An Exact Algorithm for Higher-Dimensional Orthogonal Packing*

Sándor P. Fekete
Department of Mathematical Optimization
Braunschweig University of Technology
D–38106 Braunschweig
GERMANY
s.fekete@tu-bs.de

Jörg Schepers†
IBM Germany
Gustav-Heinemann-Ufer 120/122
D–50968 Köln
GERMANY
schepers@de.ibm.com

Jan C. van der Veen
Department of Mathematical Optimization
Braunschweig University of Technology
D–38106 Braunschweig
GERMANY
j.van-der-veen@tu-bs.de

Abstract

Higher-dimensional orthogonal packing problems have a wide range of practical applications, including packing, cutting, and scheduling. Combining the use of our data structure for characterizing feasible packings with our new classes of lower bounds, and other heuristics, we develop a two-level tree search algorithm for solving higher-dimensional packing problems to optimality. Computational results are reported, including optimal solutions for all two–dimensional test problems from recent literature.

This is the third in a series of articles describing new approaches to higher-dimensional packing.

1 Introduction

Most combinatorial optimization problems that need to be solved in practical applications turn out to be members of the class of NP-hard problems. Algorithmic research of several decades has provided strong evidence that for all of these problems, it is highly unlikely that there is a *polynomial algorithm*: Such an algorithm is guaranteed to find an optimal solution in time that even in the worst case can be bounded by a polynomial in the size of the input. If no such bound can be guaranteed, the necessary time for solving instances tends to grow very fast as the instance size increases. That is why NP-hard problems have also been dubbed "intractable". See the classical monograph [20] for an overview.

When confronted with an NP-hard problem, there are several ways to deal with its computational difficulty:

We can look for a different problem.

While this way out may be quite reasonable in a theoretical context, it tends to work less well when a problem arises in practical applications that have to be solved somehow.

We can look for special properties of a problem instance or relax unimportant constraints in order to get a polynomial algorithm.

Unfortunately, practical instances and their additional constraints tend to be more difficult at a second glance, rather than simpler.

We can look for a good solution instead of an optimal one.

*A previous extended abstract version of this paper appears in *Algorithms – ESA'97* [11]

†Supported by the German Federal Ministry of Education, Science, Research and Technology (BMBF, Förderkennzeichen 01 IR 411 C7).

This approach has received an increasing amount of attention over recent years. In particular, there has been a tremendous amount of research dealing with polynomial time approximation algorithms that are guaranteed to find a solution within a fixed multiplicative constant of the optimum. See the book [26] for an overview.

We can look for an optimal solution without a bound on the runtime.

While the time for finding an optimal solution may be quite long in the worst case, a good understanding of the underlying mathematical structure may allow it to find an optimal solution (and prove it) in reasonable time for a large number of instances. A good example of this type can be found in [24], where the exact solution of a 120-city instance of the Traveling Salesman Problem is described. In the meantime, benchmark instances of size up to 13509 and 15112 cities have been solved to optimality [1], showing that the right mathematical tools and sufficient computing power may combine to explore search spaces of tremendous size. In this sense, "intractable" problems may turn out to be quite tractable.

In this paper, we consider a class of problems that is not only NP-hard, but also difficult in several other ways. Packing rectangles into a container arises in many industries, where steel, glass, wood, or textile materials are cut, but it also occurs in less obvious contexts, such as machine scheduling or optimizing the layout of advertisements in newspapers. The three-dimensional problem is important for practical applications as container loading or scheduling with partitionable resources. Other applications arise from allocating jobs to reconfigurable computer chips—see [36]. For many of these problems, objects must be positioned with a fixed orientation; this is a requirement that we will assume throughout the paper. The *d-dimensional orthogonal knapsack problem* (OKP-d) requires to select a most valuable subset S from a given set of boxes, such that S can be packed into the container. Being a generalization of the one-dimensional bin packing problem, the OKP-d is NP-complete in the strict sense. Other NP-hard types of packing problems include the *strip packing problem* (SPP), where we need to minimize the height of a container of given width, such that a given set of boxes can be packed, and the *orthogonal bin packing problem* (OBPP) where we have a supply of containers of a given size and need to minimize the number of containers that are needed for packing a set of boxes. The decision version of these problems is called the *Orthogonal Packing Problem* (OPP), where we have to decide whether a given set of boxes fits into a container.

Relatively few authors have dealt with the exact solution of orthogonal knapsack problems. All of them focus on the problem in two dimensions. One of the reasons is the difficulty of giving a simple mathematical description of the set of feasible packings: As soon as one box is packed into the container, the remaining feasible space is no longer convex, excluding the straightforward application of integer programming methods. Biró and Boros (1984) [5] give a characterization of non–guillotine patterns using network flows but derived no algorithm. Dowsland (1987) [9] proposes an exact algorithm for the case that all boxes have equal size. Arenales and Morabito (1995) [2] extend an approach for the guillotine problem to cover a certain type of non–guillotine patterns. So far, only three exact algorithms have been proposed and tested for the general case. Beasley (1985) [3] and Hadjiconstantinou and Christofides (1995) [25] give different 0–1 integer programming formulations of this problem. Even for small problem instances, they have to consider very large 0–1 programs, because the number of variables depends on the size of the container that is to be packed. The largest instance that is solved in either article has 9 out of 22 boxes packed into a 30×30 container. After an initial reduction phase, Beasley gets a 0-1 program with more than 8000 variables and more than 800 constraints; the program by Hadjiconstantinou and Christofides still contains more than 1400 0–1 variables and over 5000 constraints. ¿From Lagrangean relaxations, they derive upper bounds for a branch–and–bound algorithm, which are improved using subgradient optimization. The process of traversing the search tree corresponds to the iterative generation of an optimal packing. More recent work by Caprara and Monaci (2004) [6] on the two-dimensional knapsack problem uses our previous results (cited as [11, 17] as the most relevant reference for comparison; we compare our results and approaches later in this paper and discuss how a combination of our methods may lead to even better results.

Other research on the related problem of two- and three-dimensional bin packing has been presented: Martello and Vigo (1996) [30] consider the two-dimensional case, while Martello, Pisinger, and Vigo (1997) [28] deal with three-dimensional bin packing. We discuss aspects of those papers in [14], when considering bounds for higher-dimensional packing problems. Padberg (2000) [33] gives a mixed integer programming formulation for three-dimensional packing problems, similar to the one anticipated by the second author in his thesis [35]. Padberg expresses the hope that using a number of techniques from branch-and-cut will be

useful; however, he does not provide any practical results to support this hope.

In our papers [13, 14], we describe a different approach to characterizing feasible packings and constructing optimal solutions. We use a graph-theoretic characterization of the relative position of the boxes in a feasible packing [13]. Combined with good heuristics for dismissing infeasible subsets of boxes that are described in [14], this characterization can be used to develop a two-level tree search. In this third paper of the series, we describe how this exact algorithm can be implemented. Our computational results show that our code outperforms previous methods by a wide margin. It should be noted that our approach has been used and extended in the practical context of reconfigurable computing [36], which can be interpreted as packing in three-dimensional space, with two coordinates describing chip area and one coordinate describing time. Order constraints for the temporal order are of vital importance in this context; as it turns out, our characterization of feasible packings is particularly suited for taking these into account. See our followup paper [10] for a description of how to deal with higher-dimensional packing with order constraints.

The rest of this paper is organized as follows: after recalling some basics from our papers [13, 14, 15, 16] in Section 2, we give detailed account of our approach for handling OPP instances in Section 3. This analysis includes a description of how to apply graph theoretic characterizations of interval graphs to searching for optimal packings. Section 4 provides details of our branch-and-bound framework and the most important subroutines. In Section 5, we discuss our computational results. Section 6 gives a brief description of how our approach can be applied to other types of packing problems.

2 Preliminaries

We are given a finite set V of d-dimensional rectangular boxes with "sizes" $w(v) \in \mathbb{R}_0^{+d}$ and "values" $c(v) \in \mathbb{R}_0^+$ for $v \in V$. As we are considering fixed orientations, $w_i(v)$ describes the size of box v in the x_i-direction. The objective of the *d-dimensional orthogonal knapsack problem (OKP-d)* is to maximize the total value of a subset $V' \subseteq V$ fitting into the container C and to find a complying packing. Closely related is the *d-dimensional orthogonal packing problem (OPP-d)*, which is to decide whether a given set of boxes B fits into a unit size container, and to find a complying packing whenever possible.

For a d-dimensional packing, we consider the projections of the boxes onto the d coordinate axes x_i. Each of these projections induces a graph G_i: two boxes are adjacent in G_i, if and only if their x_i projections overlap. A set of boxes $S \subseteq V$ is called *x_i–feasible*, if the boxes in S can be lined up along the x_i–axis without exceeding the x_i-width of the container.

As we show in [13, 15], we have the following characterization of feasible packings:

Theorem 1 *A set of boxes V can be packed into a container, if and only if there is a set of d graphs $G_i = (V, E_i), i = 1, \ldots, d$ with the following properties:*

$$
\begin{aligned}
&P1: && \text{the graphs } G_i := (V, E_i) \text{ are interval graphs.}\\
&P2: && \text{each stable set } S \text{ of } G_i \text{ is } x_i\text{–feasible.}\\
&P3: && \bigcap_{i=1}^{d} E_i = \emptyset.
\end{aligned}
$$

A set $E = (E_1, \ldots, E_d)$ of edges is called a *packing class* for (V, w), if and only if it satisfies the conditions $P1$, $P2$, $P3$. Note that when constructing a packing from a packing class, some edges may be added in case of a degenerate packing; see Ferreira and Oliveira [19] for such an example. This does not impede the correctness of the theorem or its applicability.

3 Solving Orthogonal Packing Problems

For showing feasibility of any solution to a packing problem, we have to prove that a particular set of boxes fits into the container. This subproblem is called the *orthogonal packing problem (OPP)*.

In order to get a fast positive answer, we can try to find a packing by means of a heuristic. A fast way to get a negative answer has been described in our paper [14]: Using a selection of bounds (conservative scales), we can try to apply the volume criterion to show that there cannot be a feasible packing.

In this section, we discuss the case that both of these easy approaches fail. Because the OPP is NP-hard in the strong sense, it is reasonable to use enumerative methods. As we showed in our paper [13], the existence of a packing is equivalent to the existence of a packing class. Furthermore, we have shown that a feasible packing can be constructed from a packing class in time that is linear in the number of edges. This allows us to search for a packing class, instead of a packing. As we will see in the following, the advantage of this approach lies not only in exploiting the symmetries discussed in [13, 15], but also in the fact that the structural properties of packing classes give rise to very efficient rules for identifying irrelevant portions of the search tree.

3.1 Basic Idea of the Enumeration Scheme

The enumeration of packings described by Beasley in [3] emulates the intuitive idea of packing objects into a box: Each branching corresponds to placing a box at a particular position in the container, or disallowing this placement. Thus, each search node corresponds to a partial packing that is to be augmented to a complete packing. Our enumeration of packing classes is more abstract than that: At each branching we decide whether two boxes b and c overlap in their projection onto the i-axis, so that the edge $e := bc$ is contained in the ith component graph of the desired packing class E. Accordingly, in the first resulting subtree, we only search for packing classes E with $e \in E_i$; in the second, we only search for E with $e \notin E_i$. Hence, the resulting "incomplete packing classes" do *not* correspond to packing classes of subsets of boxes, instead they are (almost) arbitrary tuples of edges.

More precisely, we will store the "necessary" and "excluded " edges for each node N of the search tree and each coordinate direction i in two data structures $\mathcal{E}^N_{+,i}$ and $\mathcal{E}^N_{-,i}$. Therefore, the search space for N contains precisely the packing classes that satisfy the condition

$$\mathcal{E}^N_{+,i} \subseteq E_i \subseteq \overline{\mathcal{E}^N_{-,i}},\ i \in \{1, \ldots, d\}, \tag{1}$$

where $\overline{\mathcal{E}^N_{-,i}}$ is the complement of $\mathcal{E}^N_{-,i}$. Summarizing, we write

$$\mathcal{E}^N_+ := (\mathcal{E}^N_{+,1}, \ldots, \mathcal{E}^N_{+,d}), \quad \mathcal{E}^N_- := (\mathcal{E}^N_{-,1}, \ldots, \mathcal{E}^N_{-,d}), \quad \mathcal{E}^N := (\mathcal{E}^N_+, \mathcal{E}^N_-)$$

and denote by $\mathcal{L}(\mathcal{E}^N)$ the search space for N; by virtue of (1), this search space is only determined by $\mathcal{E}^N$. $\mathcal{E}^N$ is called the *search information* of node N, because this tuple of data structures represents the information that is currently known about the desired packing class.

An important part of the procedure consists in using the characteristic properties $P1$, $P2$, $P3$ for increasing the information on the desired packing class that is contained in $\mathcal{E}^N$. For example, let the edge e be contained in $\mathcal{E}_{+,i}$ for all $i \neq k$. Furthermore, let $E \in \mathcal{L}(\mathcal{E}^N)$. For $i \neq k$, we have $e \in E_i$ because of $\mathcal{E}^N_{+,i} \subseteq E_i$. Because of $P3$, the intersection of all E_i must be empty, implying $e \notin E_k$. Therefore, $\mathcal{E}^N_{-,k}$ can be augmented by e without changing the search space. Similar augmentation rules can be described for $P1$ and $P2$.

Depending on whether the edge e is added to $\mathcal{E}^N_{+,i}$ or to $\mathcal{E}^N_{-,i}$, we describe augmentations of $\mathcal{E}^N$ by the triples $(e, +, i)$ or $(e, -, i)$.

Because it suffices to find a single packing class, these augmentations may reduce the search space, as long as it is guaranteed that not all packing classes are removed from it. This fact allows us to exploit certain symmetries. Thus, we use *feasible augmentations* of $\mathcal{E}^N$ in the sense that a nonempty search space $\mathcal{L}(\mathcal{E}^N)$ stays nonempty after the augmentation.

When augmenting $\mathcal{E}^N$ we follow two objectives:

1. obtain a packing class in $\mathcal{E}^N_+$,
2. prove that every augmentation of $\mathcal{E}^N_+$ to a packing class has to use "excluded " edges from $\mathcal{E}^N_-$.

In the first case, our tree search has been successful. In the second case, the search on the current subtree may be terminated, because the search space is empty. Otherwise, we have to continue branching until one of the two objectives is reached.

3.2 Excluded Induced Subgraphs

For our algorithm, we need three components: a test "Is $\mathcal{E}_+^N$ a packing class?", a sufficient criterion that $\mathcal{E}_+^N$ has no feasible augmentation, and a construction method for feasible augmentations. As we describe in our paper [13, 15], all three of these components can be reduced to identifying or avoiding particular induced subgraphs in the portions of E that are fixed by $\mathcal{E}^N$.

As we have already seen, it is easy to determine all edges that are excluded by condition $P3$. By performing these augmentations of $\mathcal{E}_-^N$ immediately, we can guarantee that $P3$ is satisfied. Thus we will assume in the following that $P3$ is satisfied. Furthermore we will implicitly refer to the current search node N and abbreviate $\mathcal{E}^N$ by $\mathcal{E}$.

$P2$ explicitly excludes certain induced subgraphs: i-infeasible stable sets, i. e., i-infeasible cliques in the complement of each component graph.

In order to formulate $P1$ in terms of excluded induced subgraphs, we recall the following Theorems 3 and 4 – see the book by Golumbic[23], as well as a resulting linear-time algorithm by Korte and Möhring [27]. The following terminology is used:

Definition 2 *For a graph $G := (V, E)$, a set $F \subseteq V^2$ of directed edges is an* orientation of G, *iff*

$$\forall b, c \in V : bc \in E \iff (\vec{bc} \in F \wedge \vec{cb} \notin F) \vee (\vec{cb} \in F \wedge \vec{bc} \notin F)$$

holds. An orientation F of a graph (V, E) is called transitive, *if in addition,*

$$\forall b, c, z \in V : \vec{bc} \in F \wedge \vec{cz} \in F \Rightarrow \vec{bz} \in F$$

holds.

A graph is called a comparability graph, *iff it has a transitive orientation.*

For a cycle $C := [b_0, \ldots, b_{k-1}, b_k = b_0]$ of length k, the edges $b_ib_j, i, j \in \{0, \ldots, k-1\}$ with $(|i-j| \bmod k) > 1$ are called chords*; the chords $b_ib_j, i, j \in \{0, \ldots, k-1\}$ with $(|i-j| \bmod k) = 2$ are called* 2-chords *of C. A cycle is (*2*-)* chordless*, iff it does not have any (*2*-) chords.*

A graph $G = (V, E)$ is a cocomparability graph, *if the complement graph $G = (V, \overline{E})$ is a comparability graph.*

Theorem 3 (Gilmore and Hoffman 1964) *A cocomparability graph is an interval graph, iff it does not contain the chordless cycle C_4 of length 4 as an induced subgraph.*

Theorem 4 (Ghouilà-Houri 1962, Gilmore and Hoffman 1964) *A graph is a comparability graph, iff it does not contain a 2-chordless cycle of odd length.*

Thus, $\mathcal{E}_+$ is a packing class, if for all $i \in \{1, \ldots, d\}$ the following holds (recall that $P3$ is assumed to be satisfied):

1. $(V, \mathcal{E}_{+,i})$ does not contain a C_4 as an induced subgraph.
2. $(V, \overline{\mathcal{E}_{+,i}})$ does not contain an odd 2-chordless cycle.
3. $(V, \overline{\mathcal{E}_{+,i}})$ does not contain an i-infeasible clique.

With the help of this characterization, we get a stop criterion for subtrees. Because only those edges can be added to $\mathcal{E}_{+,i}$ that are not in $\mathcal{E}_{-,i}$, $\mathcal{E}_+$ cannot be augmented to a packing class, if for $i \in \{1, \ldots, d\}$ one of the following conditions holds:

1. $(V, \mathcal{E}_{+,i})$ contains a C_4 as an induced subgraph, with both chords lying in $\mathcal{E}_{-,i}$.
2. $(V, \mathcal{E}_{-,i})$ contains an odd 2-chordless cycle, with all its 2-chords lying in $\mathcal{E}_{+,i}$.
3. $(V, \mathcal{E}_{-,i})$ contains a i-infeasible clique.

Suppose that except for one edge e, one of these excluded configurations is contained in $\mathcal{E}$. Because of condition 1., the corresponding incomplete induced subgraph is contained in the ith component graph of each packing class $E \in \mathcal{L}(\mathcal{E})$. Because completing the excluded subgraph would contradict condition $P1$ or $P2$, the membership of e in E_i is determined. The resulting forced edges can be added to $\mathcal{E}_{+,i}$ or to $\mathcal{E}_{-,i}$ without decreasing the search space.

Example: If $(V, \mathcal{E}_{+,i})$ contains an induced C_4, for which one chord is contained in $\mathcal{E}_{-,i}$, for any packing class of the search space, the other chord e must be contained in the ith component graph. Thus the augmentation $(e, +, i)$ is feasible, but not the augmentation $(e, -, i)$.

In this way, we can reduce the search for a feasible augmentation to the search for incomplete excluded configurations. In the next section, we will relax the condition that only one edge is missing from a configuration, and only require that the missing edges are equivalent in a particular sense.

3.3 Isomorphic Packing Classes

When exchanging the position of two boxes with identical sizes in a feasible packing, we obtain another feasible packing. Similarly, we can permute equal boxes in a packing class:

Theorem 5 *Let E be a packing class for (V, w) and $\pi : V \to V$ be a permutation with*

$$\forall b \in V : \; w(b) = w(\pi(b)). \tag{2}$$

Then the d-tuple of edges in E^π that is given by

$$\forall b, c \in V \, \forall i \in \{1, \ldots, d\} \quad bc \in E_i \; \Leftrightarrow \; \pi(b)\pi(c) \in E_i^\pi,$$

is a packing class.

Proof: Because the structure of the component graphs does not change, conditions $P1$ and $P3$ remain valid. Because of (2), the weight of stable sets remains unchanged, so that $P2$ remains valid as well. □

We get a notion of isomorphism that is similar to the isomorphism of graphs:

Definition 6 *Two packing classes E and E' are called isomorphic, iff there is a permutation $\pi : V \to V$ satisfying (2), such that $E' = E^\pi$.*

Keeping only one packing class from each isomorphism class in search space avoids unnecessary work. For this purpose, we may assume that the ordering of equal boxes in a packing class follows the lexicographic order of their position vectors. As a result, in the two-dimensional case, the leftmost and bottommost box of a box type will have the lowest index. This corresponds to generating packings according to "leftmost downward placement" in [25]. This approach cannot be used for packing classes, because there are no longer any orientations (left/right, up/down, etc.)

Until now, no algorithm has been found that can decide in polynomial time whether two graphs are isomorphic, and it has been conjectured that no such algorithm exists (see [34], p. 291) When deciding whether two packing classes are isomorphic, this decision has to be made repeatedly. In addition, packing classes may only be known partially. This makes it unlikely that there is an efficient method for achieving optimal reduction of isomorphism. Therefore we are content with exploiting certain cases that occur frequently.

In the initializing phase, we may conclude by Theorem 16 from our paper [14] (corresponding to Theorem 11 in [16]) that for a box type T, there is a component graph $(V, E_i)[T]$, for which there is a clique of size $k \geq 2$. Then we can choose the numbering of T, such that the first k boxes from T belong to the clique. Thus, the corresponding $\binom{k}{2}$ edges can be fixed in $\mathcal{E}_{+,i}$. This restriction of numbering T corresponds to excluding isomorphic packing classes.

In the following, we only consider isomorphic packing classes for which the permutation in Definition 6 exchanges precisely two boxes, while leaving all other boxes unchanged. This restricted isomorphism can be checked easily. We have to search for pairs of boxes that can be exchanged in the following way:

Definition 7 *Let (V, w) be an OPP instance with search information $\mathcal{E}$. Two boxes $b, c \in V$ with $w(b) = w(c)$ are called* indistinguishable (with respect to $\mathcal{E}$)*, iff all adjacencies of b and c have identical search information, i. e.,*

$$\forall i \in \{1, \ldots, d\} \quad \forall \sigma \in \{+, -\} \quad \forall z \in V \setminus \{b, c\} : \quad bz \in \mathcal{E}_{\sigma,i} \Leftrightarrow cz \in \mathcal{E}_{\sigma,i}. \tag{3}$$

Two edges $e, e' \in E_V$ are called indistinguishable (with respect to $\mathcal{E}$)*, if there are representations $e = bc$ and $e' = b'c'$, such that the boxes b and b', as well as c and c' are indistinguishable (with respect to $\mathcal{E}$).*

The property of being indistinguishable is an equivalence relation for boxes as well as for edges.

The following lemma allows it to exploit the connection between indistinguishable edges and isomorphic packing classes:

Lemma 8 *Let (V, w) be an OPP instance with search information $\mathcal{E}$. Let A be set of of indistinguishable edges on V. Let e be an arbitrary $e \in A$. Then for any packing class $E \in \mathcal{L}(\mathcal{E})$ that satisfies $A \cap E_i \neq \emptyset$, there is an isomorphic packing class $E' \in \mathcal{L}(\mathcal{E})$ with $e \in E'_i$.*

Proof: Let e' be an edge from the set $A \cap E_i$. Then e and e' are indistinguishable. Hence there is a representation $e = bc$ and $e' = b'c'$, such that b, b' and c, c' are pairs of indistinguishable boxes. Let π be the permutation of V that swaps b and b', and c and c', and let $E' := E^\pi$. Applying (3) twice, it follows from $E \in \mathcal{L}(\mathcal{E})$ that $E' \in \mathcal{L}(\mathcal{E})$. □

Lemma 8 can be useful in two situations:

1. If we branch with $e \in A$ with respect to the i-direction, then we may assume for all $e' \in A$ in the subtree "$e \notin E_i$" that $e' \notin E_i$: For each packing class excluded in this way, there is an isomorphic packing class that is contained in the search space of the subtree $e \in E_i$.
2. If during the course of our computations, we get $A \cap \mathcal{E}_{+,i} \neq \emptyset$, then for an arbitrary $e \in A$, the augmentation $(e, +, i)$ is feasible, because only isomorphic duplicates are lost.

3.4 Pruning by Conservative Scales

By Lemma 20 from our paper [14] (Lemma 15 in [16]), we can use the information given by $\mathcal{E}$ to modify a given conservative scale, such that the resulting total volume of V is increased. Before branching, we try to apply the lemma repeatedly, such that the transformed volume exceeds the volume of the container. In this case, the search can be stopped.

Because this reduction heuristic requires the computation of several one-dimensional knapsack problems, it only pays off to use it on nodes where it may be possible to cut off large subtrees. Therefore, we have only used it on nodes of depth at most 5.

4 Detailed Description of the OPP Algorithm

In this section, we give a detailed description of our implementation of the OPP algorithm. We will omit the description of standard techniques like efficient storage of sets, lists, graphs, or the implementation of graph algorithms. The interested reader can find these in [31] and [23].

4.1 Controlling the Tree Search

The nodes of the search tree are maintained in a list $\mathcal{N}$. For each node $N \in \mathcal{N}$, there is the search information $\mathcal{E}^N$ (see above) and a triple $(e, \sigma, i)^N$ with $e \in \binom{V}{2}$, $\sigma \in \{+, -\}$ and $i \in \{1, \ldots, d\}$. This triple represents the new information when branching at N, i. e., $e \in E_i$ for $\sigma = +$, or $e \notin E_i$ for $\sigma = -$.

Figure 1 shows the course of the tree search. Lines 1. through 3. initialize $\mathcal{N}$ with the root node N_0. Initially, the components of $\mathcal{E}^{N_0}$ do not contain any edges. $(e, \sigma, i)^{N_0}$ is assigned a special value of "NULL".

In the while loop of lines 5. through 28., individual nodes are processed; if necessary, their children are added to $\mathcal{N}$. The particular branching strategy (breadth first or depth first) can be specified by a selection mechanism in line 7. If line 30. is reached, then the whole search tree was checked without finding a packing pattern.

Each node N of the search tree is processed as follows:

In routine **Update_searchinfo**, the augmentation of $\mathcal{E}^N$ described by (e, σ, i) is carried out, as long as there are feasible augmentations. If it is detected that the search on N can be stopped, **Update_searchinfo** outputs the value "EXIT". Otherwise, the routine terminates with "OK".

Call: **Solve_OPP**(P)

Input: An OPP-n instance $P := (V, w)$.

Output: A packing class for (V, w), if there is one, and SUCCESS, otherwise NULL.

1. $\mathcal{N} := \{N_0\}$.
2. initialize $\mathcal{E}^{N_0}$.
3. $(e, \sigma, i)^{N_0}$:= NULL.
4.
5. **while** ($\mathcal{N} \neq \emptyset$) **do**
6.
7. choose $N \in \mathcal{N}$.
8. $\mathcal{N} := \mathcal{N} \setminus \{N\}$.
9. $(e, \sigma, i) := (e, \sigma, i)^N$.
10.
11. **repeat**
12. **if** (**Update_searchinfo** (P, (e, σ, i), $\mathcal{E}^N$) = EXIT) **then**
13. result := EXIT.
14. **else**
15. result := **Packingclass_test** (P, $\mathcal{E}^N$, (e, σ, i)).
16. **end if**
17. **until** (result $\neq$ FIX)
18.
19. **if** (result = SUCCESS) **then** **return** $\mathcal{E}^N_+$.
20.
21. **if** (result = BRANCH) **then**
22. Create two new nodes N', N''.
23. $\mathcal{E}^{N'} := \mathcal{E}^N, \quad (e, \sigma, i)^{N'} := (e, +, i)$.
24. $\mathcal{E}^{N''} := \mathcal{E}^N, \quad (e, \sigma, i)^{N''} := (e, -, i)$.
25. $\mathcal{N} := \mathcal{N} \cup \{N', N''\}$.
26. **end if**
27.
28. **end while**
29.
30. **return** NULL.

Figure 1: OPP tree search

If **Update_searchinfo** was terminated with "OK", then the routine **Packingclass_test** checks whether $\mathcal{E}_+^N$ already is a packing class. In case of a positive answer, "SUCCESS" is output, and the algorithm terminates in line 19. Otherwise, there are three possibilities:

1. "FIX": The triple (e, σ, i) was updated in **Packingclass_test** to a new feasible augmentation that was returned to **Update_searchinfo**.
2. "EXIT": $\mathcal{E}_+^N$ cannot be augmented to a packing class without using edges from $\mathcal{E}_-^N$. The search on this subtree is stopped.
3. "BRANCH": In lines 22. through 25., two children of N are added to $\mathcal{N}$. The triple (e, σ, i) that was set in **Packingclass_test** contains the branching edge e and the branching direction i.

4.2 Testing for Packing Classes

Figure 2 shows routine **Packingclass_test**. As we have seen, $\mathcal{E}_+$ is a packing class, iff no excluded configuration occurs in any coordinate direction. In this case, in each iteration of the i loop, we keep $A = \emptyset$, and the routine terminates in line 33. with value $SUCCESS$. Otherwise, A contains a set of edges, out of which at least one has to be added to $\mathcal{E}_{+,i}$ in order to remove the excluded configuration. This edge must not be from $\mathcal{E}_{-,i}$. Thus, the search on the subtree can be stopped for $A \setminus \mathcal{E}_{-,i} = \emptyset$, and for $|A \setminus \mathcal{E}_{-,i}| = 1$, the only edge must be added to $\mathcal{E}_{+,i}$.

Otherwise, an arbitrary edge from $A \setminus \mathcal{E}_{-,i}$, together with a coordinate direction i is returned in the triple $(e, \sigma, i)^{out}$ and used for branching.

In line 4. it is tested with the help of the decomposition algorithm from [23] p. 129f. whether we have a comparability graph. The runtime is $O(\delta|E|)$, where δ is the maximal degree of a vertex, and E is the edge set of the examined graph. It is simple to modify the algorithm, such that it returns a 2-chordless cycle in case of a negative result.

With the help of the algorithm from [23] p. 133f., we can determine a maximal weighted clique in a comparability graph in time that is linear in the number of edges. This algorithm is called in line 7., as graphs at this stage have already passed the test for comparability graphs.

The search for a C_4 in line 10. can be realized by two nested loops that enumerate possible pairs of opposite edges in a potential C_4.

4.3 Updating the Search Information

Figure 3 gives an overview over Routine **Update_searchinfo**. In the following, we will always refer to the current search node N and denote the search information $\mathcal{E}^N$ by $\mathcal{E}$.

The input triple $(e, \sigma, i)^{in}$ either describes an augmentation of the search information (e is fixed in $\mathcal{E}_{\sigma,i}$), or it contains the value "NULL" on the root node.

On the root node, the search information is initialized as follows: First the edges are fixed for which the vertices form an infeasible stable set with two elements. For $i \in \{1, \ldots, d\}$, this means that all edges $e = bc$ with $w_i(b) + w_i(c) > 1$ are added to $\mathcal{E}_{+,i}$. Furthermore, we use Theorem 16 from our paper [14] (corresponding to Theorem 11 in [16]) in order to fix cliques within the subgraphs induced by the individual box types.

The augmentation $(e, \sigma, i)^{in}$ has either been fixed in the last branching step, or it was returned by the routine **Packingclass_test** together with the value "FIX". In the latter case, $\sigma = +$ holds, so we know in case of $\sigma = -$ that the augmentation results from a branching step. In Section 3.3 we concluded from Lemma 8 that in this case, all edges in $\mathcal{E}_{-,i}$ that are indistinguishable from e can be fixed. This is done in lines 10. through 13.

In the main loop (lines 17. through 26.), for each augmentation of $\mathcal{E}$ it is checked whether it arises from a configuration that allows it to fix further edges, or the search is stopped. This recursive process is controlled by the list L.

The crucial work is done by the subroutines **Check_P3**, **Avoid_C4**, and **Avoid_cliques**.

Call: **Packingclass_test**(P, $\mathcal{E}$, $(e, \sigma, i)^{out}$)

Input: An OPP-n instance $P := (V, w)$, search information $\mathcal{E}$.

Output: $(e, \sigma, i)^{out}$, and
EXIT, FIX, BRANCH, or SUCCESS

```
 1.  for  i ∈ {1, ..., d} do
 2.
 3.      A := ∅.
 4.      if  (V, complement(E_{+,i})) is not a comparability graph then
 5.          A := set of edges of the 2-chordless odd cycle.
 6.      else
 7.          if  a maximal weighted clique in (V, complement(E_{+,i})) is i-infeasible then
 8.              A := edge set of this clique.
 9.          else
10.              if  (V, E_{+,i}) contains an induced C_4 then
11.                  A := set of chords of this C_4.
12.              end if
13.          end if
14.      end if
15.
16.
17.      if  (A ≠ ∅) then
18.          if  ( A \ E_{-,i} = ∅ ) then
19.              return EXIT.
20.          else
21.              Choose an edge e from A \ E_{-,i}.
22.              (e, σ, i)^{out} := (e, +, i).
23.              if  ( A \ E_{-,i} = {e} ) then
24.                  return FIX.
25.              else
26.                  return BRANCH.
27.              end if
28.          end if
29.      end if
30.
31.  end for i
32.
33.  return SUCCESS.
```

Figure 2: Routine Packingclass_test

Call: **Update_searchinfo**(P, $(e,\sigma,i)^{in}$, $\mathcal{E}$)

Input: An OPP-n instance $P := (V, w)$, an augmentation $(e,\sigma,i)^{in}$, the search information $\mathcal{E}$.

Output: The updated search information $\mathcal{E}$, and EXIT or OK.

```
 1.  (e,σ,i) := (e,σ,i)^in.
 2.
 3.  if ( (e,σ,i) = NULL) then
 4.      initialize E and mark this augmentation in L
 5.  else
 6.      if ( σ = +) then
 7.          E_{+,i} := E_{+,i} ∪ {e}.
 8.          L       := {(e,σ,i)}.
 9.      else
10.          for  f ∈ E_V cannot be distinguished from e do
11.              E_{-,i} := E_{-,i} ∪ {f}.
12.              L       := {(f,-,i)}.
13.          end for f
14.      end if
15.  end if
16.
17.  while ( L ≠ ∅ ) do
18.
19.      choose (e,σ,i) ∈ L.
20.      L := L \ {(e,σ,i)}.
21.
22.     if (Check_P3    (P (e,σ,i) E L) ≠ OK ) then return EXIT.
23.     if (Avoid_C4        (P (e,σ,i) E L) ≠ OK ) then return EXIT.
24.     if (Avoid_cliques(P (e,σ,i) E L) ≠ OK ) then return EXIT.
25.
26.  end while
27.
28.  return OK.
```

Figure 3: Routine Update_searchinfo

Checking Condition $P3$

In subroutine **Check_P3**, for an augmentation $(e, +, i)$ the set of "free" coordinate directions

$$F := \{j \in \{1, \dots, d\} \mid e \notin \mathcal{E}_{+,j}\}$$

is computed. If this set only has one element k, then for all $E \in \mathcal{L}(\mathcal{E})$ the condition $e \notin E_k$ must hold because of $\cap_{i=1}^{d} E_i = \emptyset$. In this case, e can be fixed in $\mathcal{E}_{-,k}$, and **Check_P3** terminates with the value "OK". If there is no "free" coordinate direction left, then the search space is empty, and the routine terminates with the value "EXIT".

Avoiding Induced C_4s

Routine **Avoid_C4** tries to detect edges that can be used for completing an induced C_4 in $(V, \mathcal{E}_{+,i})$, with chords lying in $\mathcal{E}_{-,i}$. Such an edge f is then added to $\mathcal{E}_{+,i}$ or to $\mathcal{E}_{-,i}$, such that this excluded induced subgraph is avoided.

Because this configuration must have been caused by the augmentation (e, σ, i) that was given to **Avoid_C4**, e must either be an edge of the cycle, or a chord. Because f can occur as an edge of the cycle as well as a chord, we have to check a total of four cases. Figure 4 shows the routine in detail.

Avoiding Infeasible Cliques

The subroutine **Avoid_cliques** checks whether an edge $e = bc$ that has been added to $\mathcal{E}_{-,i}$ completes one of the following configurations:

1. an i-infeasible clique in $(V, \mathcal{E}_{-,i})$,
2. an i-infeasible clique in $(V, \overline{\mathcal{E}_{+,i}})$, with edges not in $\mathcal{E}_{-,i}$ being indistinguishable.

As we have seen in 3.2, in the first case, the search space is empty. The routine terminates with value "EXIT". In the second case, we can find a feasible augmentation by virtue of Lemma 8.

Computing S_0':
We search for a clique in $(V, \mathcal{E}_{-,i})$ that contains $e = bc$ and has large weight. Trivially, the box set of such a clique can only contain b, c, and boxes from

$$S_0 := \{z \in V \mid bz \in \mathcal{E}_{-,i} \wedge cz \in \mathcal{E}_{-,i}\}.$$

Now our approach depends on whether $(V, \mathcal{E}_{-,i})[S_0]$ is a comparability graph. In the positive case, we can use the linear time algorithm from [23] (just like for the test of packing classes) to determine a set of boxes S_0' that induces a maximal weighted clique in $(V, \mathcal{E}_{-,i})[S_0]$. Then $\{b, c\} \cup S_0'$ induces a clique in $(V, \mathcal{E}_{-,i})$ that has maximal weight among all cliques containing e.
If on the other hand, $(V, \mathcal{E}_{-,i})[S_0]$ is not a comparability graph, then we skip the computation of a maximal weighted clique. (As a generalization of the CLIQUE problem (problem [GT19] in [20]), this problem is NP-hard.) Instead, we compute S_0' by using a greedy strategy. Starting with $S_0' = \emptyset$, we keep augmenting S_0' by the box with the largest weight w_i, as long as the property $E_{S_0'} \subseteq \mathcal{E}_{-,i}$ remains valid. The clique induced by $\{b, c\} \cup S_0'$ in $(V, \mathcal{E}_{-,i})$ may be suboptimal.
In both cases, the routine terminates in case of an i-infeasible $S_0' \cup \{b, c\}$ with the value "EXIT".

In order to determine whether $(V, \mathcal{E}_{-,i})[S_0]$ is a comparability graph, we use the decomposition algorithm from [23] in **Packingclass_test**. In the implementation, it is worthwhile taking into account that in the case $|S_0| \leq 4$, the testing for a comparability graph can be omitted. The corresponding induced subgraphs must be comparability graphs, because a 2-chordless cycle must contain at least five different vertices. In our numerical experiments, this turned out to be a common situation.

Finding an augmenting edge by computing B:
We test whether an edge $e' \in \overline{\mathcal{E}_{+,i}} \cap \overline{\mathcal{E}_{-,i}}$ can be fixed. A sufficient condition is the existence of a set $B \subseteq V$ that satisfies the following conditions:

Call: **Avoid_C4**(P, $(e,\sigma,i)^{in}$, $\mathcal{E}$, L)

Input: An OPP-n instance $P := (V, w)$, an augmentation $(e,\sigma,i)^{in}$, the search information $\mathcal{E}$, the augmentation list L.

Output: The updated search information $\mathcal{E}$, the updated augmentation list L, and the value EXIT or OK.

```
1.  (e,σ,i) := (e,σ,i)^in.
2.
3.  if ( σ = + ) then
4.
5.      for  f ∉ E_{−,i} completes a C_4 in E_{+,i}
6.           that contains e and has chords in E_{−,i}.   do
7.        if ( f ∈ E_{+,i} ) then return EXIT.
8.        E_{−,i} := E_{−,i} ∪ {f}.
9.        L    := L    ∪ {(f,−,i)}.
10.     end for f
11.
12.     for  f ∉ E_{+,i} is chord of a C_4 in E_{+,i}
13.          that contains e and has its other chord in E_{−,i}. do
14.       if ( f ∈ E_{−,i} ) then return EXIT.
15.       E_{+,i} := E_{+,i} ∪ {f}.
16.       L    := L    ∪ {(f,+,i)}.
17.     end for f
18.
19. else (σ = −)
20.
21.     for  f ∉ E_{−,i} completes a C_4 in E_{+,i} that has e as a chord
22.          and that has its other chord also in E_{−,i}.   do
23.       if ( f ∈ E_{+,i} ) then return EXIT.
24.       E_{−,i} := E_{−,i} ∪ {f}.
25.       L    := L    ∪ {(f,−,i)}.
26.     end for f
27.
28.     for  f ∉ E_{+,i} is chord of a C_4 in E_{+,i} that has e as its other chord. do
29.       if ( f ∈ E_{−,i} ) then return EXIT.
30.       E_{+,i} := E_{+,i} ∪ {f}.
31.       L    := L    ∪ {(f,+,i)}.
32.     end for f
33.
34. end if
35.
36. return OK.
```

Figure 4: Routine Avoid_C4

Call: **Avoid_cliques**$(P,\ (e,\sigma,i)^{in},\ \mathcal{E},\ L)$

Input: An OPP-n instance (V,w), an augmentation $(e,\sigma,i)^{in}$, the search information $\mathcal{E}$, the augmentation list L.

Output: The updated search information $\mathcal{E}$, the updated augmentation L, and EXIT or OK.

1. $(bc,\sigma,i) := (e,\sigma,i)^{in}$.
2.
3. **if** ($\sigma = +$) **then return** OK.
4.
5. compute S'_0 as described.
6. **if** ($w_i(S'_0 \cup \{b,c\}) > 1$) **then return** EXIT.
7.
8. **if** (b and c are indistinguishable) **then**
9. **if** $\exists b' \in V$ with $bb' \in \overline{\mathcal{E}_{+,i}} \cap \overline{\mathcal{E}_{-,i}}$ **then**
10. compute $B := \{b,c\} \cup S' \cup X$ as described.
11. **if** ($w_i(B) > 1$) **then** $\mathcal{E}_{+,i} := \mathcal{E}_{+,i} \cup \{bb'\}$, $L := L \cup \{(bb',+,i)\}$.
12. **end if**
13. **end if**
14.
15. **for** $b' \in V$ with $bb' \in \overline{\mathcal{E}_{+,i}} \cap \overline{\mathcal{E}_{-,i}}$ and $cb' \in \mathcal{E}_{-,i}$ **do**
16. compute $B := \{b,c\} \cup S' \cup X$ as described.
17. **if** ($w_i(B) > 1$) **then** $\mathcal{E}_{+,i} := \mathcal{E}_{+,i} \cup \{bb'\}$, $L := L \cup \{(bb',+,i)\}$.
18. **end for** b'
19.
20. **for** $b' \in V$ with $bb' \in \mathcal{E}_{-,i}$ and $cb' \in \overline{\mathcal{E}_{+,i}} \cap \overline{\mathcal{E}_{-,i}}$ **do**
21. compute $B := \{b,c\} \cup S' \cup X$ as described.
22. **if** ($w_i(B) > 1$) **then** $\mathcal{E}_{+,i} := \mathcal{E}_{+,i} \cup \{cb'\}$, $L := L \cup \{(cb',+,i)\}$.
23. **end for** b'
24.
25. **for** $b' \in V$ with $bb' \in \mathcal{E}_{-,i}$ and $cb' \in \mathcal{E}_{-,i}$ **do**
26. **for** $c' \in V$ with $bc' \in \mathcal{E}_{-,i}$, $cc' \in \mathcal{E}_{-,i}$ and $b'c' \in \overline{\mathcal{E}_{+,i}} \cap \overline{\mathcal{E}_{-,i}}$ **do**
27. compute $B := \{b,c\} \cup S' \cup X$ as described.
28. **if** ($w_i(B) > 1$) **then** $\mathcal{E}_{+,i} := \mathcal{E}_{+,i} \cup \{b'c'\}$, $L := L \cup \{(b'c',+,i)\}$.
29. **end for** c'
30. **end for** b'
31.
32. **return** OK.

Figure 5: Routine Avoid_cliques

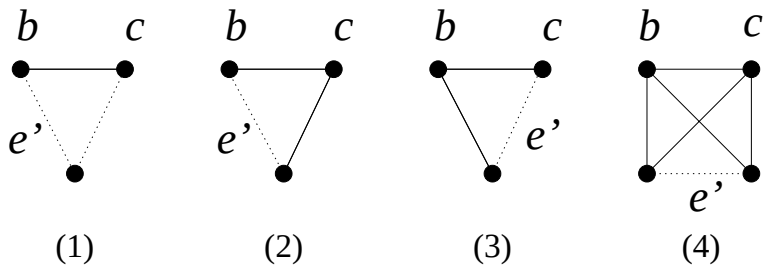

Figure 6: Relative position of e' and $e = bc$.

1. B contains all vertices of e and e'.
2. All edges in $E_B \setminus \mathcal{E}_{-,i}$ are indistinguishable.
3. B is i-infeasible.

Because of $P2$, an edge in E_B must be in the ith component graph of the desired packing class. Because this edge must not be in $\mathcal{E}_{-,i}$, it must be indistinguishable from $e' \in E_B \setminus \mathcal{E}_{-,i}$ by virtue of 2. In other words, Lemma 8 means that augmentation with indistinguishable edges leads to isomorphic packing classes. This implies the feasibility of augmentation $(e', +, i)$.

The requirement that the vertices of e lie in B results from the fact that we only search for incomplete excluded configurations that arise from adding e to $\mathcal{E}_{-,i}$.

When identifying edges that are candidates for e', we get four cases for the position of e' relative to $e = bc$ in E_B, as shown in Figure 6. Dashed lines represent the (unfixed) edges in $\overline{\mathcal{E}_{+,i}} \cap \overline{\mathcal{E}_{-,i}}$, while solid lines represent edges in $\mathcal{E}_{-,i}$. The second requirement for B implies that b and c are indistinguishable. Note that after the first resulting augmentation, b and c are indistinguishable with respect to the current search information. Thus, the cases (1), (2), (3), and (4) in the figure correspond to lines (8.–13.), (15.–18.), (20.–23.), and (25.–30.):

Therefore, constructing the set B for a given edge e' is done as follows. Let S be the set of boxes that are adjacent in $(V, \mathcal{E}_{-,i})$ to all vertices of e and e'. Similar to the above construction of S'_0 from S_0, we construct a set of boxes S' from S that induces a clique in $(V, \mathcal{E}_{-,i})$. The comparability graph test is skipped, if $(V, \mathcal{E}_{-,i})[S_0]$ has been recognized as a comparability graph: If $S \subseteq S_0$, then $(V, \mathcal{E}_{-,i})[S]$ is an induced subgraph and inherits its property of being a comparability graph.

By adding the vertices of e and e' to S', we get a set that satisfies the first two conditions that B needs to satisfy. e' is the only set in the complete graph on this set that does not belong to $\mathcal{E}_{-,i}$. Now we add boxes that provide edges indistinguishable from e'.

Let the set X contain the vertices of e'. The indistinguishable boxes for each vertex form a stable set in $(V, \mathcal{E}_{-,i})$, or they induce a clique in this graph. Only in the latter case do we add these boxes to X. With the help of this construction, any edge in the complete graph on $B := \{b, c\} \cup S' \cup X$ is either in the set $\mathcal{E}_{-,i}$, or it is indistinguishable from e'. If this set is i-infeasible, then we fix e' in $\mathcal{E}_{+,i}$ by virtue of Lemma 8.

5 A Tree Search Algorithm for Orthogonal Knapsack Problems

In this section, we elaborate how the data structure introduced in the paper [13, 15], the lower bounds described in [14, 16], and the exact algorithm for the OPP from Section 3 can be used as building blocks for new exact methods for orthogonal packing problems.

We concentrate on the most difficult problem, the OKP. After a detailed description of the new branch-and-bound approach, the following Section 6 gives evidence that our algorithm allows it to solve considerably larger instances than previous methods. In particular, we present the first results for 3-dimensional instances.

Similar exact algorithms for the SPP and the OBPP are sketched in Section 7.

5.1 The Framework

For solving the OKP, we have to determine a subset $S \subseteq V$ of boxes that has maximum value among all subsets of boxes fitting into the container. Like Beasley [3], and Christofides/Hadjiconstantinou [25], we will prove feasibility of a particular set S by displaying a feasible packing for (V, w, W). For most practical applications, this is of key importance.

In the branch-and-bound algorithms [3] and [25], the iterative choice of a subset and the corresponding packing are treated simultaneously: With each branching step, it is decided whether a particular position in the container is occupied by a particular box type.

In contrast to this, our approach works in two levels. Only after the first level has determined the subset $S \subseteq V$ will the OPP algorithm from Section 3 try to find a feasible packing. This allows us to use the lower bounds described in our paper [14, 16] for excluding most of the first level search tree without having to consider the particular structure of a packing. Our numerical results show that only in a small fraction of search nodes, the second level search has to be used. Note that the main innovation of our approach lies in this second level; it is to be expected that tuning the outer level (as was done by Caprara and Monaci [6]) yields even better results.

5.2 Branch-and-Bound Methods

We assume that the reader is familiar with the general structure of a branch-and-bound algorithm. (A good description can be found in [32].) We start by introducing some notation.

We remind the reader of the partitioning of the set V of boxes into classes of boxes with identical size and value, called *box types*:

$$V = \dot{\bigcup}_{t=1}^{m} T_t.$$

For box type t, we set $n_t := |T_t|$ and denote the elements by

$$T_t =: \{b_{t,1}, \ldots, b_{t,n_t}\}.$$

For ease of notation, we write $w^{(t)}$ instead of $w(b_{t,1})$, and $v^{(t)}$ instead of $v(b_{t,1})$.

In the test instances that we will be dealing with, all sizes of boxes and containers are integers. We denote measures of the container by $W \in \mathbb{N}^d$; when discussing mathematical arguments, we will assume without loss of generality that the container is a unit cube.

5.3 Search Nodes at Level One

The first-level search tree enumerates the subsets $S \subseteq V$ that are candidates for a solution subset for the OKP. Each node N of the search tree corresponds to an OKP instances with the additional constraint that each box type T_t has upper and lower bounds for the number of boxes that are used. These bounds are denoted by $\overline{n}_t^N$ and $\underline{n}_t^N$.

For a search node N, we set

$$\underline{S}^N := \dot{\bigcup}_{t=1}^{m} \{b_{t,1}, \ldots, b_{t,\underline{n}_t^N}\}$$

and similarly

$$\overline{S}^N := \dot{\bigcup}_{t=1}^{m} \{b_{t,1}, \ldots, b_{t,\overline{n}_t^N}\}.$$

For a partial search tree with root N, only subsets S will be considered that satisfy

$$\mathcal{S}(N) := \{S \subseteq V \mid \underline{S}^N \subseteq S \subseteq \overline{S}^N.\}$$

Thus, for a search node N, the corresponding restricted OKP is given by

$$\begin{array}{ll} \textbf{Maximize} & v(S), \\ \textbf{such that} & \text{there is a feasible packing for } (S, w), \\ & \underline{S}^N \subseteq S \subseteq \overline{S}^N. \end{array} \tag{4}$$

On the root node N_0, we start with the original problem, i. e., $\underline{n}_t^{N_0} = 0$ and $\overline{n}_t^{N_0} = n_t$ for $t \in \{1, \ldots, m\}$. Then $\underline{S}^{N_0} = \emptyset$ and $\overline{S}^{N_0} = V$.

Enumerating the first level search tree is done by *best first search*: Each node N is assigned a preliminary local upper bound, given by the minimum of $v(\overline{S}^N)$ and the local upper bound of its parent node. (A better upper bound will be determined while evaluating the partial tree at N.) At each stage, we choose a new node where this local upper bound is maximal.

5.4 Branching

When a subset S has been uniquely determined by the condition $S \in \mathcal{S}(N)$, we have reached a leaf of the first level search tree. In this case, we have

$$\underline{S}^N = S = \overline{S}^N$$

and

$$\forall t \in \{1, \ldots, m\} : \quad \underline{n}_t^N = \overline{n}_t^N .$$

Then the problem (4) is an OPP that is solved by the second level tree search.

Otherwise, we have box types T_t, with $\underline{n}_t^N < \overline{n}_t^N$. We choose the one with largest size $\max_{1 \le i \le d} w_i^{(t)}$ for an arbitrary coordinate direction. By our experience, boxes that are "bulky" in this sense have the biggest influence on the overall solution of the problem.

Now let T_{t^*} be the box type chosen in this way. We branch by splitting $\mathcal{S}(N)$ into subspaces, where the number of boxes in T_{t^*} is constant. For each $\nu \in \{\underline{n}_{t^*}^N, \ldots, \overline{n}_{t^*}^N\}$, we determine a child node N_ν. For this node, we set

$$\underline{n}_t^{N_\nu} := \begin{cases} \nu, & t = t^*, \\ \underline{n}_t^N & t \neq t^*, \end{cases}$$

and

$$\overline{n}_t^{N_\nu} := \begin{cases} \nu, & t = t^*, \\ \overline{n}_t^N & t \neq t^*. \end{cases}$$

A different branching strategy builds a binary search tree, where the two children of N each get one half of $\{\underline{n}_{t^*}^N, \ldots, \overline{n}_{t^*}^N\}$ as a range for the number of boxes in T_{t^*}. For technical reasons, we have used the first variant.

5.5 Lower Bounds

On each node N, the container is filled with boxes from $\overline{S}^N$ by using the following greedy heuristic. The best objective value of the OKP for any of these feasible solutions is stored in v_{lb}. Trivially, v_{lb} is a lower bound for the optimal value of the OKP.

In our heuristic, we build a sequence of packings, where each lower coordinate of a box equals 0 (i. e., the boundary of the container), or the upper coordinate of a preceding box. These positions are called *placement points*. Placement points are maintained in a list that is initialized by the container origin. At each step, a placement point is removed from the list, as we try to use it for placing another box. Following a given ordering, we use the first box type that fits at the chosen placement point without overlapping any of the boxes that are already packed into the container. In case of success, we compute the new placement points and add them to the list. This step is repeated until the list is empty, or all boxes have been packed.

This construction of a packing is repeated for several orderings of box types. In the first round, we use the order of decreasing value. Following rounds use a random weighting of values before sorting; weights are chosen from a uniform distribution on $[0, 1]$.

In our implementation, 50 iterations of this heuristic are performed at the root, and 10 iterations at all other nodes.

5.6 Upper Bounds

The upper bound v_{ub}^N refers to the set of boxes from $\mathcal{S}(N)$. As we showed in the paper [14, 16], for any conservative scale w' for $(\overline{S}^N, w)$, a relaxation of (4) is given by

$$\begin{array}{ll} \textbf{Maximize} & v(S), \\ \textbf{such that} & \sum_{b \in S} \otimes w'(b) \leq 1, \\ & \underline{S}^N \subseteq S \subseteq \overline{S}^N, \end{array} \tag{5}$$

where $\otimes w'(b) := \prod_{i=1}^d w_i'(b)$ denotes the volume of the modified box $w'(b)$. In order to avoid technical difficulties, we only consider conservative scales that are constant for each box type. For the benefit of later generalizations, we formulate the problem (5) explicitly as a restricted one-dimensional knapsack problem:

$$\begin{array}{ll} \textbf{Maximize} & \displaystyle\sum_{t=1}^{m} v(b_{t,1})\xi_t, \\ \textbf{such that} & \displaystyle\sum_{t=1}^{m} \otimes w'(b_{t,1})\xi_t \leq 1, \\ & \underline{n} \leq \xi \leq \overline{n}, \\ & \xi \text{ integer}. \end{array} \tag{6}$$

A problem of this type can be solved by the routine Routine MTB2 from [29], Appendix A.3.1. This transforms the restricted knapsack problem into a 0-1 knapsack problem to which the algorithm of Martello and Toth ([29], pp. 61ff.) is applied.

In our implementation, we use as an upper bound the minimum of the optimal values of the relaxations (5) for the conservative scales

$$w' \in \{(w_1, \ldots, u^{(k)} \circ w_i, \ldots, w_d) \,|\, i = 1, \ldots, d, \quad k = 1, 2, 3, 4\}$$

from our paper [14, 16].

5.7 Removal of Partial Search Trees

We can stop the search on the current search tree, if one of the following conditions is satisfied:

1. $v_{ub}^N \leq v_{lb}$.
2. $\overline{S}^N$ fits into the container.
3. $\underline{S}^N$ does not fit into the container.

The first stop criterion is used in any branch-and-bound procedure. In this case, the currently best solution cannot be improved on the current search tree.

In the second case, $\overline{S}^N \in \mathcal{S}(N)$ is a best feasible solution in $\mathcal{S}(N)$. Because we are always trying to pack all of $\overline{S}^N$ when updating the lower bound v_{lb}, this condition is checked when performing the update.

In the third case, $\underline{S}^N \subseteq S$ implies that no set $S \in \mathcal{S}(N)$ can be packed into the container, so $\mathcal{S}(N)$ cannot contain a feasible solution. This means we have to solve another OPP.

5.8 Solving Orthogonal Packing Problems

In order to solve the OPPs that occur on the leaves of the search tree and when checking the stop criterion "$\underline{S}^N$ does not fit into the container", we use the following strategy:

First we try to use the volume criterion for a selection of conservative scales. Other than the original weight function w, we use the conservative scales

$$w' \in \{(w_1, \ldots, u^{(k)} \circ w_i, \ldots, w_d) \,|\, i = 1, \ldots, d, \quad k = 1, \ldots, W_i/2\}.$$

If this does not produce a (negative) result, we try to find a packing pattern by 10 iterations of our search heuristic. If this fails as well, we use our algorithm from Section 3 to decide the OPP.

5.9 Problem Reduction

There are several ways to decrease the gap between the bounds $\underline{n}_t^N$ and $\overline{n}_t^N$ on a search node. These rules are based upon corresponding reduction tests of Beasley [3]. If the areas of boxes and container are used, we generalize the tests from two to d dimensions. By using conservative scales, we get a generalization of these tests, with markedly increased efficiency.

We start with the rule *Free Value* that remains unchanged. An optimal solution S can have at most value v_{ub}^N. Because $\underline{S} \subseteq S$, further boxes from T_t in S can contribute at most a value of $v_{ub}^N - v(\underline{S})$. Because each of these boxes has value $v^{(t)}$, we set

$$\overline{n}_t := \min \left\{ \overline{n}_t, \underline{n}_t + \left\lfloor \frac{v_{ub}^N - v(\underline{S})}{v^{(t)}} \right\rfloor \right\}. \tag{7}$$

A similar argument, used on volumes, is the basis for Beasley's reduction test *Free Area*. The volume used by $\underline{S}$ is at least as big as the sum of the volumes of the individual boxes in $\underline{S}$. Further boxes from T_t can use at most the volume of the container, reduced by this amount. Because each of these boxes uses a volume of $\otimes w(t)$, we can use the following update for $t \in \{1, \ldots, m\}$:

$$\overline{n}_t := \min \left\{ \overline{n}_t, \underline{n}_t + \left\lfloor \frac{1 - \otimes w(\underline{S})}{\otimes w(t)} \right\rfloor . \right\} \tag{8}$$

With the help of Corollary 8 from our paper [14], we can generalize *Free Area* by replacing w in (8) by an arbitrary conservative scale w' for (V, w) that is constant on T_t. In our implementation, we use w, and the conservative scales

$$w' \in \{(w_1, \ldots, u^{(k)} \circ w_i, \ldots, w_d) \,|\, i = 1, \ldots, d, \quad k = 1, \ldots, W_i/2\}.$$

To allow for further possible improvement of a bound $\underline{n}_{t^*}^N$, $t^* \in \{1, \ldots, m\}$, we expand the relaxation (6) by the additional constraint $\xi_{t^*} = \underline{n}_{t^*}$. The optimal values of the resulting knapsack problems are upper bounds for the value of those solutions $S \in \mathcal{S}(N)$ that contain precisely $\underline{n}_{t^*}$ boxes from T_{t^*}. If the minimum of these bounds does not exceed v_{lb}, a solution $S \in \mathcal{S}(N)$ with a better value than the current best must contain more than $\underline{n}_{t^*}^N$ boxes from T_{t^*}. In this case, we can increment $\underline{n}_{t^*}^N$. This test is repeated for each box type t, until no bound can be improved. If in this process, we get $\underline{n}_t^N > \overline{n}_t^N$, then the search on the partial search tree with root N can be stopped. Thus, we have derived a generalization of the reduction test *Area Program* with the help of conservative scales.

6 Computational Results

The above OKP procedure has been implemented in C++ and was tested on a Sun workstation with Ultra SPARC processors (175 MHz), using the compiler gcc. To allow for a wider range of comparisons with other two-dimensional efforts, we also tested an implementation on a PC with Pentium IV processor (2,8 GHz) with 1 GB memory using g++3.2.

6.1 Results for Benchmark Instances from the Literature

The only benchmark instances for the OKP that have been documented in the literature can be found in the articles by Beasley [3], and Hadjiconstantinou and Christofides [25]. These are restricted to the two-dimensional case. We ran our algorithm on all of these instances that were available. Like Caprara and Monaci [6], we also use a number of other instances that were originally designed for guillotine-cut instances.

The twelve instances *beasley1* through *beasley12* are taken from Beasley's OR library (see [4]). They can be found on the internet at

`http://mscmga.ms.ic.ac.uk/jeb/orlib/ngcutinfo.html`

The data for *hadchr3* and *hadchr11* is given in [25].

problem	container size	box types	# boxes	OKP nodes	OPP calls	OPP nodes	opt. boxes	opt. sol.
beasley1	(10, 10)	5	10	19	1	1	5	164
beasley2	(10, 10)	7	17	5	0	0	5	230
beasley3	(10, 10)	10	21	25	6	36	7	247
beasley4	(15, 10)	5	7	1	0	0	6	268
beasley5	(15, 10)	7	14	1	0	0	6	358
beasley6	(15, 10)	10	15	15	5	5	7	289
beasley7	(20, 20)	5	8	0	0	0	8	430
beasley8	(20, 20)	7	13	53	23	301	8	834
beasley9	(20, 20)	10	18	3	0	0	11	924
beasley10	(30, 30)	5	13	1	0	0	6	1452
beasley11	(30, 30)	7	15	36	10	16	9	1688
beasley12	(30, 30)	10	22	48	14	105	9	1865
hadchr3	(30, 30)	7	7	1	0	0	5	1178
hadchr7	(30, 30)	10	22	48	14	105	9	1865
hadchr8	(40, 40)	10	10	7	0	0	6	2517
hadchr11	(30, 30)	15	15	30	1	1	5	1270
hadchr12	(40, 40)	15	15	5	0	0	7	2949
wang20	(70, 40)	20	42	794	176	1003	8	2726
chrwhi62	(40, 70)	20	62	356	102	7991	10	1860
3	(40, 70)	20	62	356	102	7991	10	1860
3s	(40, 70)	20	62	757	166	3050	8	2726
A1	(50, 60)	20	62	935	254	19283	11	2020
A1s	(50, 60)	20	62	4291	504	8156	7	2956
A2	(60, 60)	20	53	267	70	35747	11	2615
A2s	(60, 60)	20	53	8598	2365	143002	8	3535
CHL2	(62, 55)	10	19	688	317	225011	9	2326
CHL2s	(62, 55)	10	19	1419	557	158450	10	3336
CHL3	(157, 121)	15	35	0	0	0	35	5283
CHL3s	(157, 121)	15	35	0	0	0	35	7402
CHL4	(207, 231)	15	27	0	0	0	27	8998
CHL4s	(207, 231)	15	27	0	0	0	27	13932
CHL5	(30, 20)	10	18	363	194	57115	11	589
cgcut1	(15, 10)	7	16	14	1	1	8	244
cgcut2	(40, 70)	10	23				12	2892
cgcut3	(40, 70)	20	62	356	102	7991	10	1860
gcut01	(250, 250)	10	10	33	0	0	3	48368
gcut02	(250, 250)	20	20	519	51	78	6	59798
gcut03	(250, 250)	30	30	2234	235	742	6	61275
gcut04	(250, 250)	50	50	72159	18316	145057	4	61380
gcut05	(500, 500)	10	10	52	13	13	5	195582
gcut06	(500, 500)	20	20	278	22	22	4	236305
gcut07	(500, 500)	30	30	852	124	152	4	240143
gcut08	(500, 500)	50	50	55485	9037	15970	4	245758
gcut09	(1000,1000)	10	10	12	2	8	5	939600
gcut10	(1000,1000)	20	20	335	31	40	5	937349
gcut11	(1000,1000)	30	30	1616	212	463	6	969709
gcut12	(1000,1000)	50	50	8178	593	1236	5	979521
gcut13	(3000,3000)	32	32					$\geq$8622498 $\leq$9000000
okp1	(100, 100)	15	50	3244	661	35523	11	27718
okp2	(100, 100)	30	30	23626	7310	8721	11	22502
okp3	(100, 100)	30	30	8233	816	921	11	24019
okp4	(100, 100)	33	61	1458	15	50	10	32893
okp5	(100, 100)	29	97	5733	643	13600	8	27923

Table 1: Two-dimensional benchmark instances from previous literature.

problem	time/s	B85	HC95	CM04			
				A_0	A_1	A_2	A_3
beasley1	<0.01	0.9					
beasley2	<0.01	4.0					
beasley3	<0.01	10.5					
beasley4	<0.01	0.1	0.04				
beasley5	<0.01	0.4					
beasley6	<0.01	55.2	45.20				
beasley7	<0.01	0.5	0.04				
beasley8	0.02	218.6					
beasley9	<0.01	18.3	5.20				
beasley10	<0.01	0.9					
beasley11	<0.01	79.1					
beasley12	0.02	229.0	>800				
hadchr3	<0.01		532				
hadchr7	0.01		>800				
hadchr8	<0.01		>800				
hadchr11	<0.01		>800				
hadchr12	<0.01		65.2				
wang20	0.67			6.75	6.31	17.84	2.72
chrwhi62	0.54						
3	0.54						
3s	0.46						
A1	1.12						
A1s	1.51						
A2	1.62						
A2s	8.35						
CHL2	10.36						
CHL2s	6.84						
CHL3	<0.01						
CHL3s	<0.01						
CHL4	<0.01						
CHL4s	<0.01						
CHL5	4.66						
cgcut1	<0.01			0.30	1.47	1.46	1.46
cgcut2	>1800			>1800	>1800	533.45	531.93
cgcut3	0.54			23.76	23.68	4.59	4.58
gcut1	0.01			0.00	0.00	0.01	0.01
gcut2	0.47			0.52	0.19	25.75	0.22
gcut3	4.34			>1800	2.16	276.37	3.24
gcut4	195.62			>1800	346.99	>1800	376.52
gcut5	0.02			0.00	0.50	0.03	0.50
gcut6	0.38			0.06	0.09	9.71	0.12
gcut7	2.24			1.31	0.63	354.50	1.07
gcut8	253.54			1202.09	136.71	>1800	168.50
gcut9	0.01			0.01	0.09	0.05	0.08
gcut10	0.67			0.01	0.13	6.49	0.14
gcut11	8.82			16.72	14.76	>1800	16.30
gcut12	109.81			63.45	16.85	>1800	25.39
gcut13	>1800			>1800	>1800	>1800	>1800
okp1	10.82			24.06	25.46	72.20	35.84
okp2	20.25			>1800	>1800	1535.95	1559.00
okp3	5.98			21.36	1.91	465.57	10.63
okp4	2.87			40.40	2.13	0.85	4.05
okp5	11.78			40.14	>1800	513.06	488.27

Table 2: Runtimes of our implementation, compared to other methods. The columns "B85", "HC95", "CM04" show the runtimes as reported in [3, 25, 6].

Tables 1 and 2 show our results for these OKP-2 instances. For all instances, we found an optimal solution in at most 0.02 seconds. The small number of OKP search nodes (at most 65) as well as OPP search nodes (at most 294) shows the high efficiency of the rules for reducing the search tree. It is also remarkable that on less than a quarter of the search nodes, the enumeration procedure for the OPP had to be used. The majority of reductions resulted from transformed volumes.

Instances *wang20* and *chrwhi62* are considerably larger. They were taken from Wang [37], and Christofides and Whitlock [7] and originally designed for testing the efficiency of guillotine cut algorithms, as were the next three sets: Instances 3 through CHL5 are taken from the benchmark sets by Hifi, which can be found at

`ftp://panoramix.univ-paris1.fr/pub/CERMSEM/hifi/2Dcutting/`.

The sets *cgcut* and *gcut* are also guillotine-cut type instances and can be found at the OR library. Finally, we created five new OKP instances *okp* that are listed in detail in Table 3. They are random instances generated in the same way as *beasley 1–12* after applying initial reduction.

A detailed listing of our results for these OKP-2 instances can be found in Table 1. The first column lists the names of the instances; the second shows the size of the container, followed by the number of different box types and the total number of boxes. The fifth column shows the number of nodes in the outer search tree, followed by the total number of calls to the inner search tree, i.e., the times an OPP had to be resolved. The last two columns show the number of boxes in an optimal solution, and the optimal value. (Instance *gcut13* is still unsolved; our lower bound corresponds to the best solution found so far.) Results are shown in Table 2, where the first column lists the instance names, the second column shows the runtime on a PC with Pentium IV processor. Columns 3 and 4 give the runtimes as reported by Beasley [3] on a CYBER 855, by Hadjiconstantinou and Christofides [25] on a CDC 7600. In [6] Caprara and Monaci give runtimes for four different algorithms. None of these algorithms dominates all the others; the best of them (called A_3) uses a clever hybrid strategy for checking feasibility during branching.

The comparison in Table 2 should be considered with some care, because different computers with different compilers were used for the tests. Some indication for the relative performance of the different machines can be found at

`http://www.netlib.org/benchmark/performance.ps`,

where the results of the Linpack100 benchmark are presented (see [8]). According to these results, a CDC 7600 manages 120 Mflop/s, a CYBER 875 (Cyber 855 does not appear on the list) gets 480 Mflop/s, a Sun Ultra SPARC achieves 7000 Mflop/s, an Intel Pentium III (750 MHz) 13,800 Mflop/s, while an Intel Pentium IV (2,8 GHz) manages 131,700 Mflop/s. Note that these speeds may not be the same for other applications. Furthermore, there is always a certain amount of chance involved when comparing branch-and-bound procedures on individual instances.

Despite of these difficulties in comparison, it is clear that our new method constitutes significant progress. One indication is the fact that the ratio of running times between "large" and "small" instances is smaller by several orders of magnitude: As opposed to our two-level algorithm, the search trees in the procedures by Beasley and Hadjiconstantinou/Christofides appear to be reaching the threshold of exponential growth for some of the bigger instances. After 800 seconds, the procedure by Hadjiconstantinou/Christofides timed out on instances *beasley12*, *hadchr8*, and *hadchr11*, without finding a solution. The comparison with Caprara and Monaci [6] is less conclusive: Both implementations fair pretty well on medium-sized instances, with different behavior for large instances. (Comparing a previous implementation of our algorithm with A_3, Caprara and Monaci concluded that "... the algorithm of [Fekete and Schepers] appears to be more stable ...") This behavior may also be the result of some differences in branching strategies, which can always turn out differently on individual instances. It should be noted that the main basis for the success of our method is the underlying mathematical characterization, and tuning of branching strategies and bounds can be expected to provide further progress. Promising may also be a combination of the first-level strategy of [6] with our second-level strategy.

At this stage, an instance like *cgcut13* is still out of reach, even though we were able to improve the best known solution to 8,622,498, from 8,408,316 in [6], with an upper bound of 9,000,000, leaving a gap of about 4%. It should be interesting to develop long-running, special-purpose exact algorithms, just like Applegate et al. [1] did for the Traveling Salesman Problem.

Table 3: The new problem instances okp1–okp5.

Problem	okp1: container = (100,100), 15 box types (50 boxes)
$size =$	[(4,90),(22,21),(22,80),(1,88),(6,40),(100,9),(46,14),(10,96), (70,27),(57,18),(10,84),(100,1),(2,41),(36,63),(51,24)]
$value =$	[838,521,4735,181,706,2538,1349,1685,5336,1775,1131,129,179, 6668,3551]
$n =$	[5,2,3,5,5,5,3,1,3,1,1,5,5,2,4]
Problem	okp2: container = (100,100), 30 box types (30 boxes)
$size =$	[(8,81),(5,76),(42,19),(6,80),(41,48),(6,86),(58,20),(99,3),(9,52), (100,14),(7,53),(24,54),(23,77),(42,32),(17,30),(11,90),(26,65), (11,84),(100,11),(29,81),(10,64),(25,48),(17,93),(77,31),(3,71), (89,9),(1,6),(12,99),(33,72),(21,26)]
value=	[953,389,1668,676,3580,1416,3166,537,1176,3434,676,1408,2362, 4031,1152,2255,3570,1913,1552,4559,713,1279,3989,4850,299, 1577,12,2116,2932,1214]
$n_j =$	$1, \quad j \in \{1, \ldots, 30\}$
Problem	okp3: container = (100,100), 30 box types (30 boxes)
$size =$	[(3,98),(34,36),(100,6),(49,26),(14,56),(100,3),(10,90),(23,95), (10,97),(50,47),(41,45),(13,12),(19,68),(50,46),(23,70),(28,82), (12,65),(9,86),(21,96),(19,64),(21,75),(45,26),(19,77),(5,84), (16,21),(23,69),(5,89),(22,63),(41,6),(76,30)]
value=	[756,2712,1633,2332,2187,470,1569,4947,2757,4274,4347,396,3866, 5447,2904,6032,1799,929,5186,2120,1629,2059,2583,953,1000, 2900,1102,2234,458,5458]
$n_j =$	$1, \quad j \in \{1, \ldots, 30\}$
Problem	okp4: container = (100,100), 33 box types (61 boxes)
$size =$	[(48,48),(6,85),(100,14),(17,85),(69,20),(12,72),(5,48),(1,97), (66,36),(15,53),(29,80),(19,77),(97,7),(7,57),(63,37),(71,14),(3,76), (34,54),(5,91),(14,87),(62,28),(6,7),(20,71),(92,7),(10,77),(99,4), (14,44),(100,2),(56,40),(86,14),(22,93),(13,99),(7,76)]
$value =$	[5145,874,2924,3182,2862,1224,531,249,6601,1005,6228,3362,907, 473,6137,1556,313,4123,581,1999,5004,2040,3143,795,1460,841, 1107,280,5898,2096,4411,3456,1406]
$n =$	[1,2,1,1,1,1,3,3,2,1,3,1,1,2,2,1,3,1,2,1,3,3,1,1,2,3,2,3,2,1,1,3,3]
Problem	okp5: container = (100,100), 29 box types (97 boxes)
$size =$	[(8,81),(5,76),(42,19),(6,80),(41,48),(6,86),(58,20),(99,3),(9,52), (100,14),(7,53),(24,54),(23,77),(42,32),(17,30),(11,90),(26,65), (11,84),(100,11),(29,81),(10,64),(25,48),(17,93),(77,31),(3,71), (89,9),(1,6),(12,99),(21,26)]
$value =$	[953,389,1668,676,3580,1416,3166,537,1176,3434,676,1408,2362, 4031,1152,2255,3570,1913,1552,4559,713,1279,3989,4850,299, 1577,12,2116,1214]
$n =$	[3,4,4,4,1,5,5,5,5,4,5,1,1,5,5,4,2,3,1,1,2,1,4,1,5,4,5,2,5]

6.2 Generating New Test Instances for 2D and 3D

In order to get a broader test basis, and also include the three-dimensional case, we generated 300 new test instances. We followed the method described in [30] and [28].

Our test instances are characterized by three parameters:

1. type of the instance (I, II, III) (see Tables 4 and 5),
2. number m of box types,
3. number ν of boxes for each box type.

Each of the instances consists of a container of size 100 in each coordinate direction, and m box types, which are obtained as follows:

There are four (OKP-2) or five (OKP-3) classes of box types. The type of the instance determines the probability of each new box type T_t to belong to one of these classes. We use the distributions shown in Table 4 and Table 5.

Depending on its class, the sizes of a box type are generated randomly, according to the distributions in Table 4 and Table 5. We round up to integer values. In order to get the value of a box type, the volume is multiplied with a random number from $\{1, 2, 3\}$. The number of boxes in a new box type it determined by the parameter ν, independent of t.

In this manner, we generated (for two as well as for three dimensions) ten OKP instances for each of the three instance types and each of the five parameter combinations:

$$(m, \nu) \in \{(20, 1), (30, 1), (40, 1), (20, 3), (20, 4).\}$$

6.3 Results for New Test Instances

Tables 6, 7, 8 and 9 show the results for test runs on two- and three-dimensional instances. For ten test instances of any combination of parameters, we show how many of these instances we could solve within a time limit of 1000 seconds on a Sun Ultra SPARC with 175 MHz. From the solved instances, we show the minimum (Min), the average (Av), and the maximum (Max) of the number of OKP and OPP nodes, as well as the resulting runtimes.

It is evident that the difficulty grows with the percentage of "small" boxes. This is not very surprising, because these boxes do not restrict the possibilities for the rest of a selected subset as much as large or bulky boxes do.

The large difference in difficulty for instances with identical parameterization does not arise from our method of generation instances, but is characteristic for instances of hard combinatorial optimization problems. This effect has been known even for one-dimensional packing problems, which have a much simpler structure. Because of this spread, the number of nodes and runtimes are only significant for combinations of parameters where most instances could be solved.

For the OKP-2 with $m \leq 40$ and $|V| \leq 40$, we could find an optimal solution in tolerable runtime for almost all instances. For 60 and more boxes, classes II and III started to have higher numbers of instances that could not be solved within the time limit. Only for instances with 80 boxes and about 70 % of small boxes, our algorithm seemed to reach its limits for the current implementation.

Even when taking into account that classes of three-dimensional instances vary more with respect to the percentage of small boxes than those for two dimensions, it is remarkable that this percentage makes a huge difference with respect to the difficulty of the resulting instances. For an average of 20 % of small boxes (Class I), all instances (with the exception of a single one with $|V| = 80$) could be solved. For an average of 40 % of small boxes (Class II), our method works well, at least for instances with $m \leq 40, |V| \leq 40$. If the percentage of small boxes rises to 60 % in Class III, then even for $m = 30, |V| = 30$, our program does not find an optimal solution for a large number of instances.

Summarizing, we can say that our new method has greatly increased the size of instances that are practically solvable. In particular, the size of the container is no longer a limiting factor. It should be noted that even for three-dimensional instances with $m = 20$, the 0-1 programs following the approach by Beasley and Hadjiconstantinou/Christofides contain several 100,000 variables, even making the generous assumption of a grid reduction to 10 %.

Class of box types	$w_1(x)$ evenly distributed on	$w_2(x)$ evenly distributed on
1 (Bulky in 2)	$[1, 50]$	$[75, 100]$
2 (Bulky in 1)	$[75, 100]$	$[1, 50]$
3 (Large)	$[50, 100]$	$[50, 100]$
4 (Small)	$[1, 50]$	$[1, 50]$

Instance type	Classes of box types			
	1	2	3	4
I	20 %	20 %	20 %	40 %
II	15 %	15 %	15 %	55 %
III	10 %	10 %	10 %	70 %

Table 4: Random generation of OKP-2 test instances

Class of box types	$w_1(x)$ evenly distributed on	$w_2(x)$ evenly distributed on	$w_3(x)$ evenly distributed on
1 (bulky in 2,3)	$[1, 50]$	$[75, 100]$	$[75, 100]$
2 (bulky in 1,3)	$[75, 100]$	$[1, 50]$	$[75, 100]$
3 (bulky in 1,2)	$[75, 100]$	$[75, 100]$	$[1, 50]$
4 (large)	$[50, 100]$	$[50, 100]$	$[50, 100]$
5 (small)	$[1, 50]$	$[1, 50]$	$[1, 50]$

Instance type	Class of box types				
	1	2	3	4	5
I	20 %	20 %	20 %	20 %	20 %
II	15 %	15 %	15 %	15 %	40 %
III	10 %	10 %	10 %	10 %	60 %

Table 5: Random generation of OKP-3 test instances

Class	m	$\|V\|$	Solved (out of 10)	# OKP nodes			# OPP nodes		
				Min	Av	Max	Min	Av	Max
I	20	20	10	5	57	174	0	8	23
	30	30	10	19	307	914	0	158	969
	40	40	9	40	933	3826	0	431	2411
	20	60	9	31	231	677	1	287	1018
	20	80	8	82	336	943	234	147864	727719
II	20	20	10	6	139	1038	0	305	2699
	30	30	10	28	548	1568	0	2873	26866
	40	40	8	32	5062	28754	2	14910	84975
	20	60	7	36	297	571	5	237144	1633573
	20	80	6	62	536	1110	83	168530	280688
III	20	20	10	3	117	516	0	299	1169
	30	30	10	82	737	1860	3	10588	53510
	40	40	9	342	3865	10655	745	62065	416200
	20	60	8	31	1006	4064	3	345130	1174938
	20	80	2	96	196	296	241	85729	171218

Table 6: Results for randomly generated OKP-2 instances

Class	m	$\|V\|$	Solved (out of 10)	runtime/s		
				Min	Av	Max
I	20	20	10	0.06	0.56	1.26
	30	30	10	0.29	4.48	13.59
	40	40	9	1.31	22.02	76.30
	20	60	9	0.43	2.35	5.95
	20	80	8	0.99	62.46	243.41
II	20	20	10	0.06	2.18	17.69
	30	30	10	0.36	10.64	39.59
	40	40	8	0.55	51.12	152.46
	20	60	7	0.45	95.44	640.47
	20	80	6	1.86	112.89	267.90
III	20	20	10	0.08	1.48	5.77
	30	30	10	1.07	17.67	53.00
	40	40	9	6.66	103.10	313.91
	20	60	8	0.36	191.98	719.67
	20	80	2	2.18	34.52	66.86

Table 7: Runtimes for randomly generated OKP-2 instances using a Sun Ultra SPARC with 175MHz.

Class	m	$\|V\|$	Solved (out of 10)	# OKP nodes			# OPP nodes		
				Min	Av	Max	Min	Av	Max
I	20	20	10	1	73	352	0	22	82
	30	30	10	11	276	1190	1	59	291
	40	40	10	73	953	2848	5	2684	20975
	20	60	10	20	541	2961	3	19896	198091
	20	80	9	42	414	1511	14	145	399
II	20	20	10	11	75	328	1	35	166
	30	30	10	5	327	972	0	6579	62827
	40	40	8	59	2197	13064	20	85465	671934
	20	60	5	1	292	719	0	232	912
	20	80	3	142	149	161	23	46	65
III	20	20	10	5	57	138	0	4433	36747
	30	30	6	1	859	2250	1	3794	10063
	40	40	3	17	652	1715	7	1326	3885
	20	60	3	51	3728	10842	27	55164	165276
	20	80	1	73	73	73	38	38	38

Table 8: Results for randomly generated OKP-3 instances.

Class	m	$\|V\|$	Solved (out of 10)	runtime/s		
				Min	Av	Max
I	20	20	10	0.06	1.63	7.76
	30	30	10	0.36	9.15	43.58
	40	40	10	2.66	44.99	121.96
	20	60	10	0.50	18.33	125.76
	20	80	9	0.67	10.76	37.04
II	20	20	10	0.26	1.76	6.92
	30	30	10	0.37	18.94	81.69
	40	40	8	2.26	133.48	845.70
	20	60	5	0.28	12.00	38.94
	20	80	3	4.13	5.43	6.95
III	20	20	10	0.29	4.73	21.69
	30	30	6	0.26	35.69	101.61
	40	40	3	2.01	29.66	78.53
	20	60	3	1.21	211.63	607.83
	20	80	1	2.06	2.06	2.06

Table 9: Runtimes for randomly generated OKP-3 instances on a Sun Ultra SPARC with 175MHz, timeout after 1000s.

6.4 A New Library of Benchmark Instances

We are in the process of setting up a new library for multi-dimensional packing problems, called PackLib^2 [18]. The idea is to have one place where benchmark instances, results, and solution history can be found. For this purpose, we are using a universal XML-format that allows inclusion of all this information. We provide parsers for conversion directly into C formats, and converters for all standard data formats. Results include visualization of solutions by drawings of the feasible packings. Finally, we hope to provide a number of algorithms at the website. Interested researchers are encouraged to contact the first author.

7 Solving Other Types of Packing Problems

7.1 Strip Packing Problems

In an exact SPP procedure, we start with a heuristic for generating a packing; its "height" is used as an upper bound. A lower bound $\underline{h}$ can be obtained with the help of the methods described in our paper [16]. If there is a gap between these bounds, we have to use enumeration. Because the OPP is the decision version of the SPP for a fixed objective value, an obvious approach would be binary search in combination with the OPP algorithm from Section 3.

A more efficient method can be obtained by observing that any OPP node that did not find a solution for height h cannot possibly find a solution for height $h' < h$. Thus, we can solve the SPP with the help of a modified version of our OPP routine.

For finitely many boxes, there are only finitely many possibilities for the minimal height of a packing. The set H of these values can be determined by using the method from [7] for computing normalized coordinates.

The height of the packing obtained by the heuristic is stored under h. The variable h' is initialized with the largest value from H below h. We start the OPP tree search for the container with height h'. If the algorithm finds a feasible packing, then h is updated to the value h', and h' is replaced by the next smaller value of H. Now the OPP search is done for container height h'. As noted above, no search node that was dismissed before has to be considered again. The search is performed until all search nodes have been checked, or h reaches the value $\underline{h}$ of the lower bound.

7.2 Orthogonal Bin Packing Problems

The basic scheme of our exact method follows the outline by Martello and Vigo [30], and Martello, Pisinger, and Vigo [28]:

Within a branch-and-bound framework, a packing (for a number of containers) is produced iteratively. A list L maintains all containers that are used. In the beginning, L is empty. At each branching step, a box b is either assigned to a container C in L, or a new container is generated for b, and added to L. The crucial step is to check whether a container C can hold all boxes that are assigned to it.

We get upper bounds by packing the unassigned boxes heuristically. Our new suggestions concern the other steps of the approach, which cause the largest computational effort:

1. computing lower bounds
2. solving the resulting OPPs.

The improvement of the lower bounds from [30] and [28] have already been discussed in our paper [16].

In [30], the resulting OPPs are enumerated by using the method of Hadjiconstantinou/Christofides. For solving the three-dimensional OPPs in [28], there is a special enumeration scheme, using the principle of placement points described in Section 5.1. As discussed in our paper [13, 15], we get a drastic improvement by using our method from Section 3 that is based on packing classes.

8 Conclusion

In this paper, we have shown that higher-dimensional packing problems of considerable size can be solved to optimality in reasonable time, by making use of a structural characterization of feasible packings. Further

progress may be achieved by refined lower bounds and by using a more sophisticated outer tree search, as in the recent paper by Caprara and Monaci [6]. Currently, we are working on such a more advanced implementation, motivated by ongoing research on reconfigurable computing, We expect this work to lead to progress for other problem variants.

Acknowledgments

We thank Eleni Hadjiconstantinou who provided the data for instances *hadchr8* and *hadchr12*. Thank you to Alberto Caprara for some interesting discussions. Two anonymous referees helped with a number of useful comments that improved the overall presentation of this paper.

References

[1] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. On the solution of traveling salesman problems. *Documenta Mathematica Journal der Deutschen Mathematiker-Vereinigung*, ICM III (1998), pp. 645–656.

[2] M. Arenales and R. Morabito. An AND/OR–graph approach to the solution of two–dimensional non–guillotine cutting problems. *European Journal of Operations Research*, **84** (1995), pp. 599–617.

[3] J. E. Beasley. An exact two–dimensional non–guillotine cutting stock tree search procedure. *Operations Research*, **33** (1985), pp. 49–64.

[4] J. E. Beasley. OR-Library: Distributing test problems by electronic mail. *Journal of the Operations Research Society*, **41** (1990), pp. 1069–1072.

[5] M. Biró and E. Boros. Network flows and non–guillotine cutting patterns. *European Journal of Operations Research*, **16** (1984), pp. 215–221.

[6] A. Caprara and M. Monaci. On the 2–dimensional knapsack problem. *Operations Research Letters*, **32** (2004), pp. 5–14.

[7] N. Christofides and C. Whitlock. An algorithm for two–dimensional cutting problems. *Operations Research*, **25** (1977), pp. 31–44.

[8] J. J. Dongarra. Performance of various computers using standard linear equations software. Working paper, University of Tennessee, 2004. Continuous updates are available at `http://www.netlib.org/benchmarks/performance.ps`.

[9] K. A. Dowsland. An exact algorithm for the pallet loading problem. *European Journal of Operations Research*, **31** (1987), pp. 78-84.

[10] S. P. Fekete, E. Köhler, and J. Teich. Higher-dimensional packing with order constraints. In *Algorithms and Data Structures (WADS 2001)*, Springer Lecture Notes in Computer Science, vol. 2125, 2001, pp. 192–204. Full version to appear in *SIAM Journal of Discrete Mathematics*.

[11] S. P. Fekete and J. Schepers. A new exact algorithm for general orthogonal d-dimensional knapsack problems, *Algorithms – ESA '97*, Springer Lecture Notes in Computer Science, vol. 1284, 1997, pp. 144-156.

[12] S. P. Fekete and J. Schepers. New classes of lower bounds for bin packing problems. *Mathematical Programming*, **91** (2001), pp. 11–31.

[13] S. P. Fekete and J. Schepers. A combinatorial characterization of higher-dimensional orthogonal packing. Journal version of [15]. *Mathematics of Operations Research*, **29** (2004), pp. 353–368.

[14] S. P. Fekete and J. Schepers. A general framework for bounds for higher-dimensional orthogonal packing problems. Journal version of [16]. *Mathematical Methods of Operations Research*, **60** (2004), pp. 81–94.

[15] S. P. Fekete and J. Schepers. On more-dimensional packing I: Modeling. *ZPR Technical Report 97–288.* Available at `http://www.zpr.uni-koeln.de/~paper`.

[16] S. P. Fekete and J. Schepers. On more-dimensional packing II: Bounds. *ZPR Technical Report 97–289.* Available at `http://www.zpr.uni-koeln.de/~paper`.

[17] S. P. Fekete and J. Schepers. On more-dimensional packing III: Exact Algorithms. *ZPR Technical Report 97–290.* Available at `http://www.zpr.uni-koeln.de/~paper`. Tech report version of this paper.

[18] S. P. Fekete and J. van der Veen. Packlib^2: A packing benchmark library. To appear in *European Journal on Operational Research.* Available at `http://www. math.tu-bs.de/packlib2`.

[19] E. P. Ferreira and J. F. Oliveira. A note on Fekete and Schepers' algorithm for the non-guillotinable two-dimensional packing problem. Technical report, available at `http://paginas.fe.up.pt/~jfo/techreports/Fekete%20and%20Schepers%200PP%20degeneracy.pdf`.

[20] M. R. Garey and D. S. Johnson. *Computers and Intractability: A guide to the theory of NP-completeness.* Freeman, San Francisco, 1979.

[21] A. Ghouilà-Houri. Caractérization des graphes non orientés dont on peut orienter les arrêtes de manière à obtenir le graphe d'une relation d'ordre. *C.R. Acad. Sci. Paris*, **254** (1962), pp. 1370–1371.

[22] P.C. Gilmore and A.J. Hoffmann. A characterization of comparability graphs and of interval graphs. *Canadian Journal of Mathematics*, **16** (1964), pp. 539–548.

[23] M. C. Golumbic. *Algorithmic graph theory and perfect graphs.* Academic Press, New York, 1980.

[24] M. Grötschel. On the symmetric travelling salesman problem: solution of a 120-city problem. *Mathematical Programming Study*, **12** (1980), pp. 61–77.

[25] E. Hadjiconstantinou and N. Christofides. An exact algorithm for general, orthogonal, two–dimensional knapsack problems. *European Journal of Operations Research*, **83** (1995), pp. 39–56.

[26] D. Hochbaum. *Approximation Algorithms for NP-hard Problems.* PWS Publishing, 1996.

[27] N. Korte and R. H. Möhring. An incremental linear–time algorithm for recognizing interval graphs. *Siam Journal of Computing*, **18** (1989), pp. 68-81.

[28] S. Martello, D. Pisinger, and D. Vigo. The three-dimensional bin packing problem. *Operations Research*, **48** (2000) pp. 256–267.

[29] S. Martello and P. Toth. *Knapsack Problems – Algorithms and Computer Implementations.* Wiley, Chichester, 1990.

[30] S. Martello and D. Vigo. Exact solution of the two-dimensional finite bin packing problem. *Management Science* **44** (1998), pp. 388–399.

[31] K. Mehlhorn. *Data Structures and Efficient Algorithms, Vol.1: Sorting and Searching.* Springer Verlag, 1984.

[32] G. Nemhauser und L. Wolsey. *Integer and Combinatorial Optimization.* Wiley, Chichester, 1988.

[33] M. Padberg. Packing small boxes into a big box. *Math. Methods Oper. Res.* **52** (2000), pp. 1–21.

[34] C. H. Papadimitriou. *Computational Complexity.* Addison Wesley, New York, 1994.

[35] J. Schepers. *Exakte Algorithmen für orthogonale Packungsprobleme.* Dissertation, Universität zu Köln, 1997. Available at `http://www.zpr.uni-koeln.de/~paper`.

[36] J. Teich, S. P. Fekete, and J. Schepers. Optimal hardware reconfigurations techniques. *Journal of Supercomputing*, **19** (2001), pp. 57–75.

[37] P. Y. Wang. Two algorithms for constrained two–dimensional cutting stock problems. *Operations Research*, **31** (1983), pp. 573-586.

[38] M. Wottawa. *Struktur und algorithmische Behandlung von praxisorientierten dreidimensionalen Packungsproblemen.* Dissertation, Mathematisches Institut, Universität zu Köln, 1996. Available at `http://www.zpr.uni-koeln.de/~paper`.