

Where are the hard knapsack problems?

David Pisinger

*Dept. of Computer Science, University of Copenhagen,
Universitetsparken 1, DK-2100 Copenhagen, Denmark*

February 2003*

Abstract

The knapsack problem is believed to be one of the “easier” \mathcal{NP} -hard problems. Not only can it be solved in pseudo-polynomial time, but also decades of algorithmic improvements have made it possible to solve nearly all standard instances from the literature. The purpose of this paper is to give an overview of all recent exact solution approaches, and to show that the knapsack problem still is hard to solve for these algorithms for a variety of new test problems. These problems are constructed either by using standard benchmark instances with larger coefficients, or by introducing new classes of instances for which most upper bounds perform badly. The first group of problems challenge the dynamic programming algorithms while the other group of problems are focused towards branch-and-bound algorithms. Numerous computational experiments with all recent state-of-art codes are used to show that the (KP) is still difficult to solve for a wide number of problems. One could say that the previous benchmark tests were limited to a few highly structured instances, which do not show the full characteristics of knapsack problems.

1 Introduction

The classical knapsack problem is defined as follows: We are given a set of n items, each item j having an integer profit p_j and an integer weight w_j . The problem is to choose a subset of the items such that their overall profit is maximized, while the overall weight does not exceed a given capacity c . We may formulate the model as the following integer programming model

$$\text{(KP)} \quad \text{maximize} \quad \sum_{j=1}^n p_j x_j \quad (1)$$

$$\text{subject to} \quad \sum_{j=1}^n w_j x_j \leq c, \quad (2)$$

$$x_j \in \{0, 1\}, \quad j = 1, \dots, n. \quad (3)$$

*Technical Report 2003/08, DIKU, University of Copenhagen, Denmark

where the binary decision variables x_j are used to indicate whether item j is included in the knapsack or not. Without loss of generality it may be assumed that all profits and weights are positive, that all weights are smaller than the capacity c , and that the overall weight of the items exceeds c .

From practical experience it is known that many (KP) instances of considerable size can be solved within reasonable time by exact solution methods. This fact is due to several algorithmic refinements which emerged during the last two decades. These include *advanced dynamic programming recursions*, the concept of solving a *core*, and the separation of *cover inequalities* to tighten the formulation. For a recent survey of the latest techniques see Martello, Pisinger, Toth [6] or the monograph by Kellerer, Pferschy, Pisinger [4].

The knapsack problem is \mathcal{NP} -hard in the weak sense, meaning that it can be solved in pseudo-polynomial time through dynamic programming. Meyer auf der Heide [10] showed that for the linear decision tree model (LDT) of computation no super-polynomial lower bound can exist. This negative result was extended by Fournier and Koiran [3] who showed that for even less powerful models of computation no super-polynomial lower bound is likely to exist.

The lack of theoretical upper and lower bounds on the computational complexity of (KP) leaves plenty of space for practical algorithmic development. Although none of these algorithms can ensure efficient solution times for all instances, progress is made on reaching acceptable running times for all “practically occurring” instances. In the following section we will give a short overview of the latest exact algorithms for the (KP), and state their worst-case complexity where possible. In Section 3 we experimentally measure the performance of these algorithms for a large variety of instance types, including two groups of new, difficult instances. The first group of difficult instances is based on large coefficients, while the second group contains six categories of structurally difficult instances with small coefficients. Although the classical problem instances are quite easy to solve for the most recent algorithms, it is interesting to see that the instances do not need to be changed much before the algorithms get a significantly different performance. In our search for algorithms which are able to solve all “practically occurring” instances it is important to be aware of these problems, and to extend our algorithms and in particular upper bounds to more robust variants. These thoughts are summarized in Section 4.

2 Exact algorithms for the knapsack problem

For the following discussion we need a few definitions. Assume that the items are sorted according to non-increasing efficiencies $\frac{p_j}{w_j}$, so that we have

$$\frac{p_1}{w_1} \geq \frac{p_2}{w_2} \geq \dots \geq \frac{p_n}{w_n}, \quad (4)$$

The LP-relaxation of (KP) can be solved through the greedy algorithm by simply filling the knapsack until item $s = \min\{h : \sum_{j=1}^h w_j > c\}$, which is also known as the *split item*. The LP-bound is then defined as

$$U_1 = \sum_{j=1}^{s-1} p_j + \left(c - \sum_{j=1}^{s-1} p_j \right) \frac{p_s}{w_s}$$

The split item s and hence also the LP-solution can be found in $O(n)$ time using a median-search algorithm presented by Balas and Zemel [1]. In the following we will denote z^* the integer optimal solution value and x^* the corresponding solution vector.

Various branch-and-bound algorithms for (KP) have been presented. The more recent of these solve a *core* problem, i.e. a (KP) defined on a subset of the items where there is a large probability of finding an optimal solution. The MT2 algorithm [7] is the most advanced of these algorithms. It starts by solving the core problem obtaining a lower bound z on the (KP) as well as an enumerative upper bound U . If $z = U$ it stops, otherwise it reduces the size of the instance by fixing variables at their optimal value. In the last phase, the reduced (KP) is solved to optimality.

Realizing that the core size is difficult to estimate in advance, Pisinger [11] proposed to use an expanding core algorithm, which simply starts with a core consisting of the split item only, and then adds more items to the core when needed. The proposed Expknapsack makes use of branch-and-bound where computationally cheap upper bounds from LP-relaxation are used.

Martello and Toth [9] proposed a special variant of the MT2 algorithm which was developed to deal with hard knapsack problems. The resulting algorithm MThard makes use of a new family of upper bounds based on the generation of additional *cardinality constraints* which again are Lagrangian relaxed to reach an ordinary (KP).

Straightforward use of dynamic programming leads to the well-known *Bellman recursion* [2] which solves the (KP) in pseudo-polynomial time $O(nc)$. Reversing the roles of profits and weights in the dynamic programming recursion leads to an algorithm with running time $O(nz^*)$.

Pisinger [14] presented an improved variant of the Bellman recursion which runs in time $O(nm/\log m)$ on a word RAM, where $m = \max\{c, z^*\}$. For the most difficult instances where c and z^* are of the same magnitude, this leads to a logarithmic improvement over the Bellman recursion.

A balanced solution is a solution which can be reached from the greedy solution through a number of inserts and removals which maintain the knapsack filled close to the capacity. Using this concept, Pisinger [13] introduced a dynamic programming recursion with time and space complexity $O(nw_{\max}\Gamma)$ where $p_{\max} = \max_{j=1,\dots,n} p_j$, $w_{\max} = \max_{j=1,\dots,n} w_j$ and $\Gamma = z^* - z$ for any lower bound z . Since $\Gamma \leq p_{\max}$ — as we may choose z equal to the greedy solution — the complexity is bounded by the term $O(nw_{\max}p_{\max})$ which is linear in n when the magnitude of the weights and profits is bounded by a constant. Note that the practical running times may be improved by finding a lower bound z of good quality. If we are so lucky that the optimal solution has been found by some initial heuristic so that the balanced dynamic programming algorithm is used only for proving optimality, the time and space complexity is limited to $O(nw_{\max})$.

A final dynamic programming recursion was presented by Pisinger [12] in the framework of an *expanding core* algorithm. The Minknapsack algorithm is based on the ordinary Bellman recursion, but the recursion starts from the greedy solution and alternate between an insertion or a removal of an item. The running time is bounded by $O(n(b-a))$ where $\{a, \dots, b\}$ is the set of items which needed to be enumerated. This set is also known as the “*minimal*” core of the problem.

The currently most successful algorithm for (KP) was presented by Martello, Pisinger and

Toth [5]. The algorithm can be seen as a combination of many different concepts and is hence called Combo. The enumeration part of Combo is based on the same dynamic programming algorithm as Minknap but additional techniques are gradually introduced if the problem seems to be hard to solve. A good indication of the hardness of the problem is the number of states in the dynamic programming algorithm — if it surpasses a given threshold value, we take it as an indication of need for additional techniques. Depending on the hardness of the problem, rudimentary divisibility techniques are used, followed by bounds from cardinality constraints, and finally followed by improved lower bounding through merging of items with the states in the dynamic programming. The three techniques are introduced gradually, and in most cases they result in a decrease of the states in the dynamic programming algorithm. In the worst case, we proceed with the ordinary Minknap recursion, but hopefully having tighter upper and lower bounds available.

3 Computational Experiments

Despite the \mathcal{NP} -hardness of the knapsack problem we strive towards developing algorithms which efficiently can solve a large variety of problems occurring in practice. Unfortunately, very few real-life instances of (KP) are reported in the literature, hence algorithm design has focused on a set of synthetic benchmark tests. These tests are all based on randomly distributed profits or weights, and hence implicitly contain some structure which may not be present in real-life instances.

In the following study we will consider nearly all the knapsack algorithms presented in Section 1. This includes the MT2, MThard, Combo, Expknap, and Minknap algorithm. The dynamic programming algorithm based on balancing has never been implemented, although its worst-case complexity makes it a promising candidate. The word-RAM algorithm is very complicated and the constants hidden in the big-Oh notation are large. Hence it is mainly interesting from a theoretical point of view. Concerning the performance of various older algorithms the comparison with MT2 reported in [8] should be consulted.

We will consider several groups of randomly generated instances of (KP) which have been constructed to reflect special properties that may influence the solution process. In all instances the weights are uniformly distributed in a given interval with *data range* $R = 1000$ and 10000 . The profits are expressed as a function of the weights, yielding the specific properties of each group. The instance groups are graphically illustrated in Figure 1.

- **Uncorrelated data instances:** p_j and w_j are chosen randomly in $[1, R]$. In these instances there is no correlation between the profit and weight of an item. Such instances illustrate those situations where it is reasonable to assume that the profit does not depend on the weight. Uncorrelated instances are generally easy to solve, as there is a large variation between the profits and weights, making it easy to fathom numerous variables by upper bound tests or by dominance relations.
- **Weakly correlated instances:** weights w_j are chosen randomly in $[1, R]$ and the profits p_j in $[w_j - R/10, w_j + R/10]$ such that $p_j \geq 1$. Despite their name, weakly correlated in-

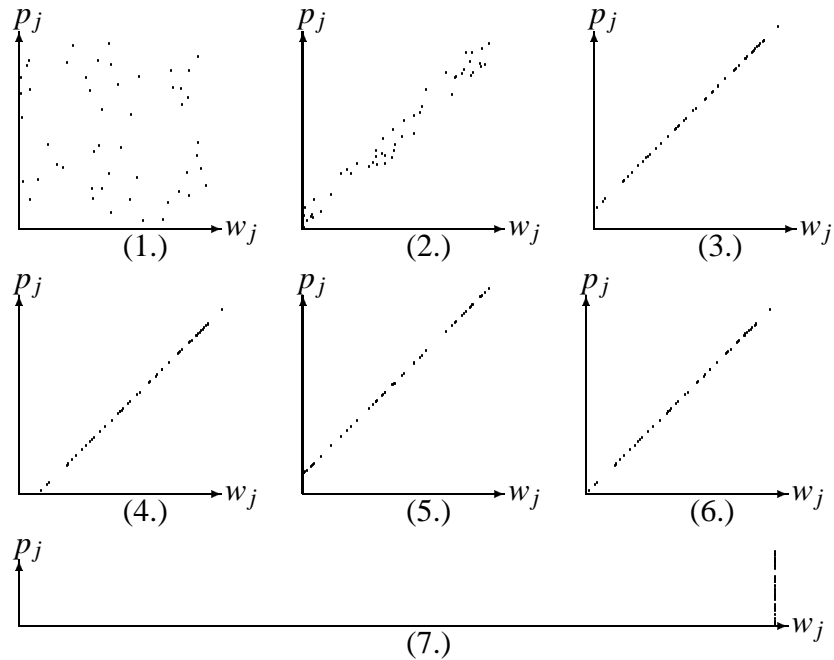


Figure 1: Classical test instances. (1.) Uncorrelated instances. (2.) Weakly correlated instances. (3.) Strongly correlated instances. (4.) Inverse strongly correlated instances. (5.) Almost strongly correlated instances. (6.) Subset sum instances. (7.) Uncorrelated instances with similar weights. Note that instances (3.) and (5.) look very similar since the extra “noise” in almost strongly correlated instances is very small.

stances have a very high correlation between the profit and weight of an item. Typically the profit differs from the weight by only a few percent. Such instances are perhaps the most realistic in management, as it is well-known that the return of an investment is generally proportional to the sum invested within some small variations.

- **Strongly correlated instances:** weights w_j are distributed in $[1, R]$ and $p_j = w_j + R/10$. Such instances correspond to a real-life situation where the return is proportional to the investment plus some fixed charge for each project. The strongly correlated instances are hard to solve for two reasons:
 - a) The instances are *ill-conditioned* in the sense that there is a large gap between the continuous and integer solution of the problem.
 - b) Sorting the items according to decreasing efficiencies correspond to a sorting according to the weights. Thus for any small interval of the ordered items (i.e. a “core”) there is a limited variation in the weights, making it difficult to satisfy the capacity constraint with equality.
- **Inverse strongly correlated instances:** profits p_j are distributed in $[1, R]$ and $w_j = p_j + R/10$. These instances are like strongly correlated instances, but the fixed charge is negative.

- **Almost strongly correlated instances:** weights w_j are distributed in $[1, R]$ and the profits p_j in $[w_j + R/10 - R/500, w_j + R/10 + R/500]$. These are a kind of fixed-charge problem with some noise added. Thus they reflect the properties of both strongly and weakly correlated instances.
- **Subset sum instances:** weights w_j are randomly distributed in $[1, R]$ and $p_j = w_j$. These instances reflect the situation where the profit of each item is equal (or proportional) to the weight. Thus, our only goal is to obtain a filled knapsack. Subset sum instances are however difficult to solve as instances to (KP) because most of the considered upper bounds return the same trivial value c , and thus we cannot use bounding rules for cutting off branches before an optimal solution has been found.
- **Uncorrelated instances with similar weights:** weights w_j are distributed in $[100000, 100100]$ and the profits p_j in $[1, 1000]$.

$n \setminus R$	uncorr.		weak.corr		str.corr		inv.str.corr		al.str.corr		subs.sum		sim.w
	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^5
50	0.1	0.0	0.1	0.0	5.8	4.7	1.2	1.8	2.9	3.0	0.0	0.8	1.7
100	0.0	0.1	0.1	0.1	2239.9	2235.9	359.7	15083.3	582.7	1607.7	0.1	0.8	328.7
200	0.1	0.1	0.3	0.4	—	—	—	—	—	—	0.2	2.9	—
500	0.1	0.4	0.7	1.1	—	—	—	—	—	—	0.5	1.8	—
1000	0.3	1.0	0.8	2.5	—	—	—	—	—	—	0.4	2.2	—
2000	1.0	1.4	0.9	4.1	—	—	—	—	—	—	0.7	2.3	—
5000	2.3	3.8	1.5	9.5	—	—	—	—	—	—	0.9	3.0	—
10000	2.9	4.9	2.3	14.5	—	—	—	—	—	—	1.4	3.2	—

Table I: Average solution times in milliseconds, MT2 (Intel Pentium III, 933 MHz).

$n \setminus R$	uncorr.		weak.corr		str.corr		inv.str.corr		al.str.corr		subs.sum		sim.w
	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^5
50	0.0	0.0	0.0	0.0	10.7	6.0	2.1	10.0	6.5	3.3	0.1	0.7	2.6
100	0.1	0.0	0.2	0.2	1438.4	3795.8	3517.3	11232.4	3604.9	2452.2	0.0	0.8	2289.8
200	0.0	0.1	0.2	0.4	—	—	—	—	—	—	0.1	0.8	—
500	0.1	0.2	0.6	1.2	—	—	—	—	—	—	0.1	1.0	—
1000	0.5	0.1	0.4	2.9	—	—	—	—	—	—	0.0	0.9	—
2000	0.4	0.7	0.6	7.4	—	—	—	—	—	—	0.2	1.4	—
5000	1.0	2.3	0.7	21.1	—	—	—	—	—	—	0.3	1.5	—
10000	1.3	5.3	1.5	34.3	—	—	—	—	—	—	0.7	1.8	—

Table II: Average solution times in milliseconds, Expknapsack (Intel Pentium III, 933 MHz).

We will compare the performance of MT2, Expknapsack, Minknapsack, MThard and Combo. The behavior of the algorithms will be considered for different problem sizes n , different problem types, and two data ranges. For each instance type a series of $H = 100$ instances is performed. The capacity in each instance is chosen as $c = \frac{h}{H+1} \sum_{j=1}^n w_j$, for test instance number $h = 1, \dots, H$.

$n \setminus R$	uncorr.		weak.corr		str.corr		inv.str.corr		al.str.corr		subs.sum		sim.w
	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^5
50	0.1	0.0	0.0	0.1	0.4	2.3	0.4	2.5	0.4	0.3	0.3	5.4	0.0
100	0.1	0.0	0.1	0.2	1.9	31.1	1.4	33.5	1.5	3.9	0.2	6.1	0.1
200	0.0	0.0	0.1	0.2	5.6	163.2	4.4	122.0	5.1	20.9	0.2	6.9	0.6
500	0.2	0.1	0.0	0.2	20.0	523.0	20.6	594.7	17.1	155.6	0.3	5.9	3.1
1000	0.1	0.1	0.2	0.9	53.9	1829.0	51.7	1655.1	43.6	625.0	0.1	6.0	12.1
2000	0.5	0.5	0.3	1.2	113.9	2899.1	141.6	3196.0	86.2	1863.0	0.3	6.7	40.7
5000	0.8	1.0	0.4	1.9	542.6	12403.1	454.0	12571.5	216.7	6103.8	0.3	7.5	150.5
10000	1.2	1.3	1.0	3.4	1301.9	26713.6	954.9	27972.2	250.5	11274.7	1.1	7.8	167.4

Table III: Average solution times in milliseconds, Minknap (Intel Pentium III, 933 MHz).

$n \setminus R$	uncorr.		weak.corr		str.corr		inv.str.corr		al.str.corr		subs.sum		sim.w
	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^5
50	0.1	0.0	0.0	0.1	0.7	1.1	0.6	0.6	1.2	5.5	0.1	0.7	0.5
100	0.1	0.0	0.1	0.1	1.4	2.0	1.1	1.8	3.5	25.5	0.1	0.7	1.1
200	0.1	0.1	0.5	0.4	8.0	9.1	4.0	5.2	11.5	60.1	0.3	2.2	5.4
500	0.4	0.5	1.0	1.0	18.9	18.9	11.0	12.4	20.5	131.6	0.4	1.8	12.4
1000	1.0	1.1	1.4	2.1	32.0	39.6	18.6	24.1	39.0	182.3	0.5	3.5	23.6
2000	1.7	1.8	1.6	4.4	42.5	66.4	26.2	33.7	63.5	182.5	0.4	2.0	40.1
5000	2.9	6.0	2.8	9.1	54.8	225.0	42.9	92.2	316.2	348.6	0.6	2.9	54.5
10000	4.9	10.4	3.2	17.9	92.7	847.6	64.3	338.5	—	—	1.4	147.9	76.6

Table IV: Average solution times in milliseconds, MThard (Intel Pentium III, 933 MHz).

All tests were run on an Intel Pentium III, 933 MHz, and a time limit of 1 hour was assigned to each instance type for all H instances. If any of the instances were not solved within the time limit, this is indicated by a dash in the table.

Tables I to V compare the solution times of the five algorithms. The simple branch-and-bound codes MT2 and Expknap, have the overall worst performance, although they are quite fast on easy instances like the uncorrelated and weakly correlated ones. Moreover they are the fastest codes for the subset sum instances. For the different variants of strongly correlated instances, MT2 and Expknap are able to solve only tiny instances.

The dynamic programming algorithm Minknap has an overall stable performance, as it is able to solve all instances within reasonable time. It is the fastest code for uncorrelated and weakly correlated instances and it has almost the same good performance for subset sum instances as the branch-and-bound algorithms. The strongly correlated instances take considerably more time to be solved but although Minknap uses only simple bounds from LP-relaxation, the pseudo-polynomial time complexity gets it safely through these instances.

The MThard algorithm has a very good overall performance, as it is able to solve nearly all problems within split seconds. The short solution times are mainly due to its cardinality bounds which make it possible to terminate the branching after having explored a small number of nodes. There are however some anomalous entries for large-sized almost strongly correlated problems, where the cardinality bounds somehow fail. Also for some large-sized subset sum problems MThard takes unproportionally long time.

The best performance is obtained with the Combo algorithm. This “hybrid” algorithm solves

$n \setminus R$	uncorr.		weak.corr		str.corr		inv.str.corr		al.str.corr		subs.sum		sim.w
	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^3	10^4	10^5
50	0.0	0.0	0.1	0.0	0.2	1.4	0.3	1.3	0.2	0.4	0.2	4.1	0.1
100	0.1	0.0	0.1	0.1	0.8	4.2	0.7	4.6	0.8	1.4	0.1	3.6	0.2
200	0.1	0.2	0.3	0.2	1.7	5.5	1.2	5.7	1.4	3.7	0.0	5.9	0.5
500	0.1	0.3	0.2	0.2	1.3	8.3	1.8	5.6	1.7	2.0	0.0	5.3	2.4
1000	0.1	0.3	0.4	0.7	1.6	10.4	2.4	5.6	3.1	1.4	0.3	4.3	4.7
2000	0.0	0.6	0.4	1.2	2.8	5.3	3.0	7.6	2.2	2.5	0.3	5.2	6.2
5000	0.7	1.0	0.3	3.2	4.2	5.8	4.4	6.4	3.6	4.6	0.4	4.0	10.2
10000	1.6	2.0	0.9	4.8	9.0	8.7	8.8	11.3	8.3	9.6	0.7	3.0	13.1

Table V: Average solution times in milliseconds, Combo (Intel Pentium III, 933 MHz).

all the considered instances within 15 milliseconds on average and the running times have a very systematic growth rate with small variations. In particular, it is worth noting that the ratio between solving a difficult instance and an easy instance is within a factor of 10, thus showing that the additional work needed to derive tight bounds can be done very fast.

Based on the results in Table V one could draw the wrong conclusion that (KP) is easy to solve. One should however notice that we consider instances where all the coefficients are of moderate size. If a real-life instance has this property (or can be scaled down to satisfy the property without significantly affecting the solution), then the problems are indeed easy to solve. However, there still exist many applications where more or less intractable knapsack instances with very large coefficients occur.

3.1 Difficult Instances

There are two directions to follow when constructing difficult instances: One may either consider the traditional instances with larger coefficients. This will obviously make the dynamic programming algorithms run slower, but also the upper bounds get weakened since the gap to the optimal solution is scaled and cannot be closed by simply rounding down the upper bound to the nearest smaller integer. A second direction to follow is to construct instances where the coefficients are of moderate size, but where all currently used upper bounds have a bad performance.

3.2 Difficult Instances with Large Coefficients

Our first attempt to construct difficult instances is to use the “standard” instances from Section 3 for increasing data range R . The motivation being, that we know that the magnitude of the weights strictly affect the computational complexity of dynamic programming algorithms, hence in this way the effect of upper bound tests, reduction, and early termination becomes more clear.

For the following experiments, the codes were compiled using 64 bit integers, making it possible to run tests with weights of considerable size. For technical reasons it was not possible to modify the FORTRAN codes MT2 and MThard to 64 bit integers.

The outcome of the experiments is shown in Tables VI to VIII. It is interesting to observe that uncorrelated and weakly correlated instances are almost unaffected by the increased data

$n \setminus R$	uncorr.			weak.corr			str.corr			inv.str.corr			al.str.corr			subs.sum			sim.w 10^8
	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	
50	0.1	0.0	0.0	0.1	0.1	0.1	15.2	5.8	12.3	6.4	5.4	4.6	7.0	3.3	2.5	6.1	57.6	448.3	2.8
100	0.1	0.1	0.1	0.1	0.2	0.2	13689.6	2888.8	7088.6	2810.6	1223.9	4203.7	1478.1	3543.4	1126.6	8.0	58.2	677.9	2605.4
200	0.1	0.1	0.1	0.4	0.4	0.4	—	—	—	—	—	—	—	—	—	8.1	83.1	605.9	—
500	0.2	0.2	0.2	1.6	1.4	0.1	—	—	—	—	—	—	—	—	—	8.2	69.9	575.7	—
1000	0.6	0.6	0.7	3.8	4.3	1.7	—	—	—	—	—	—	—	—	—	11.1	64.0	778.8	—
2000	1.5	1.1	0.8	10.3	11.7	5.5	—	—	—	—	—	—	—	—	—	13.5	66.6	627.7	—
5000	3.3	1.7	1.6	36.5	1.6	1.5	—	—	—	—	—	—	—	—	—	15.5	77.3	659.6	—
10000	9.7	8.3	3.1	106.7	74.4	3.9	—	—	—	—	—	—	—	—	—	24.8	88.0	558.1	—

Table VI: Average solution times in milliseconds, Expknapp, for large range instances (Intel Pentium III, 933 MHz).

$n \setminus R$	uncorr.			weak.corr			str.corr			inv.str.corr			al.str.corr			subs.sum			sim.w 10^8
	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	
50	0.1	0.1	0.0	0.2	0.1	0.1	22.9	25.6	40.1	11.1	20.9	14.9	2.9	2.2	1.8	83.7	—	—	0.1
100	0.1	0.0	0.1	0.2	0.2	0.2	241.9	1183.4	—	172.1	776.1	—	13.4	15.5	12.4	86.9	—	—	0.6
200	0.2	0.1	0.2	0.4	0.6	0.4	1291.9	—	—	1294.8	—	—	66.3	80.2	66.3	96.2	—	—	1.8
500	0.2	0.3	0.3	1.1	1.3	0.3	5114.3	—	67.9	5119.4	—	0.0	398.0	463.4	1.3	96.7	—	2376.3	17.8
1000	0.5	0.8	0.6	2.5	2.7	1.6	—	—	—	14715.6	—	—	1123.5	1409.5	183.2	90.5	—	—	72.0
2000	1.6	1.1	1.0	5.0	5.1	2.8	—	—	—	31093.9	—	0.4	3532.1	4462.7	494.6	100.4	—	—	0.5
5000	3.4	1.9	1.8	12.5	2.0	2.2	—	—	—	—	0.9	—	12629.2	7.7	7.2	105.2	401.5	4656.0	1.8
10000	6.5	5.7	4.1	24.0	17.5	4.0	—	—	—	—	—	—	25760.2	4937.8	85.7	122.4	—	—	5782.3

Table VII: Average solution times in milliseconds, Minknapp, for large range instances (Intel Pentium III, 933 MHz).

range. However, for nearly all other instance types the problems become harder as the data range is increased. Several instances cannot be solved within the given time limit, and it is also interesting to note that the upper bounds used in Combo for some instances have difficulties in closing the gap between the upper and lower bound. The dynamic programming part of Minknapp and Combo ensures that the optimal solution is found in pseudo-polynomial time, but as the data range is increased this time bound starts to grow unacceptably high.

The main observations from the computational experiments with small-range instances are however still valid. First of all, we notice that the methods used for finding tighter upper and lower bounds in Combo do pay off, since in general Combo is able to solve considerably more instances than the “clean” dynamic programming version Minknapp. Next, we notice that branch-and-bound methods cannot compete with dynamic programming methods. There is actually a

$n \setminus R$	uncorr.			weak.corr			str.corr			inv.str.corr			al.str.corr			subs.sum			sim.w 10^8
	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	10^5	10^6	10^7	
50	0.1	0.1	0.1	0.2	0.2	0.2	14.1	81.8	175.0	15.0	—	—	2.4	1.9	1.9	38.9	601.8	—	0.2
100	0.0	0.1	0.1	0.2	0.2	0.4	159.7	984.6	—	19.3	454.2	—	8.2	9.7	7.2	23.3	604.5	—	0.6
200	0.1	0.2	0.2	0.6	0.6	0.5	1079.4	8705.7	—	28.8	751.2	—	29.7	35.5	20.2	14.7	346.4	—	2.1
500	0.3	0.5	0.3	1.2	1.6	0.4	4225.1	—	72.3	19.3	—	0.1	245.0	330.2	1.4	16.8	207.5	14.5	11.1
1000	0.7	0.7	0.6	2.5	2.8	1.6	—	—	—	18.8	445.9	—	284.3	485.2	144.0	15.8	145.1	—	14.6
2000	1.6	1.4	1.0	5.3	5.1	3.1	—	—	—	21.0	77.4	0.2	462.1	1113.7	140.4	15.9	62.2	—	0.7
5000	2.9	2.0	1.7	13.6	1.9	2.3	—	523.7	4237.8	33.8	0.8	—	982.4	7.9	6.2	17.3	16.0	15.7	2.2
10000	7.4	6.1	3.4	27.2	18.8	3.5	—	—	4947.6	47.1	56.4	—	4476.1	339.6	12.5	19.8	20.6	20.9	90.0

Table VIII: Average solution times in milliseconds, Combo, for large range instances (Intel Pentium III, 933 MHz).

single exception to this observation for the subset sum instances. The randomly generated subset sum instances have numerous solution to the decision problem and hence a branch-and-bound algorithm may terminate the search as soon as an optimal solution has been found.

3.3 Difficult Instances With Small Coefficients

It is more challenging to construct instances with small coefficients, where present algorithms perform badly. Obviously, the worst-case complexity of dynamic programming algorithms still holds, but one may construct the instances so that it is difficult to fathom states through upper bound tests when using dynamic programming with upper bounds. The following classes of instances have been identified as having the desired property.




Figure 2: New test instances considered. (1.) Uncorrelated spanner instances $\text{span}(2, 10)$. (2.) Weakly correlated spanner instances $\text{span}(2, 10)$. (3.) Strongly correlated spanner instances $\text{span}(2, 10)$. (4.) Multiple strongly correlated instances $\text{mstr}(3R/10, 2R/10, 6)$. (5.) Profit ceiling instances $\text{pceil}(3)$. (6.) Circle instances $\text{circle}(\frac{2}{3})$.

- **spanner instances** $\text{span}(v, m)$

These instances are constructed such that all items are multiples of a quite small set of items — the so-called spanner set. The spanner instances $\text{span}(v, m)$ are characterized by the following three parameters: v is the size of the spanner set, m is the multiplier limit, and finally we may have any distribution (uncorrelated, weakly correlated, strongly correlated, etc.) of the items in the spanner set.

More formally, the instances are generated as follows: A set of v items is generated with weights in the interval $[1, R]$, and profits according to the distribution. The items (p_k, w_k) in the spanner set are normalized by dividing the profits and weights with $m + 1$. The n items are then constructed, by repeatedly choosing an item (p_k, w_k) from the spanner set, and a

multiplier a randomly generated in the interval $[1, m]$. The constructed item has profit and weight $(a \cdot p_k, a \cdot w_k)$.

Three distributions of the spanner problems $\text{span}(s, m)$ have been considered: uncorrelated, weakly correlated, and strongly correlated. The multiplier limit was chosen as $m = 10$. Computational experiments showed that the instances became harder to solve for smaller spanner sets. Hence in the following we will consider the instances: uncorrelated $\text{span}(2, 10)$, weakly correlated $\text{span}(2, 10)$, and strongly correlated $\text{span}(2, 10)$.

- **multiple strongly correlated instances** $\text{mstr}(k_1, k_2, d)$

These instances are constructed as a combination of two sets of strongly correlated instances. Both instances have profits $p_j := w_j + k_i$ where $k_i, i = 1, 2$ is different for the two instances.

The multiple strongly correlated instances $\text{mstr}(k_1, k_2, d)$ are generated as follows: The weights of the n items are randomly distributed in $[1, R]$. If the weight w_j is divisible by d , then we set the profit $p_j := w_j + k_1$ otherwise set it to $p_j := w_j + k_2$. Notice that the weights w_j in the first group (i.e. where $p_j = w_j + k_1$) will all be multiples of d , so that using only these weights we can at most use $d \lfloor c/d \rfloor$ of the capacity. To obtain a completely filled knapsack we need to include some of the items from the second distribution.

Computational experiments showed that very difficult instances could be obtained with the parameters $\text{mstr}(3R/10, 2R/10, d)$. Choosing $d = 6$ results in the most difficult instances, but values of d between 3 and 10 can all be used.

- **profit ceiling instances** $\text{pceil}(d)$

These instances have the property that all profits are multiples of a given parameter d . The weights of the n items are randomly distributed in $[1, R]$, and the profits are set to $p_j = d \lceil w_j/d \rceil$.

The parameter d was experimentally chosen as $d = 3$, as this resulted in sufficiently difficult instances.

- **circle instances** $\text{circle}(d)$

The instances $\text{circle}(d)$ are generated such that the profits as function of the weights form an arc of a circle (actually an ellipsis). The weights are uniformly distributed in $[1, R]$ and for each weight w the corresponding profit is chosen as $p = d \sqrt{4R^2 - (w - 2R)^2}$.

Experimental results showed that difficult instances appeared by choosing $d = \frac{2}{3}$ which was chosen for the following experiments.

Since all the above “difficult” instances are generated with moderate data range, all the codes MT2, Expknapp, Minknapp, MThard and Combo could be used in the normal version using 32 bit integers. The outcome of the experiments is shown in Tables IX to XIII.

It appears, that the branch-and-bound algorithms MT2 and Expknapp are able to solve this kind of instances only for small values of n . The MThard algorithm, which is a combination of branch-and-bound and dynamic programming has a slightly better performance, being able to

n	uncorr. span(2,10)	weak. corr. span(2,10)	str. corr. span(2,10)	mstr(2,6)	pceil(3)	circle($\frac{2}{3}$)
20	0.1	0.6	0.5	0.0	0.2	0.1
50	9982.0	—	—	0.7	14.1	0.3
100	—	—	—	77.0	—	18.4
200	—	—	—	—	—	—
500	—	—	—	—	—	—
1000	—	—	—	—	—	—
2000	—	—	—	—	—	—
5000	—	—	—	—	—	—
10000	—	—	—	—	—	—

Table IX: Solution times in milliseconds, MT2, for difficult instances with small coefficients (Intel Pentium III, 933 MHz).

n	uncorr. span(2,10)	weak. corr. span(2,10)	str. corr. span(2,10)	mstr(2,6)	pceil(3)	circle($\frac{2}{3}$)
20	0.1	0.4	0.3	0.0	0.2	0.0
50	2563.7	—	—	1.3	18.7	0.2
100	—	—	—	68.5	—	50.0
200	—	—	—	—	—	29522.6
500	—	—	—	—	—	—
1000	—	—	—	—	—	—
2000	—	—	—	—	—	—
5000	—	—	—	—	—	—
10000	—	—	—	—	—	—

Table X: Solution times in milliseconds, Expknaps, for difficult instances with small coefficients (Intel Pentium III, 933 MHz).

solve instances of moderate size. Both of the two dynamic programming algorithms Minknap and Combo are able to solve also large-sized instances since the worst-case running time of $O(nc)$ is acceptable for small values of the data range. It is however interesting to observe, that the tighter bounds in Combo do not significantly improve the running times. Indeed, Minknap is able to solve some of the instances faster than Combo using only dynamic programming and computationally cheap upper bounds.

4 Conclusion

We have compared the solution times of all recent algorithms, using classical and new benchmark tests. The classical instances are generally easy to solve and hence one could faulty conclude that no more research is needed for the knapsack problem. The new classes of benchmark tests clearly show that this is not the case. There are numerous interesting instances for which all currently known upper bounds perform badly, and for which the running times of the algorithms are close to the worst-case time-bound.

Dynamic programming is one of our best approaches for solving difficult (KP), since this is

n	uncorr. span(2,10)	weak. corr. span(2,10)	str. corr. span(2,10)	mstr(2,6)	pceil(3)	circle($\frac{2}{3}$)
20	0.0	0.0	0.0	0.0	0.1	0.0
50	0.0	0.2	0.3	0.0	0.3	0.3
100	0.3	0.5	0.4	0.9	0.9	1.0
200	0.8	1.8	1.7	2.5	2.7	2.7
500	4.2	8.2	7.0	8.7	18.2	9.5
1000	17.0	30.4	25.9	17.8	48.9	20.3
2000	67.6	136.7	108.7	37.5	179.7	43.2
5000	377.9	826.9	671.7	134.7	1368.7	152.1
10000	1695.1	3036.8	2579.0	332.6	5120.9	262.9

Table XI: Solution times in milliseconds, Minknap, for difficult instances with small coefficients (Intel Pentium III, 933 MHz).

n	uncorr. span(2,10)	weak. corr. span(2,10)	str. corr. span(2,10)	mstr(2,6)	pceil(3)	circle($\frac{2}{3}$)
20	0.1	0.5	0.4	0.1	0.2	0.0
50	4.4	5.8	6.1	0.4	2.2	0.5
100	14.9	29.9	30.4	1.6	15.7	2.9
200	4618.6	6464.3	12986.1	11.6	28904.1	15.4
500	—	—	—	10945.0	—	79.0
1000	—	—	—	—	—	432.7
2000	—	—	—	—	—	—
5000	—	—	—	—	—	—
10000	—	—	—	—	—	—

Table XII: Solution times in milliseconds, MThard, for difficult instances with small coefficients.

the only solution method which gives us a worst-case guarantee on the running time, independently on whether the upper bounding tests will work well.

Research in upper bounds working on more general classes of problems should also be stimulated. The knapsack polytope is one of the best studied (see e.g. Weismantel [15]) and numerous valid inequalities have been found which can be used to strengthen the LP-formulation of the (KP). Although a branch-and-cut framework may seem as an over-kill for knapsack problems, such techniques may become necessary in order to solve the more difficult instances.

References

- [1] E. Balas and E. Zemel. An algorithm for large zero-one knapsack problems. *Operations Research*, 28:1130–1154, 1980.
- [2] R.E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- [3] H. Fournier and P. Koiran. Are lower bounds easier over the reals? In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC)*, pages 507–513, 1998.
- [4] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. Springer, 2003.

n	uncorr. span(2,10)	weak. corr. span(2,10)	str. corr. span(2,10)	mstr(2,6)	pceil(3)	circle($\frac{2}{3}$)
20	0.0	0.0	0.0	0.0	0.2	0.0
50	0.1	0.1	0.0	0.1	0.4	0.0
100	0.2	0.5	0.5	0.5	0.8	0.5
200	0.7	1.6	1.6	1.3	2.3	1.1
500	3.6	7.6	7.8	4.2	14.9	4.6
1000	13.0	29.6	29.6	8.2	45.1	12.9
2000	50.6	136.8	147.8	20.3	149.7	31.0
5000	295.6	926.8	969.7	80.2	1216.2	130.0
10000	1319.1	3450.0	4186.0	202.1	4661.4	247.4

Table XIII: Solution times in milliseconds, Combo, for difficult instances with small coefficients.

- [5] S. Martello, D. Pisinger, and P. Toth. Dynamic programming and strong bounds for the 0-1 knapsack problem. *Management Science*, 45:414–424, 1999.
- [6] S. Martello, D. Pisinger, and P. Toth. New trends in exact algorithms for the 0-1 knapsack problem. *European Journal of Operational Research*, 123:325–332, 2000.
- [7] S. Martello and P. Toth. A new algorithm for the 0-1 knapsack problem. *Management Science*, 34:633–644, 1988.
- [8] S. Martello and P. Toth. *Knapsack Problems: Algorithms and Computer Implementations*. J. Wiley, 1990.
- [9] S. Martello and P. Toth. Upper bounds and algorithms for hard 0-1 knapsack problems. *Operations Research*, 45:768–778, 1997.
- [10] F. Meyer auf der Heide. A polynomial linear search algorithm for the n -dimensional knapsack problem. *Journal of the ACM*, 31:668–676, 1984.
- [11] D. Pisinger. An expanding-core algorithm for the exact 0-1 knapsack problem. *European Journal of Operational Research*, 87:175–187, 1995.
- [12] D. Pisinger. A minimal algorithm for the 0-1 knapsack problem. *Operations Research*, 45:758–767, 1997.
- [13] D. Pisinger. Linear time algorithms for knapsack problems with bounded weights. *Journal of Algorithms*, 33:1–14, 1999.
- [14] D. Pisinger. Dynamic programming on the word RAM. *Algorithmica*, 35:128–145, 2003.
- [15] R. Weismantel. On the 0/1 knapsack polytope. *Mathematical Programming*, 77:49–68, 1997.