, Italy, 2002*, ACM: New York, 69–76.
- Sheffi, Yossi (2004): Combinatorial Auctions in the Procurement of Transportation Services, *Interfaces*, 34 (4): 245–252.
- Song, Jiongjiong and Amelia Regan (2003): Combinatorial Auctions for Transportation Service Procurement: The Carrier Perspective, *Transportation Research Record*, 1833: 40–46.
- Umetani, Shunji and Mutsunori Yagiura (2007): Relaxation Heuristics for the Set Covering Problem, *Journal of the Operations Research Society of Japan*, 50 (4): 350–375.
- Vincke, Philippe (1992): *Multicriteria Decision-Aid*, Wiley: Chichester.
- Yagiura, Mutsunori, Masahiro Kishida, and Toshihide Ibaraki (2006): A 3-Flip Neighborhood Local Search for the Set Covering Problem, *European Journal of Operational Research*, 172 (2): 472–499.
- Zitzler, Eckart, Marco Laumanns, and Lothar Thiele (2002): SPEA2: Improving the Strength Pareto Evolutionary Algorithm for Multiobjective Optimization, in: *Proceedings of the EUROGEN2001 Conference, Athens, Greece, CIMNE: Barcelona*, 95–100.
- Zitzler, Eckart and Lothar Thiele (1999): Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach, *IEEE Transactions on Evolutionary Computation*, 3 (4): 257–271.



Zitzler, Eckart, Lothar Thiele, Marco Laumanns, Carlos M. Fonseca, and V. G. da Fonseca (2003): Performance Assessment of Multiobjective Optimizers: An Analysis and Review, *IEEE Transactions on Evolutionary Computation*, 7 (2): 117–132.

Biographies

Tobias Buer is working as research and teaching assistant at the Department of Information Systems at the University of Hagen. From the same institution he also received a master's degree in Business Administration and a master's degree in Economics. His research interests are in operational planning problems related to coordination and collaboration in the logistics domain,

currently with a focus on multi-objective combinatorial optimization problems.

Giselher Pankratz is assistant professor at the Department of Information Systems at the University of Hagen. He holds a master's degree in Business Administration from the University of Siegen and a doctoral degree from the University of Hagen. His current research interests include heuristic methods for logistics planning, collaborative transportation planning, combinatorial auctions, and multi-objective optimization. In 2003, he received the Stinnes Logistics Award for his work on methods for integrated transport operations planning in freight forwarding companies.