

Branch and price for submodular bin packing


Liding Xu

LIX CNRS, École Polytechnique


With: Claudia D'Ambrosio, Sonia Haddad Vanier, Emiliano Traversi

02/02/2023

Cloud computing is a over-competitive market.



Interests of research



Resource (memory, CPU, bandwidth, disk) allocation is key to success of cloud computing providers, e.g., AWS, GOOGLE, OVH, SCALEWAY, HETZNER, ALIBABA....

On demand allocations of virtual machines: we want to preallocate virtual machines in a physical machine while considering uncertainty.

- How to model the uncertainty? Adding some nonlinearity in the classical model.
- Is Dantzig-Wolfe reformulation applicable in mixed-integer nonlinear programming (MINLP)? Yes, it works.
- How is the practical quality of well-known heuristics/approximation algorithms for resource allocation? Good.

- 1 Branch-and-price for submodular binpacking
- 2 Dantzig-Wolfe reformulation
- 3 Conclusion

Table of Contents

- 1 Branch-and-price for submodular binpacking
- 2 Dantzig-Wolfe reformulation
- 3 Conclusion

Table of Contents

- 1 Branch-and-price for submodular binpacking
- 2 Dantzig-Wolfe reformulation
- 3 Conclusion

Dantzig-Wolfe reformulation and relaxation for MILP

Dantzig-Wolfe (DW): a reformulation technique for Mixed-Integer Linear Program (MILP) with decomposable structures.

- **Decomposition:** a **main** problem + a **pricing sub** problem, and the main problem has an huge LP relaxation.

Dantzig-Wolfe reformulation and relaxation for MILP

Dantzig-Wolfe (DW): a reformulation technique for Mixed-Integer Linear Program (MILP) with decomposable structures.

- **Decomposition:** a **main** problem + a **pricing sub** problem, and the main problem has an huge LP relaxation.
- **Main LP relaxation:** solved via **column generation**.
- **Pricing sub problem:** solved via **pricing algorithm**.
- **Integer DW reformulation:** solved via **branch-and-price algorithm**.

The performance of DW reformulation

- **Decomposition:** main problem (integer optimization) + sub problem.
- **the LP relaxation** of the main problem is sometimes called configuration LP.
- **Quality of relaxations:** stronger bounds than those of the naive LP relaxations.



Figure: Decomposable structures in coefficients matrix of an MILP

Classical binpacking (BP): packing unsplittable items into a minimal number of bins (capacity usage is **linear** w.r.t. allocation of items).

- **Practical issue:** sizes of items are **uncertain and unprovisioned?**

Classical binpacking (BP): packing unsplittable items into a minimal number of bins (capacity usage is **linear** w.r.t. allocation of items).

- **Practical issue:** sizes of items are **uncertain and unprovisioned?**
- **Motivating example:** all cloud computing providers function on virtualization techniques:
 - create virtual machines (**items**) on dedicated servers (**bins**);
 - receive many **random requests for virtual machines**;
 - some statistics on demands/data.

We consider *static* but *probabilistic* models (assumption on data):

- BP with chance constraints;
- Distributionally robust BP.

These models admit an *exact* MINLP reformulation (nice for us, no need for sampling): **submodular binpacking**.

Definition

A function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ is called a submodular function, if for every $x, y \in \{0, 1\}^n$, $f(x) + f(y) \geq f(\max(x, y)) + f(\min(x, y))$, where \max, \min are element-wise minimum and maximum.

Econ. explanation: you get less and less utility when paying more resources.

The submodular binpacking formulation

Submodular binpacking:

$$\min \sum_{j \in \mathcal{M}} y_j, \quad (1)$$

$$\sum_{i \in \mathcal{N}} a_i v_{ij} + \sigma \sqrt{\sum_{i \in \mathcal{N}} b_i v_{ij}} \leq c y_j, \quad \forall j \in \mathcal{M}, \quad (2)$$

$$\sum_{j \in \mathcal{M}} v_{ij} = 1, \quad \forall i \in \mathcal{N}, \quad (3)$$

$$v_{ij} \in \{0, 1\}, \quad \forall i \in \mathcal{N}, j \in \mathcal{M}, \quad (4)$$

$$y_j \in \{0, 1\}, \quad \forall j \in \mathcal{M} \quad (5)$$

Our first contribution: DW reformulation.

DW decomposition: main problem

Let \mathcal{P} be the set of all possible configurations of a bin. DW reformulation = set cover formulation:

$$\min \sum_{p \in \mathcal{P}} \lambda_p, \quad (6a)$$

$$\sum_{p \in \mathcal{P}} d_{ip} \lambda_p \geq 1, \quad \forall i \in \mathcal{N}, \quad (6b)$$

$$\lambda_p \in \{0, 1\}, \quad \forall p \in \mathcal{P}. \quad (6c)$$

DW decomposition: main problem

Let \mathcal{P} be the set of all possible **configurations** of a bin. DW reformulation = set cover formulation:

$$\min \sum_{p \in \mathcal{P}} \lambda_p, \quad (6a)$$

$$\sum_{p \in \mathcal{P}} d_{ip} \lambda_p \geq 1, \quad \forall i \in \mathcal{N}, \quad (6b)$$

$$\lambda_p \in \{0, 1\}, \quad \forall p \in \mathcal{P}. \quad (6c)$$

$|\mathcal{P}|$ is huge, so solving the LP relaxation via **pricing**.

DW decomposition: pricing problem

Find a configuration which yields the steepest descent of the main problem:

$$v_{\text{pricing}} = \max \sum_{i \in \mathcal{N}} \pi'_i x_i, \quad (7a)$$

$$\sum_{i \in \mathcal{N}} a'_i x_i + \sigma \sqrt{\sum_{i \in \mathcal{N}} b'_i x_i} \leq c, \quad (7b)$$

$$x_i \in \{0, 1\}, \quad \forall i \in \mathcal{N}. \quad (7c)$$

DW decomposition: pricing problem

Find a configuration which yields the steepest descent of the main problem:

$$v_{\text{pricing}} = \max \sum_{i \in \mathcal{N}} \pi'_i x_i, \quad (7a)$$

$$\sum_{i \in \mathcal{N}} a'_i x_i + \sigma \sqrt{\sum_{i \in \mathcal{N}} b'_i x_i} \leq c, \quad (7b)$$

$$x_i \in \{0, 1\}, \quad \forall i \in \mathcal{N}. \quad (7c)$$

A **convex MINLP**: solvable by CPLEX.

Our second contribution: Tailored MILP-based branch-and-cut.

An exact pricing algorithm: initial MILP relaxation

$$\sum_{i \in \mathcal{N}} a'_i x_i + \sigma \sqrt{\sum_{i \in \mathcal{N}} b'_i x_i} \leq c \equiv \sum_{i \in \mathcal{N}} b'_i x_i \leq (c - \sum_{i \in \mathcal{N}} a'_i x_i)^2 / \sigma^2.$$

1d piece-wise-linear concave overestimator for $q(w) := (c - w)^2$.




Figure: Graphs of the quadratic and its PWL over-estimator

An exact algorithm for pricing problem: Relaxation refining

Relaxation refining: posterior approximation, **lazy cuts** separated at **infeasible** binary points, via **linearization** of the convex function

$$\sum_{i \in \mathcal{N}} a'_i x_i + \sigma \sqrt{\sum_{i \in \mathcal{N}} b'_i x_i^2}.$$

Combined together, an exact algorithm faster than CPLEX's branch-and-cut algorithm.

Our third contribution: a meta pricing scheduler that replace exact pricing algorithm via heuristic algorithms.

When can we avoid to use the expensive exact pricing algorithm? Some hints from the dual bound.

Maser LP (MLP) (long time to converge):

$$\text{Dual bound } v_{\text{MLP}} = \min \sum_{p \in \mathcal{P}} \lambda_p : \sum_{p \in \mathcal{P}} d_{ip} \lambda_p \geq 1, \forall i \in \mathcal{N}, \lambda_p \in [0, 1], \forall p \in \mathcal{P}. \quad (8)$$

Maser LP (MLP) (long time to converge):

$$\text{Dual bound } v_{\text{MLP}} = \min \sum_{p \in \mathcal{P}} \lambda_p : \sum_{p \in \mathcal{P}} d_{ip} \lambda_p \geq 1, \forall i \in \mathcal{N}, \lambda_p \in [0, 1], \forall p \in \mathcal{P}. \quad (8)$$

Restricted Maser LP (RMLP) (solved during column generation):

$$v_{\text{RMLP}} = \min \sum_{p \in \mathcal{P}} \lambda_p : \sum_{p \in \mathcal{P}} d_{ip} \lambda_p \geq 1, \forall i \in \mathcal{N}, \lambda_p \in [0, 1], \forall p \in \mathcal{P}' \subseteq \mathcal{P}. \quad (9)$$

$$v_{\text{MLP}} \leq v_{\text{RMLP}}.$$

Recall that v_{pricing} is the pricing objective value (a maximum problem). Then, $v_{\text{RMLP}}/v_{\text{pricing}} \leq v_{\text{MLP}} \leq v_{\text{RMLP}}$. Thus, $v_{\text{RMLP}}/v_{\text{pricing}}$ is also a lower bound (Farley bound) .

When exact pricing cannot yield better Farley bound

Assume that we have a good and fast pricing heuristic (which also generates a column), we run it first

Theorem

Let v_{heur} be the solution value of the pricing heuristic, let v_{RMLP} be the optimum of RMLP, and let v_{ld} be the current local dual bound for the master problem. If $\frac{v_{RMLP}}{v_{heur}} \leq v_{ld}$, the exact algorithm cannot yield a better local dual bound than v_{ld} .

Because $v_{heur} \geq v_{price}$, $\frac{v_{RMLP}}{v_{heur}} \leq v_{ld}$ implies that $\frac{v_{RMLP}}{v_{price}} \leq v_{ld}$ (no need for exact pricing).

Settings:

- Compact: compact formulation solved by CPLEX.
- DW: DW reformulation with pricing problem solved by CPLEX.
- DW+: DW reformulation with tailored pricing algorithm.
- DW-hybrid: DW+ with a hybrid strategy using heuristics or exact pricing algorithms.

- Randomly generated instances: number of items in $\{100, 400, 1000\}$.
- Intel Xeon @ 3.90GHz and 126GB memory, 3600 CPU second time limit.
- Every setting is given by a solution found by a greedy heuristic..

Experiment: results

Benchmarks	Solvers	Problem statistics			
		time (s)	dual gap (%)	primal improve(%)	#columns
CloudSmall ($ \mathcal{N} = 100$)	BSOCP-BC	1452	15.8	0.0	-
	DW-BC	2129	11.4	0.9	1373
	DW-PWL	633	2.4	2.7	1869
	DW-Hybrid	330	2.0	3.4	3485
CloudMedium ($ \mathcal{N} = 400$)	BSOCP-BC	3600	100.0	0.0	-
	DW-BC	3600	39.0	0.1	861
	DW-PWL	3600	17.2	0.4	3372
	DW-Hybrid	3600	11.8	0.6	6879
CloudLarge ($ \mathcal{N} = 1000$)	BSOCP-BC	3600	100.0	0.0	-
	DW-BC	3600	59.6	0.0	741
	DW-PWL	3600	43.1	0.2	2105
	DW-Hybrid	3600	34.2	0.4	4257

Table: Aggregated statistics of the main computational results

- **Experiment:** provides better dual bounds than CPLEX's bounds for compact formulation (13%+ improvement).
- **Conclusion:** greedy algorithms perform quite well (usually better than theoretical performance guarantee) (5% to optimality).

Related paper: Xu, D'Ambrosio, Haddad Vanier, and Traversi. "Branch and Price for Submodular Bin Packing", European journal on computational optimization, 2023.

Table of Contents

- 1 Branch-and-price for submodular binpacking
- 2 Dantzig-Wolfe reformulation
- 3 Conclusion

- In this paper, we explore DW reformulation for the submodular binpacking (MINLP);
- We propose several methods to generate columns efficiently;
- The results show that the DW reformulation is a promising method for the submodular binpacking (closing duality gap).

The end

Thank you for your attention!