Straddles and the Dantzig–Wolfe Bound of Cardinality-Indexed Knapsack Formulations

Md Shahrukh Anjum[1,*], Jitamitra Desai[2], and Ishwar Murthy[2]

[1]Amrut Mody School of Management, Ahmedabad University, Ahmedabad, India

[2]Decision Sciences Area, Indian Institute of Management Bangalore, Bengaluru, India

*Corresponding author: shahrukh.anjum@ahduni.edu.in

Abstract. Many integer programs assign indivisible items to capacitated resources, each contributing a knapsack constraint with a setup variable. Cardinality-indexed formulations copy the variables of each resource once for every number of items it may serve; they have been used for generalized assignment, concentrator location, capacitated Steiner arborescence and unsplittable network design. We characterize when their linear relaxation is exact. The relaxation of a resource equals its integer hull if and only if the resource has no straddle, a configuration that describes the fractional vertices of each fixed-cardinality level; in contrast to indexing by load, none of these vertices is absorbed when the levels are convexified together. If every resource is straddle-free, the relaxation attains the Dantzig–Wolfe bound under coupling through aggregated variables, and the condition is also necessary when the coupling can complete a straddle. With exact level capacities, a resource is straddle-free if and only if its weights take at most two values or its capacity is extreme at every level. This can be verified in polynomial time. When straddles occur, the integer hulls of the levels restore the Dantzig–Wolfe bound; we describe them on principal levels and characterize which rank-type covers define facets. Computational experiments illustrate the results.

Subject classifications: Programming: integer: theory: exactness of extended formulations; Programming: integer: cutting planes: facets of cardinality-constrained knapsacks; Networks/graphs: multicommodity: unsplittable network design.

Area of review: Optimization.

Keywords: extended formulation; Dantzig–Wolfe decomposition; knapsack polytope.

1 Introduction

Integer programs that assign indivisible items to capacitated resources are among the most studied models in operations research: jobs are assigned to agents in the generalized assignment problem, terminals to concentrators in capacitated concentrator location, and commodities to arcs in unsplittable network design. Each resource contributes a knapsack constraint together with a setup variable, and the linear relaxation is weak because a fractional item can fill exactly the capacity that no combination of whole items fills. A standard way to strengthen the relaxation is to replace each resource by the convex hull of its feasible configurations, which gives the Dantzig–Wolfe bound. When the resources are coupled only through their aggregated variables, this is the strongest bound obtainable by convexifying the resources independently, although not in general the strongest relaxation of the coupled problem; for generalized assignment it is the Lagrangian bound and the root bound of branch-and-price (Savelsbergh, 1997). Pseudo-polynomial arc-flow formulations attain this bound (Valério de Carvalho, 1999; de Lima et al., 2022; Yamín et al., 2026), but their size grows with the magnitude of the capacities.

A different extended formulation copies the variables of each resource once for every number of items that the resource may serve. A binary level variable states that the resource serves exactly that many items, and each copy carries the knapsack constraint together with an exact count of the items it holds. Murthy and Ransbotham (2019) introduced this cardinality-indexed formulation for generalized assignment, and it has since been adapted to capacitated concentrator location (Di Francesco et al., 2021), capacitated Steiner arborescence (Contu et al., 2026) and unsplittable multicommodity network design (Anonymous, 2026a). Its size is polynomial in the numbers of resources and items, independently of the capacities. In branch-and-cut procedures that combine reduced sets of levels with cardinality-based valid inequalities, these studies report root bounds that are markedly stronger than those of the compact formulation and smaller search trees, in some cases by orders of magnitude, on hard instances of standard benchmark sets. These studies establish that the formulation improves on the compact one; they do not establish when it attains the Dantzig–Wolfe bound. What has been missing is a structural explanation: on which resources is the cardinality-indexed relaxation as strong as possible, and where, and why, does it fall short of the Dantzig–Wolfe bound?

This paper answers that question. Our contribution is theoretical; the computational study illustrates and verifies the theoretical results and does not propose a new solution algorithm. We address two questions. First, when is the cardinality-indexed relaxation of a resource exact, and what does exactness imply for the bound of the whole model? Second, when it is not exact, what must be added to recover the Dantzig–Wolfe bound, and how can the integer hull of a single level be described?

1.1 Contributions

Exactness. Our first result characterizes when the cardinality-indexed relaxation of a resource equals its integer hull: this is the case if and only if the resource has no straddle. Fixing the number of selected items turns the linear relaxation of a knapsack into a level, the base polytope of a uniform matroid cut by one halfspace. It is known that its fractional vertices have at most two fractional components, and they admit a simple combinatorial description: a set of items selected fully together with two further items selected fractionally, the residual capacity left by the full items lying strictly between the weights of the two fractional ones. We call such a configuration a straddle (Figure 1 in Section 3 shows an example). We show that taking the convex hull over all levels never removes these vertices, so the relaxation of a resource has, as vertices, exactly its feasible loads and one point per straddle. This holds for every admissible choice of level capacities; the analogous property fails for other disjunctions. If the variables are indexed by the load instead of the count, a fractional vertex of one term can be a convex combination of integer points of other terms: with two items of weights one and two and a capacity of three, the term for load two contains the point that takes the first item fully and half of the second, and this point is the midpoint of the loads one and three. Along every edge of the unit cube the count changes by exactly one, whereas the load changes by an item weight, and this is why a fractional vertex of one level cannot be absorbed by the others. Although the proof is elementary, it implies that indexing by the attainable loads is exact only when all weights are equal, whereas indexing by the count is exact on a large class of resources that can be recognized efficiently.

Consequences for the Dantzig–Wolfe bound. Straddle-freeness of every resource guarantees the Dantzig–Wolfe bound whenever the resources are coupled only through their aggregated variables, as in the generalized assignment, concentrator location and unsplittable network design formulations; for the Steiner arborescence formulation, whose coupling retains level information, our conclusions concern individual arcs rather than the bound of the complete formulation. Conversely, a straddle that the coupling constraints can complete yields an objective for which the cardinality-indexed bound is strictly weaker. For generalized assignment with an uncapacitated outside agent, straddle-freeness of every agent is therefore necessary and sufficient for the bound to hold for every cost vector. For the other applications we state, application by application, which conclusions are sufficient, which are necessary, and which are both (Table 2 in Section 3).

Recognition. The level capacities that make the relaxation tightest are the exact ones, the largest weights of feasible sets with a given number of items; computing them is a cardinality-constrained subset-sum problem. We show that, under exact level capacities, a resource is straddle-free if and only if its weights take at most two distinct values or, at every level, its capacity either admits the heaviest sets or only the lightest ones. The proof is constructive: when the condition fails, it returns a straddle, and with it a violated cover inequality, in linear time from one maximum-weight feasible set. The test does not require the exact level capacities: straddle-freeness can be

decided in polynomial time from the weights and the original capacity alone, so the resources on which the relaxation is exact can be identified before the model is built. For two-valued resources we also characterize straddle-freeness under the original capacity, which three of the four applied formulations use, and the condition is more restrictive: a two-valued resource on which two or more levels bind is straddle-free under the original capacity if and only if its weights are consecutive multiples of their difference and its capacity is a multiple of that difference. Exactness does not make the integer problem easy: generalized assignment remains strongly NP-hard when every agent has two distinct weights.

Recovering the bound when exactness fails. The results for resources with straddles differ in scope, from a general reduction to complete descriptions of special cases. In general, the same disjunctive argument shows that replacing each level by its integer hull restores the Dantzig–Wolfe bound in the space of the cardinality-indexed formulation, which reduces the question to describing the integer hull of a single level. On a principal level, on which one feasible set dominates all others class by class, we give a complete description: the feasible sets are the bases of a nested matroid, the hull has a totally unimodular description with fewer inequalities than distinct weights, and principal levels can be recognized in polynomial time; if every level is principal, a polynomial-size linear program in the original level space attains the Dantzig–Wolfe bound without column generation. On the remaining levels, we characterize, by four exchange conditions, which rank-type covers define facets, these being the inequalities with zero-one coefficients that bound the number of items taken from a set, and we turn the characterization into a procedure that strengthens a cover toward a facet. This is a characterization within one family of inequalities, not a complete facet description, since rank-type covers do not describe every level. Finally, for three distinct weights one of whose gaps divides the other, a level is the cardinality slice of a two-coefficient knapsack polytope and inherits its known description (Section EC.4). Section 6.4 combines these results into a procedure that uses only the data: test every resource for straddles, add the principal-level inequalities where they apply, and separate over the level hulls, or use column generation, only on the remaining levels. No inequality on a single resource goes beyond the Dantzig–Wolfe bound, so any further strength must come from inequalities that couple several resources.

Computational results. On generalized assignment and network design instances, the cardinality-indexed bound coincides with the Dantzig–Wolfe bound on every class that the classification identifies as straddle-free, and elsewhere falls further below it as the share of levels with a straddle increases. On the network design instances of Table 6, the level-hull inequalities of Section 6 attain the Dantzig–Wolfe bound on every instance, and the principal-level inequalities suffice whenever all levels with a straddle are principal. On the standard GAP benchmarks, the recognition test finds no straddle-free agent, and the cardinality-indexed relaxation alone closes about one quarter of the gap, compared with four fifths for the Dantzig–Wolfe bound. The results are consistent with the pattern reported in the applied studies: the formulation is strongest when resources serve few items whose weights take few distinct values.

1.2 Related Literature

The object we study lies where two lines of work meet. The first concerns knapsack polytopes and their cardinality-constrained variants: the facial structure of knapsack sets with cardinality constraints (de Farias and Nemhauser, 2003; Zeng and Richard, 2011), surveyed by Hojny et al. (2020), and the property that the vertices of a knapsack relaxation with a cardinality constraint have at most two fractional components (Caprara et al., 2000). The second uses a count or value disjunction to build an extended formulation with a provably strong relaxation (Balas, 1998; Köppe et al., 2008); by the principle of variable redefinition (Martin, 1987), an extended formulation that describes the convex hull of every subproblem attains the Dantzig–Wolfe bound (Conforti et al., 2014), and for hop-constrained tree problems such a reformulation attains the Lagrangian bound (Gouveia, 1998). Closest in aim, Chen et al. (2025) recover the Dantzig–Wolfe bound in the original variable space with one cut per block, derived from the pricing problem at optimal dual multipliers, and show that the resulting formulation is less dual degenerate and solves faster on knapsack-based problems. Their cuts are tied to one objective and require the Dantzig–Wolfe dual to be solved first; the level-hull inequalities of Section 6 describe single levels independently of the objective, and straddle-free blocks need none. Value disjunctions are exact because they keep the integer hull of each term (Köppe et al., 2008); the cardinality-indexed formulation keeps only the linear relaxation of each level, and our results say when this suffices.

de Farias and Nemhauser (2003) show that adding a cardinality inequality to a continuous support-constrained knapsack gives the convex hull if and only if an exchange condition on the weights holds, a condition of the same exchange type as the absence of straddles; their set has a single capacity and the count enters as an inequality, whereas our levels are binary, the count is an equation imposed level by level, and the capacities depend on the level. The classification of hypersimplex splits by Herrmann and Joswig (2008) gives an alternative, nonconstructive route to the necessity part of our recognition theorem. For discretized and capacity-indexed formulations of location and tree problems, the literature compares the linear-programming bound with that of the compact formulation (Gouveia, 1995; Gouveia and Saldanha-da-Gama, 2006; Correia et al., 2010; Uchoa et al., 2008), and Frangioni and Gendron (2009) relate multiple-choice reformulations of capacities to residual capacity inequalities; none characterizes when a disjunction that keeps only the linear relaxation of each term is exact. The feasible sets of a level form a shifted family, and shifted matroids, also called nested or Schubert matroids, are well studied (Ardila, 2003; Hampe, 2017); matroid methods, including the description of cardinality-constrained matroid polytopes by Maurras and Stephan (2011), apply to a level exactly when it is principal. A complete linear description of the knapsack polytope with two distinct weights is given by Weismantel (1996), and a totally dual integral description for the weights one and p by Dahl and Foldnes (2003); we use the adjacency criterion of Geist and Rodin (1992) and, for more than two coefficients, Graver bases (Graver, 1975; De Loera et al., 2013) and primitive partition identities (Diaconis et al., 1996). In network design, single-arc sets for unsplittable flows are knapsack sets with a binary design variable (Atamtürk and Rajan, 2002; Chen et al., 2021), and multi-arc inequalities and reformulations are

studied by Raack et al. (2011), Chouman et al. (2017) and Greening et al. (2026); see Crainic and Gendron (2021) for a review of exact methods.

Relation to prior results. To the best of our knowledge, the main new results are the following three: that a disjunction over the count keeps every fractional vertex of every level, so that the cardinality-indexed relaxation of a resource is exact if and only if the resource is straddle-free, for every admissible level capacity (Theorem 2(b)); that indexing by the attainable loads is exact only when all weights are equal (Corollary 2(b)); and that, under exact level capacities, straddle-freeness is equivalent to at most two distinct weights or an extreme capacity at every level, with a constructive certificate and a polynomial test on the original data (Theorem 4 and Proposition 2). We do not claim as new the two-fractional-component property of the levels (Caprara et al., 2000), the fact that level hulls restore the Dantzig–Wolfe bound, which is variable redefinition once the levels are described (Martin, 1987), the matroid structure of shifted families, or the two-coefficient knapsack descriptions that the three-weight results transfer (Weismantel, 1996; Dahl and Foldnes, 2003). Table 1 lists every main result, the setting in which it holds, and how it relates to known results.

Table 1: Main results, the setting in which each holds, and their relation to known results.

Result	Setting	Relation to known results
Straddles are the fractional vertices of a level (Thm. 2a)	every admissible level capacity	follows from the two-fractional-component property (Caprara et al., 2000)
No straddle vertex is absorbed; exact iff straddle-free (Thm. 2b)	every admissible level capacity	new; elementary proof, whose consequences are the contribution; fails for load indexing, exact only for equal weights (Prop. 1, Cor. 2)
Dantzig–Wolfe bound attained; necessity (Cor. 1, Thm. 3)	aggregated coupling; necessity needs a realizable straddle	sufficiency by variable redefinition once exactness is known; necessity new
Classification and polynomial recognition (Thm. 4, Prop. 2)	exact level capacities	sufficiency also by total unimodularity; necessity also derivable, nonconstructively, from Herrmann and Joswig (2008); classification, constructive certificate and recognition new
Two weights under any admissible capacity (Prop. EC.2(c), Cor. 4)	two distinct weights, including $\bar{u} = u$	new; exact test for the formulations that use the original capacity
Load versus cardinality indexing (Prop. 3, Cor. 5)	single block; GAP with an outside agent	new comparison; load indexing equals capacity tightening in GAP
Level hulls restore the Dantzig–Wolfe bound (Thm. 5, Cor. 6)	any coupling through aggregated variables	variable redefinition (Martin, 1987) via the disjunctive argument of Thm. 1; objective-independent, unlike the per-block cuts of Chen et al. (2025)
Principal levels are nested matroids (Thm. 6, Prop. 4)	any level	matroid structure known for shifted families; recognition and use for the bound new
Facet-defining rank-type covers; strengthening (Thm. 7, Cor. 8)	full-dimensional level; rank-type covers only	new for exact-cardinality levels; not a complete facet description; differs from knapsack covers (Balas, 1975)
Unit cardinality edges; lifted three-weight hulls (Thms. EC.2–EC.4, Section EC.4)	two coefficients; three weights with dividing gaps	new edge characterization; transfers Weismantel (1996); arithmetic case equivalent to Dahl and Foldnes (2003) given Thm. EC.3

Organization. Section 2 introduces the model and writes the cardinality-indexed relaxation as a disjunctive hull. Section 3 characterizes exactness through straddles and derives the consequences for the Dantzig–Wolfe bound. Section 4 classifies the straddle-free levels and gives the recognition procedure, and Section 5 compares cardinality with load indexing. Section 6 treats resources with straddles through the integer hulls of their levels and ends with the consequences for formulation choice (Section 6.4). Section 7 reports the computational study, and Section 8 concludes. Proofs that are omitted or sketched are given in full in the electronic companion, together with further results and experimental details; results and sections numbered with the prefix EC refer to it.

2 Model and the Cardinality-Indexed Relaxation

Let K be a set of items and R a set of *blocks*, each modeling one resource: an agent in generalized assignment (GAP), a concentrator in concentrator location, an arc in network design. Block r has a set $K_r \subseteq K$ of eligible items with integer weights $0 < d_{rk} \le u_r$ and an integer capacity u_r. The binary variable x_{rk} indicates that item k uses block r, and y_r that block r is installed. For $\chi \in \mathbb{R}^{K_r}$ write $|\chi| = \sum_k \chi_k$ and $d_r(\chi) = \sum_k d_{rk}\chi_k$, and identify a set $T \subseteq K_r$ with its incidence vector $\mathbf{1}_T$. The feasible set of block r is $F_r = \{(\chi,\eta) \in \{0,1\}^{K_r} \times \{0,1\} : d_r(\chi) \le u_r\eta,\ \chi \le \eta\mathbf{1}\}$, and $F_r^\circ = F_r \setminus \{(\mathbf{0},1)\}$. The blocks are coupled by a polyhedron Γ in the space of (x,y,w) that involves the block variables only through (x,y), where w collects auxiliary variables such as flows, so the model is

$$\min\{c^\top x + f^\top y + \xi^\top w :\ (x,y,w) \in \Gamma,\ (x_r,y_r) \in F_r\ \forall r \in R\}, \qquad f \ge 0. \tag{C}$$

GAP and concentrator location couple the blocks by assignment equations, network design by flow conservation, and capacitated Steiner arborescence adds in-degree and node-balancing constraints.

Assumption 1. Every feasible solution of (C) can be modified, without increasing its cost, so that every installed block serves at least one item.

The assumption holds whenever Γ does not involve y, since an empty installed block can be closed ($f \ge 0$); Lemma EC.1 in Section EC.2 verifies it for capacitated Steiner arborescence.

Let $L_r = \max\{|T| : T \subseteq K_r,\ d_r(T) \le u_r\}$ and $\Lambda_r = \{1,\dots,L_r\}$. For each *level* $\ell \in \Lambda_r$, the *exact level capacity* is $u^*_{r\ell} = \max\{d_r(T) : |T| = \ell,\ d_r(T) \le u_r\}$. A level capacity $\bar u_{r\ell}$ is *admissible* if it is an integer with $u^*_{r\ell} \le \bar u_{r\ell} \le u_r$. The studies on GAP, concentrator location and Steiner arborescence use $\bar u = u$, and the network design study uses $\bar u = u^*$. The *level polytope* of block r is

$$B_r = \Big\{(x_{r\cdot},y_{r\cdot}) \ge 0 :\ d_r(x_{r\cdot\ell}) \le \bar u_{r\ell} y_{r\ell},\ |x_{r\cdot\ell}| = \ell\, y_{r\ell},\ x_{rk\ell} \le y_{r\ell}\ \ (\ell \in \Lambda_r,\ k \in K_r),\ \textstyle\sum_\ell y_{r\ell} \le 1\Big\}.$$

The *cardinality-indexed formulation* (CI) replaces each constraint $(x_r,y_r) \in F_r$ by $(x_{r\cdot},y_{r\cdot}) \in B_r$ with binary components, and replaces (x,y) in Γ and in the objective by its image under the aggregation map π, $x_{rk} = \sum_\ell x_{rk\ell}$ and $y_r = \sum_\ell y_{r\ell}$. Under Assumption 1, (C) and (CI) have the

same optimal value z^*, since a block serving a set T is lifted to level $|T|$ and $d_r(T) \le u^*_{r|T|} \le \bar{u}_{r|T|}$. We write $\mathrm{LP}(\cdot)$ and $z_{\mathrm{LP}}(\cdot)$ for the linear relaxation of a formulation and its value. With $Q_r = \pi(B_r)$, since Γ sees the level variables only through π,

$$\pi(\mathrm{LP}(\mathrm{CI})) = \{(x, y, w) \in \Gamma : \ (x_r, y_r) \in Q_r \ \forall r\}. \tag{1}$$

The Dantzig–Wolfe (DW) bound z_{DW} is the minimum of the objective over $\{(x, y, w) \in \Gamma : (x_r, y_r) \in \mathrm{conv}(F_r^\circ) \ \forall r\}$; when Γ involves only (x, y), it is the Lagrangian bound obtained by dualizing Γ (Geoffrion, 1974). We use F_r° rather than F_r because (CI) cannot represent an installed empty block; when Γ does not involve y, the two definitions give the same bound.

We now fix a block and drop the index r. For $\ell \in \Lambda$ let

$$P_\ell = \{\chi \in [0,1]^K : \ |\chi| = \ell, \ d(\chi) \le \bar{u}_\ell\}, \qquad X_\ell = P_\ell \cap \{0,1\}^K.$$

Since $u^*_\ell \le \bar{u}_\ell \le u$, the set X_ℓ consists of the feasible loads with ℓ items, and P_ℓ is nonempty because the ℓ lightest items fit. We call $Q = \pi(B)$ the *cardinality relaxation* of the block.

Theorem 1. *$Q = \mathrm{conv}\big(\{(\mathbf{0}, 0)\} \cup \bigcup_{\ell\in\Lambda} P_\ell \times \{1\}\big)$, and every vertex of B has $y \in \{0, e_1, \dots, e_L\}$. Equivalently, Q is the disjunctive hull of the compact block relaxation $B^C = \{(\chi, \eta) \in [0,1]^{K+1} : d(\chi) \le u\eta, \ \chi \le \eta\mathbf{1}\}$ with respect to the disjunction $\{y = 0\} \vee \bigvee_\ell \{y = 1, \ |x| = \ell, \ d(x) \le \bar{u}_\ell\}$.*

Proof. For fixed ℓ, the constraints of B on $(x_{\cdot\ell}, y_\ell)$ homogenize the description of the polytope P_ℓ, so they define $\{(x_{\cdot\ell}, y_\ell) : y_\ell \ge 0, \ x_{\cdot\ell} \in y_\ell P_\ell\}$. With $\lambda_\ell = y_\ell$ and $\lambda_0 = 1 - \sum_\ell y_\ell$, every point of B is a convex combination of 0 and of points of the sets $S_\ell = \{y = e_\ell, \ x_{\cdot\ell} \in P_\ell, \ x_{\cdot\ell'} = 0 \ (\ell' \ne \ell)\}$, and conversely, because these sets occupy disjoint coordinates. Applying π gives the formula. □

Consequently, disjunctive cuts derived from the compact relaxation with the setup split or with any split on the count are implied by Q, and splits on a level indicator give no cut for a single block (Proposition EC.1 in Section EC.1.1).

3 Straddles and Exactness

Definition 1. A *straddle at level ℓ* is a triple (A, i, j) with $A \subseteq K$, $|A| = \ell - 1$, and $i, j \notin A$, such that $d_j < \bar{u}_\ell - d(A) < d_i$. Its *vertex* is $v^{Aij} = \mathbf{1}_A + t_{Aij} e_i + (1 - t_{Aij}) e_j$, where $t_{Aij} = (\bar{u}_\ell - d(A) - d_j)/(d_i - d_j) \in (0, 1)$. A block is *straddle-free* if it has no straddle at any level.

The quantity $\bar{u}_\ell - d(A)$ is the capacity left for one more item once A has been chosen; a straddle is a choice of A for which this residual separates the weights of two remaining items.

Example 1. Take weights $d = (1, 2, 3)$, capacity $u = 4$ and level 2. The feasible pairs are $\{1, 2\}$ and $\{1, 3\}$, so $u^*_2 = 4$. With $A = \{2\}$, the residual 2 lies strictly between $d_1 = 1$ and $d_3 = 3$, so $(\{2\}, 3, 1)$ is a straddle with vertex $v = (\frac{1}{2}, 1, \frac{1}{2})$, which has two items and weight 4. It lies in

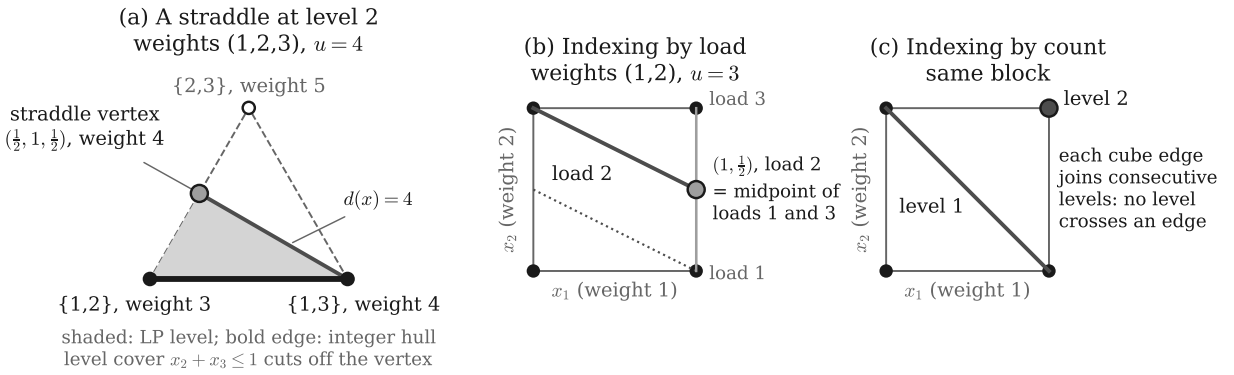


Figure 1: (a) Level 2 of a block with weights $(1, 2, 3)$ and $u = 4$ (Example 1). The shaded part of the triangle is the LP level P_2; the bold edge is its integer hull $\mathrm{conv}(X_2)$. The capacity line $d(x) = 4$ cuts the edge from $\{1, 2\}$ to $\{2, 3\}$ at the straddle vertex $(\frac{1}{2}, 1, \frac{1}{2})$, which the level cover $x_2 + x_3 \le 1$ cuts off. (b) Indexing by load, weights $(1, 2)$, $u = 3$: the load-2 point $(1, \frac{1}{2})$ is the midpoint of the load-1 and load-3 vertices, so load indexing absorbs it. (c) Indexing by count for the same block: each cube edge joins consecutive levels, so a fractional vertex of one level is not a convex combination of points of other levels (Proposition 1).

P_2 but not in $\mathrm{conv}(X_2)$, the segment between $(1, 1, 0)$ and $(1, 0, 1)$, and it violates the level cover $x_2 + x_3 \le 1$ of the infeasible pair $\{2, 3\}$ by $\frac{1}{2}$.

Figure 1(a) draws this level; panels (b) and (c) show the example of the introduction, in which load indexing absorbs a fractional point and cardinality indexing cannot.

Theorem 2. *For every admissible $\bar{u}$:*

(a) every vertex of P_ℓ is integral or has exactly two fractional components, and the fractional vertices of P_ℓ are exactly the straddle vertices v^{Aij} at level ℓ, all of which satisfy $d(v^{Aij}) = \bar{u}_\ell$;

(b) the vertices of Q are $(\mathbf{0}, 0)$, the points $(\mathbf{1}_T, 1)$ for the nonempty feasible loads T, and the points $(v^{Aij}, 1)$ for the straddles. In particular, $Q = \mathrm{conv}(F^\circ)$ if and only if the block is straddle-free.

Proof. (a) A vertex v of P_ℓ satisfies $|K|$ linearly independent constraints with equality, at most two of which are not bounds, so v has at most two fractional components, and not exactly one because $|v| = \ell$ is an integer. If v_i and v_j are fractional, then $|\chi| = \ell$ and $d(\chi) \le \bar{u}_\ell$ are tight and independent on $\{i, j\}$, so $d_i \ne d_j$, say $d_i > d_j$. With $A = \{k : v_k = 1\}$, $v_i + v_j = \ell - |A|$ is an integer in $(0, 2)$, so $v_i + v_j = 1$ and $|A| = \ell - 1$, and $d(A) + v_i d_i + (1 - v_i) d_j = \bar{u}_\ell$ shows that (A, i, j) is a straddle with $t_{Aij} = v_i$. Conversely, $v^{Aij} \in P_\ell$ is the unique solution of the bounds on $K \setminus \{i, j\}$ and the two tight constraints.

(b) By Theorem 1, every vertex of Q is $(\mathbf{0}, 0)$ or $(v, 1)$ with v a vertex of some P_ℓ, and the integral ones are vertices of the cube. Let (A, i, j) be a straddle at level ℓ, $v = v^{Aij}$, and suppose $(v, 1) = \sum_s \mu_s (p_s, 1)$ with $\mu_s > 0$, $\sum_s \mu_s = 1$ and $p_s \in P_{\ell_s}$. The coordinates of v equal to 0 or 1 force $p_s = \mathbf{1}_A + \alpha_s e_i + \beta_s e_j$ with $\alpha_s, \beta_s \in [0, 1]$, and $\alpha_s + \beta_s = \ell_s - \ell + 1$ is an integer. The ℓ-set $A \cup \{i\}$ has weight $d(A) + d_i > \bar{u}_\ell \ge u^*_\ell$, so it is infeasible (every feasible ℓ-set has weight at most u^*_ℓ), that is, $d(A \cup \{i\}) > u$. Hence $\alpha_s + \beta_s = 2$ is impossible, since $\mathbf{1}_{A \cup \{i,j\}}$ has weight above

$u \geq \bar{u}_{\ell+1}$ and lies in no level. Thus $\alpha_s + \beta_s \leq 1$, and $\sum_s \mu_s(\alpha_s + \beta_s) = v_i + v_j = 1$ forces $p_s \in P_\ell$ for all s; since v is a vertex of P_ℓ, every $p_s = v$. A fractional vertex of Q is not in $\operatorname{conv}(F^\circ) \subseteq Q$, which gives the last statement. □

Corollary 1. *For every admissible $\bar{u}$, $z_{\mathrm{LP}}(\mathrm{C}) \leq z_{\mathrm{LP}}(\mathrm{CI}) \leq z_{\mathrm{DW}}$, with $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$ if every block is straddle-free. Under Assumption 1, $z_{\mathrm{DW}} \leq z^*$.*

Proof. Since $P_\ell \times \{1\} \subseteq B^C$ and $X_\ell \subseteq P_\ell$, Theorem 1 gives $\operatorname{conv}(F_r^\circ) \subseteq Q_r \subseteq B_r^C$, and (1) gives the inequalities, with equality when $Q_r = \operatorname{conv}(F_r^\circ)$ for all r. Under Assumption 1, some optimal solution of (C) has $(x_r, y_r) \in F_r^\circ$ for all r and lies in the region defining z_{DW}. □

A straddle shows that $Q_r \neq \operatorname{conv}(F_r^\circ)$, not that the bound is lost: the coupling may make the straddle vertex unreachable, and the objective may make it irrelevant. The next result shows that every straddle the coupling can complete is exposed by some objective.

Theorem 3. *Suppose that Γ involves the blocks only through (x, y), and let (A, i, j) be a straddle of block r. Call it* realizable *if Γ contains two points in which every block $s \neq r$ lies in $\operatorname{conv}(F_s^\circ)$ and block r takes the values $(\mathbf{1}_{A\cup\{i\}}, 1)$ and $(\mathbf{1}_{A\cup\{j\}}, 1)$, respectively. If the straddle is realizable, some cost vector that vanishes outside x_r gives $z_{\mathrm{LP}}(\mathrm{CI}) < z_{\mathrm{DW}}$. Hence, if every block with a straddle has a realizable one, $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$ for every cost vector if and only if every block is straddle-free.*

Proof. Let $v = v^{Aij}$, $t = t_{Aij}$, and let $W = \operatorname{conv}(\{\mathbf{0}\} \cup \{\mathbf{1}_T : T \text{ a feasible load of } r\})$, the projection of $\operatorname{conv}(F_r^\circ)$ onto x_r. A straddle has $\ell \geq 2$, so $A \neq \emptyset$; in a convex combination of points of W equal to v, the coordinates in A, which equal 1, force every point to be $\mathbf{1}_T$ with $T \supseteq A$, and the coordinates outside $A \cup \{i, j\}$, which equal 0, force $T \subseteq A \cup \{i, j\}$. Since $v_i = t > 0$, some such T contains i, so $T = A \cup \{i\}$ or $T = A \cup \{i, j\}$. Neither is a feasible load: as in the proof of Theorem 2(b), $d(A \cup \{i\}) > \bar{u}_\ell \geq u^*_\ell$, so $A \cup \{i\}$ is infeasible and $d(A \cup \{i\}) > u$, and then $d(A \cup \{i, j\}) > u$ as well. So $v \notin W$, and some c_r has $c_r^\top v < \min_{x\in W} c_r^\top x$. Give x_r the cost c_r and all other variables cost 0. If $z^i, z^j \in \Gamma$ are the two points of the definition, $z = tz^i + (1-t)z^j \in \Gamma$ has block r at $(v, 1) \in Q_r$ and every other block in $\operatorname{conv}(F_s^\circ) \subseteq Q_s$, so $z \in \pi(\mathrm{LP(CI)})$ by (1) and $z_{\mathrm{LP}}(\mathrm{CI}) \leq c_r^\top v$. Every point of the region defining z_{DW} has $x_r \in W$, and that region contains z^j, so $z_{\mathrm{DW}} \geq \min_W c_r^\top x > c_r^\top v$. □

For GAP, a straddle of agent r is realizable whenever the remaining jobs can be served by the other agents both when r serves $A \cup \{i\}$ and when it serves $A \cup \{j\}$, in particular when an outside agent exists that can serve every job at once without a fixed cost, as in Section EC.1.8; an uncapacitated block is itself straddle-free (Theorem 4(i) below). Shifting all assignment costs of a job by a constant changes both bounds equally, so the costs can be taken nonnegative. For a straddle (A, i, j), the *level cover* $\sum_{k\in A\cup\{i\}} x_{k\ell} \leq (\ell - 1)y_\ell$ is valid and violated at the straddle vertex by t_{Aij}; level covers are the ℓ-item cases of the cardinality-based generalizations of cover inequalities in Murthy and Ransbotham (2019); Di Francesco et al. (2021); Contu et al. (2026).

Realizability depends on the coupling and not on the costs. When Γ does not involve y, as in GAP, concentrator location and network design, a sufficient test is that each endpoint can be

completed to a point of Γ in which block r serves $A \cup \{i\}$, respectively $A \cup \{j\}$, and every other resource takes an integer feasible configuration; the capacity of block r is deliberately not imposed, since $A \cup \{i\}$ exceeds it, and closing the empty installed blocks of the other resources puts them in F_s°. In GAP and concentrator location this is a feasibility question for assigning the remaining items, and in network design for routing the remaining commodities. Table 2 states what Corollary 1 and Theorem 3 give for each application.

Table 2: Dantzig–Wolfe conclusions for the four cardinality-indexed formulations. $\bar{u}$: level capacity used in the published formulation. "Straddle-free" refers to that capacity; under u^* more blocks are straddle-free (Proposition EC.2). Full details are in Table EC.1.

Formulation	$\bar{u}$	Sufficient for $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$	A straddle yields costs with $z_{\mathrm{LP}}(\mathrm{CI}) < z_{\mathrm{DW}}$ (Thm. 3)
GAP with an uncapacitated outside agent	u	every agent straddle-free (Cor. 1)	always; straddle-freeness is necessary and sufficient (Cor. 1 and Thm. 3)
GAP without an outside agent	u	as above	when the other jobs can be assigned to the other agents both if the agent serves $A \cup \{i\}$ and if it serves $A \cup \{j\}$
Concentrator location (Di Francesco et al., 2021)	u	every concentrator straddle-free	as for GAP; always if an uncapacitated concentrator exists
Unsplittable network design (Anonymous, 2026a)	u^*	every arc straddle-free	when the other commodities can be routed in both completions; no general characterization
Capacitated Steiner arborescence (Contu et al., 2026)	u	not a bound statement: arcs into Steiner nodes project into Q_r, and are exact if straddle-free (Section EC.2)	no statement; the node-copy formulation couples levels across arcs

3.1 Why the Count: a Criterion for Non-Absorption

Theorem 2(b) is an instance of a simple criterion for disjunctions over the values of an integer linear form that keep only a relaxation of each term. The criterion is elementary, and its converse is limited to one-dimensional cube faces; we include it because it locates the difference between the count and the load in one property of the cube.

Proposition 1. *Let $a \in \mathbb{Z}^K$, let $C \subset \mathbb{Z}$ be finite, and for $c \in C$ let $P_c \subseteq [0,1]^K \cap \{x : a^\top x = c\}$ be a polytope. Let $Q = \mathrm{conv}(\{(\mathbf{0}, 0)\} \cup \bigcup_{c \in C} P_c \times \{1\})$, let v be a vertex of P_c, and let Φ be the smallest face of $[0,1]^K$ containing v.*

(a) If all points of $\bigcup_{c' \neq c} P_{c'} \cap \Phi$ satisfy $a^\top x < c$, or all satisfy $a^\top x > c$, then $(v, 1)$ is a vertex of Q.

(b) If $\dim \Phi = 1$, the converse holds.

Proof. (a) In a proper convex combination $(v, 1) = \sum_s \mu_s (p_s, 1)$ with $p_s \in P_{c_s}$, every coordinate of v in $\{0, 1\}$ forces the same value in every p_s, so $p_s \in \Phi$; then $c = \sum_s \mu_s c_s$ forces $c_s = c$ for all s under the hypothesis, and v being a vertex of P_c forces $p_s = v$. (b) If Φ is a segment containing v

in its relative interior and the hypothesis fails, then $a^\top x$ is strictly monotone on Φ, which contains points $q^- \in P_{c^-}$ and $q^+ \in P_{c^+}$ with $c^- < c < c^+$ on opposite sides of v, and $(v, 1)$ is a proper convex combination of $(q^-, 1)$ and $(q^+, 1)$. □

For cardinality indexing, $a = \mathbf{1}$. Because the count changes by one along every cube edge, the face Φ at a straddle vertex meets other levels only at $\mathbf{1}_A$ and $\mathbf{1}_{A\cup\{i,j\}}$, and the latter is infeasible, so all such points lie on one side. For load indexing, $a = d$ and the load changes by an item weight along each cube edge; every fractional vertex of a load term lies on a cube edge, so the criterion is exact.

Corollary 2. *Let $\mathcal{V}$ be a finite set of positive integers, $H_q = \{\chi \in [0,1]^K : d(\chi) = q\}$, and $Q^{\mathrm{load}}_{\mathcal{V}} = \mathrm{conv}(\{(\mathbf{0}, 0)\} \cup \bigcup_{q\in\mathcal{V}} H_q \times \{1\})$.*

(a) The fractional vertices of H_q are the points $\mathbf{1}_A + te_k$ with $d(A) < q < d(A) + d_k$, and such a point is a vertex of $Q^{\mathrm{load}}_{\mathcal{V}}$ if and only if the elements of $\mathcal{V} \cap [d(A), d(A) + d_k]$ are all at least q or all at most q.

(b) For the set $\mathcal{V}^*$ of attainable loads, $Q^{\mathrm{load}}_{\mathcal{V}^*} = \mathrm{conv}(F^\circ)$ if and only if all item weights are equal.

Proof. (a) A vertex of H_q has at most one fractional coordinate k; with A its coordinates equal to 1, Φ is the segment from $\mathbf{1}_A$ to $\mathbf{1}_{A\cup\{k\}}$, and Proposition 1 applies. (b) If all weights equal α, no attainable load lies strictly between $d(A)$ and $d(A) + \alpha$. Otherwise let $d_k > d_{\min}$; then $(d_{\min}/d_k)e_k$ is a fractional vertex of $H_{d_{\min}}$ with $A = \emptyset$, and every element of $\mathcal{V}^*$ is at least $d_{\min}$, so by (a) it survives. □

4 Exactness Under Level Capacities and Recognition

Let $\sigma_m(S)$ and $\tau_m(S)$ denote the sums of the m smallest and of the m largest weights in a set S, and write $\sigma_\ell = \sigma_\ell(K)$ and $\tau_\ell = \tau_\ell(K)$, so that $\sigma_\ell \le u^*_\ell \le \tau_\ell$.

Theorem 4. *Let $\bar{u}_\ell = u^*_\ell$. Level ℓ contains no straddle if and only if one of the following holds: (i) $u^*_\ell = \tau_\ell$, that is, the capacity does not bind at level ℓ; (ii) $u^*_\ell = \sigma_\ell$, that is, only ℓ-sets of minimum weight are feasible; (iii) the weights of the block take at most two distinct values.*

Proof. Write $u^* = u^*_\ell$ and $\mathrm{res}(A) = u^* - d(A)$ for each $(\ell-1)$-set A. There is a straddle with first component A precisely when $\min_{k\notin A} d_k < \mathrm{res}(A) < \max_{k\notin A} d_k$.

Sufficiency. Under (i), $d(A) + d_k \le \tau_\ell = u^*$ for $k \notin A$, so $\mathrm{res}(A) \ge \max_{k\notin A} d_k$; under (ii), symmetrically, $\mathrm{res}(A) \le \min_{k\notin A} d_k$. Under (iii) with weights $\alpha < \beta$, let a feasible ℓ-set of weight u^* contain p items of weight β; if A contains q such items, $\mathrm{res}(A) = (p-q)(\beta-\alpha) + \alpha$, which is at least β if $p > q$ and at most α otherwise.

Necessity. Suppose (i)–(iii) fail. Level 1 satisfies (i), so $\ell \ge 2$. Fix a feasible ℓ-set T with $d(T) = u^*$; here the exactness of the capacity is used, since every residual is now $\mathrm{res}(A) = d(T) - d(A)$. The complement $O = K \setminus T$ is nonempty, since otherwise (i) holds; let $\underline{d}$ and $\overline{d}$ be the smallest and largest weights in O.

Step 1. For $k \in T$, the set $T \setminus \{k\}$ has residual d_k and complement $O \cup \{k\}$, so if $\underline{d} < d_k < \overline{d}$ there is a straddle. Assume therefore that $d_k \le \underline{d}$ or $d_k \ge \overline{d}$ for every $k \in T$.

Step 2. Since (i) fails, not every item of T is at least as heavy as every item of O, so some $k \in T$ and $o \in O$ have $d_k < d_o$, and Step 1 gives $d_k \le \underline{d}$. Since (ii) fails, symmetrically some $k' \in T$ and $o' \in O$ have $d_{k'} > d_{o'}$ and $d_{k'} \ge \overline{d}$. Thus $d_k < d_{k'}$, and $d_k \le d_{o''} \le d_{k'}$ for every $o'' \in O$.

Step 3, Case A: some $o'' \in O$ has $d_k < d_{o''} < d_{k'}$. The $(\ell-1)$-set $A = (T \setminus \{k, k'\}) \cup \{o''\}$ has residual $d_k + d_{k'} - d_{o''} \in (d_k, d_{k'})$, and (A, k', k) is a straddle.

Case B: every $o \in O$ has $d_o \in \{\alpha, \beta\}$, where $\alpha = d_k < \beta = d_{k'}$. The items o, o' of Step 2 give $d_o = \beta$ and $d_{o'} = \alpha$. Since (iii) fails, some $\hat{k} \in T$ has weight $\gamma \notin \{\alpha, \beta\}$, and Step 1 gives $\gamma < \alpha$ or $\gamma > \beta$. If $\gamma < \alpha$, then $A = (T \setminus \{\hat{k}, k'\}) \cup \{o'\}$ has residual $\gamma + \beta - \alpha \in (\gamma, \beta)$, and $(A, k', \hat{k})$ is a straddle. If $\gamma > \beta$, then $A = (T \setminus \{\hat{k}, k\}) \cup \{o\}$ has residual $\gamma + \alpha - \beta \in (\alpha, \gamma)$, and $(A, \hat{k}, k)$ is a straddle. □

Given a maximum-weight feasible ℓ-set T, the proof returns a straddle, and hence a violated level cover, in $O(|K|)$ time; finding T is a cardinality-constrained subset-sum problem, which is NP-hard in general and solvable by dynamic programming in pseudo-polynomial time.

Relation to hypersimplex splits. Geometrically, the hyperplane $\{d(\chi) = u^*_\ell\}$ passes through a vertex of the hypersimplex $\{\chi \in [0,1]^K : |\chi| = \ell\}$ and crosses none of its edges in the relative interior exactly in cases (i)–(iii). By the classification of hypersimplex splits of Herrmann and Joswig (2008), a hyperplane that splits a hypersimplex without crossing an edge has two-valued coefficients on its affine hull, and this gives a second, nonconstructive proof of the necessity part of Theorem 4 (Section EC.1.5). The contribution of Theorem 4 is therefore not new hypersimplex geometry. Its constructive proof, however, provides three results that the split classification does not. First, it returns a certificate, a straddle and a violated level cover, in linear time from one maximum-weight feasible set, so that a failed test yields a cutting plane. Second, it leads to a test on the original data that avoids the exact level capacities (Proposition 2 below). Third, it extends to the capacities $\bar{u} \ne u^*$ that three of the four applied formulations use (Proposition EC.2 in Section EC.1 and Corollary 4 below), to which the split classification does not apply because the hyperplane then passes through no vertex.

Corollary 3. *With exact level capacities, the cardinality relaxation of a block equals its integer hull if and only if its weights take at most two values or every level satisfies (i) or (ii).*

Exactness can be decided without computing u^*.

Proposition 2. *Fix a block and $\ell \in \Lambda$, let L_ℓ be a set of ℓ lightest items, and let $g_\ell = \min\{d_o - d_k : k \in L_\ell,\ o \notin L_\ell,\ d_o > d_k\}$, with $g_\ell = \infty$ if no such pair exists. (a) Condition (i) holds if and only if $\tau_\ell \le u$. (b) Condition (ii) holds if and only if $\sigma_\ell + g_\ell > u$. (c) If the weights take two values $\alpha < \beta$, with n_β items of weight β, then $u^*_\ell = \alpha\ell + (\beta - \alpha)\min\{n_\beta, \ell, \lfloor (u - \alpha\ell)/(\beta - \alpha) \rfloor\}$. Hence straddle-freeness under exact level capacities is decided, and on straddle-free blocks all exact level capacities are computed, in $O(|K| \log |K| + |K|L)$ time.*

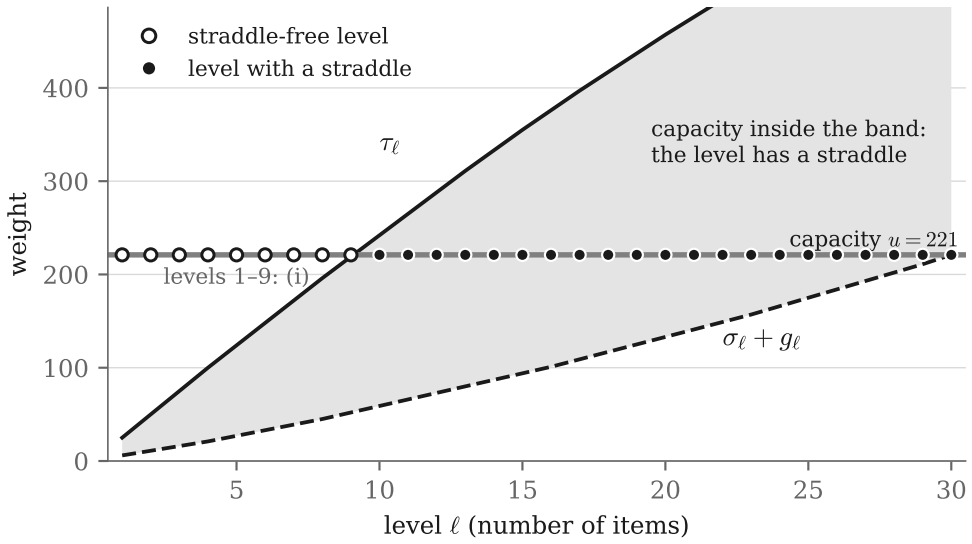


Figure 2: Where straddles occur, for agent 1 of the OR-Library instance c05100 (100 jobs, 5 agents; 21 distinct weights). τ_ℓ is the weight of the ℓ heaviest items, and $\sigma_\ell + g_\ell$ is the smallest weight of an ℓ-set heavier than the ℓ lightest items. By Theorem 4 and Proposition 2, under exact level capacities, a level of a block with three or more distinct weights has a straddle exactly when $\sigma_\ell + g_\ell \le u < \tau_\ell$. Levels 1–9 satisfy $u \ge \tau_\ell$ (condition (i)) and are straddle-free; all levels from 10 to $L = 30$ have straddles.

The proof is in Section EC.1. Figure 2 shows the two curves of Proposition 2 for a benchmark agent, which is straddle-free only at the levels where its capacity does not bind. With other admissible capacities, a level that is straddle-free for an admissible capacity remains straddle-free for every smaller admissible capacity, so a block with a straddle under u^* has one under every admissible capacity; under (i) a level is straddle-free for every admissible capacity, and with two weights $\alpha < \beta$ it is straddle-free if and only if (i) holds or $\bar{u}_\ell - u^*_\ell$ is a multiple of $\beta - \alpha$ (Proposition EC.2 in Section EC.1). In particular, the original capacity $\bar{u} = u$, which the GAP, concentrator location and Steiner arborescence formulations use, can create straddles on two-valued blocks that u^* avoids, and the test is then as follows.

Corollary 4. *Let the weights of a block take two values $\alpha < \beta$, let $\Delta = \beta - \alpha$, and let $\bar{u} = u$. The block is straddle-free if and only if Δ divides $u - \alpha\ell$ for every level $\ell \in \Lambda$ with $\tau_\ell > u$. If at least two levels have $\tau_\ell > u$, this holds if and only if Δ divides both α and u, that is, the two weights are consecutive multiples of Δ and u is a multiple of Δ.*

The proof is in Section EC.1. On blocks with two weights that are not consecutive multiples of their difference, such as weights 5 and 8, the original capacity therefore loses exactness as soon as two levels bind, whereas u^* keeps it; Section 7.2 shows the effect on the bound.

Items that belong to no feasible ℓ-set, or to every one, can be fixed at bound in $O(|K| \log |K|)$ time, and the same tests then apply to the reduced level, which can reveal exactness that the full level hides (Corollary EC.1 in Section EC.1.4). Although at most two distinct weights per resource guarantee exactness under exact level capacities, this does not make the integer problem

tractable: GAP remains strongly NP-hard when every agent sees at most two distinct weights (Proposition EC.3 in Section EC.1.6).

5 Cardinality Versus Load Indexing

The most common alternative disaggregation indexes the variables of a resource by its load (Gouveia, 1995; Correia et al., 2010): binary variables z_q for q in a set $\mathcal{V}$ of positive integers, with copies satisfying $d(x_{\cdot q}) = q\, z_q$, $x_{kq} \le z_q$ and $\sum_q z_q \le 1$. As in Theorem 1, its projection is the set $Q^{\text{load}}_{\mathcal{V}}$ of Corollary 2. It represents every feasible load only if $\mathcal{V}$ contains the set $\mathcal{V}^*$ of attainable loads. Write $Q(\bar{u})$ for the cardinality relaxation with level capacities $\bar{u}$.

Proposition 3. *For every block and every $\mathcal{V} \supseteq \mathcal{V}^*$, $Q(u^*) \subseteq Q^{\text{load}}_{\mathcal{V}}$, and*

$$\operatorname{conv}(F^\circ) \;\subseteq\; Q(u^*) \;\subseteq\; Q^{\text{load}}_{\mathcal{V}^*} \;\subsetneq\; B^C.$$

The first inclusion is an equality if and only if the block is straddle-free. The second is an equality if and only if all weights are equal; otherwise, with $d_{\min}$ the smallest weight and $d_k > d_{\min}$, the point $((d_{\min}/d_k)e_k, 1)$ lies in $Q^{\text{load}}_{\mathcal{V}^*} \setminus Q(u^*)$. The point $(\mathbf{0}, 1)$ lies in $B^C \setminus Q^{\text{load}}_{\mathcal{V}^*}$.

Proof sketch. By Theorem 1 it suffices that $P_\ell(u^*) \subseteq \operatorname{conv}(\bigcup_{q\in\mathcal{V}} H_q)$, which holds because every vertex of $P_\ell(u^*)$ has an attainable weight by Theorem 2(a). The equality conditions are Theorem 2(b), Theorem 4 and Corollary 2; the witness violates $|x| \ge y$, which holds on $Q(u^*)$ (Section EC.1). □

Under aggregated coupling, cardinality indexing with exact level capacities is thus at least as strong as any valid load indexing of the same resources; formulations with load- or level-specific coupling across resources are not covered. For GAP with an outside agent the comparison is exact (Theorem EC.1 in Section EC.1.8). Let z_{load} be the bound of load indexing over the attainable loads and $\hat{u}_r = \max \mathcal{V}^*_r$.

Corollary 5. *Consider GAP with an uncapacitated outside agent without a fixed cost, as in Section EC.1.8, and exact level capacities. For every cost vector, $z_{\text{LP}}(\text{C}) \le z_{\text{load}} \le z_{\text{LP}}(\text{CI}) \le z_{\text{DW}}$. Moreover: (i) $z_{\text{LP}}(\text{CI}) = z_{\text{DW}}$ for every cost vector if and only if every agent is straddle-free; (ii) z_{load} equals the bound of the compact formulation with each capacity u_r lowered to $\hat{u}_r$, also without the outside agent; (iii) $z_{\text{load}} = z_{\text{LP}}(\text{CI})$ for every cost vector if and only if, for every agent, all weights are equal or $u_r \ge d_r(K_r)$.*

Part (ii) says that in GAP, indexing by the attainable loads is no stronger than tightening the capacities, whereas cardinality indexing attains the Dantzig–Wolfe bound on every agent with at most two distinct weights.

6 Resources With Straddles: Level Hulls and the Dantzig–Wolfe Bound

When a block has straddles, $Q \neq \operatorname{conv}(F^\circ)$. The disjunctive argument of Theorem 1 shows exactly what is missing.

Theorem 5. *For a block and each $\ell \in \Lambda$, let $G_\ell \chi \leq \gamma_\ell$ be any linear system with $\{\chi \in [0,1]^K : |\chi| = \ell,\ G_\ell \chi \leq \gamma_\ell\} = \operatorname{conv}(X_\ell)$, and let B^H be B with the inequalities $G_\ell x_{\cdot\ell} \leq \gamma_\ell y_\ell$ added for every ℓ. Then $\pi(B^H) = \operatorname{conv}(F^\circ)$. Consequently, if these inequalities are added to* LP(CI) *for every block, the resulting bound equals z_{DW}, and it suffices to add them on the levels that contain a straddle.*

Proof. For each ℓ, the constraints of B^H on $(x_{\cdot\ell}, y_\ell)$ define $\{(x_{\cdot\ell}, y_\ell) : y_\ell \geq 0,\ x_{\cdot\ell} \in y_\ell \operatorname{conv}(X_\ell)\}$: for $y_\ell > 0$ they state that $x_{\cdot\ell}/y_\ell$ lies in the system that describes $\operatorname{conv}(X_\ell)$, and for $y_\ell = 0$ the constraints $0 \leq x_{k\ell} \leq y_\ell$ force $x_{\cdot\ell} = 0$. The argument of Theorem 1 gives $\pi(B^H) = \operatorname{conv}(\{(\mathbf{0}, 0)\} \cup \bigcup_\ell \operatorname{conv}(X_\ell) \times \{1\}) = \operatorname{conv}(F^\circ)$, and (1) applied to $\pi(B^H)$ gives z_{DW}. On a straddle-free level, $P_\ell = \operatorname{conv}(X_\ell)$ by Theorem 2(a). □

The theorem reduces the gap between $z_{\mathrm{LP}}(\mathrm{CI})$ and z_{DW} to the integer hull of single levels. Separation over $\operatorname{conv}(X_\ell)$ at a point $(x_{\cdot\ell}, y_\ell)$ with $y_\ell > 0$ is separation of $x_{\cdot\ell}/y_\ell$, so polynomial separation for every level gives z_{DW} in polynomial time by the equivalence of separation and optimization. For a fixed number q of distinct weights, optimization over X_ℓ is polynomial, by enumerating the $O(\ell^{q-1})$ count vectors and choosing the cheapest items in each weight class, so z_{DW} is then computable in polynomial time by the ellipsoid method applied to the Dantzig–Wolfe dual, with pricing as the separation oracle; for two distinct weights, pricing reduces to enumerating the number of heavy items. The value of the results below is therefore not the complexity of the bound but explicit inequality families in the space of (CI), which can be used as cutting planes in branch-and-cut, together with a class of blocks on which no separation is needed at all. This also distinguishes them from the block cuts of Chen et al. (2025), which recover z_{DW} with one cut per block but only for the objective whose dual multipliers generated them.

6.1 Principal Levels

Let $w_1 > \cdots > w_q$ be the distinct weights of a block, K_p the items of weight w_p, $P_p = K_1 \cup \cdots \cup K_p$, and $C_p(S) = |S \cap P_p|$. For ℓ-sets, summation by parts gives

$$d(S) = \sum_{p=1}^{q-1} (w_p - w_{p+1})\, C_p(S) + w_q\, \ell, \tag{2}$$

so an ℓ-set whose prefix counts are dominated by those of a feasible ℓ-set is feasible. We call level ℓ *principal* if some $T^* \in X_\ell$ has $C_p(T^*) \geq C_p(S)$ for all $p < q$ and all $S \in X_\ell$.

Theorem 6. *Let $X_\ell \neq \emptyset$. (a) If level ℓ is principal with dominating set T^*, then $\operatorname{conv}(X_\ell) = \{x \in [0,1]^K : |x| = \ell,\ x(P_p) \leq C_p(T^*)\ (p = 1, \ldots, q-1)\}$, a totally unimodular system, and X_ℓ is the*

family of bases of a nested matroid. (b) Conversely, if X_ℓ is the family of bases of a matroid, then level ℓ is principal.

Proof. (a) With the items ordered by class, every row of the system is the incidence vector of an initial segment, so the matrix is an interval matrix and totally unimodular, and the right-hand sides are integers. Its integer points are the ℓ-sets dominated by T^*, which are feasible by (2); conversely, every feasible ℓ-set is dominated by T^*. The sets I with $|I \cap P_p| \le C_p(T^*)$ and $|I| \le \ell$ are the independent sets of a laminar matroid of rank ℓ whose bases are X_ℓ. (b) Order the items by nonincreasing weight, ties broken arbitrarily, so that every P_p is an initial segment. The basis T produced by the greedy algorithm in this order contains a basis of every initial segment I, so $|T \cap I| = r(I) \ge |S \cap I|$ for every basis S, where r is the rank function (this is Gale's optimality of the greedy basis; Gale, 1968). With $I = P_p$, T dominates every member of X_ℓ. □

Proposition 4. *For $p < q$, let $m_p = \max\{j :\ \sigma_j(P_p) + \sigma_{\ell-j}(K \setminus P_p) \le u,\ j \le \min\{\ell, |P_p|\},\ \ell - j \le |K \setminus P_p|\}$, the largest number of items from P_p in a feasible ℓ-set. Level ℓ is principal if and only if $0 \le m_p - m_{p-1} \le |K_p|$ for $p = 1, \dots, q$, with $m_0 = 0$ and $m_q = \ell$, and $\sum_p w_p(m_p - m_{p-1}) \le u$. Hence principal levels, and their descriptions in Theorem 6, are found in $O(q\,\ell\,|K|)$ time after sorting.*

Proof. The lightest ℓ-set with exactly j items of P_p consists of the j lightest items of P_p and the $\ell - j$ lightest of $K \setminus P_p$, so m_p is as stated. A dominating set T^* must have $C_p(T^*) = m_p$ for all p, so the level is principal if and only if some feasible ℓ-set has prefix counts $m_1, \dots, m_{q-1}$, that is, $m_p - m_{p-1}$ items of class p; the stated conditions say that such a set exists and fits. □

Corollary 6. *If every level of every block is principal, then* LP(CI) *with the at most $q_r - 1$ inequalities $\sum_{k \in P_{rp}} x_{rk\ell} \le C_p(T^*_{r\ell})\, y_{r\ell}$ added on every level ℓ of every block r is a linear program whose value is z_{DW} and whose size is polynomial in $|R|$ and $|K|$ and independent of the capacities. More generally, adding these inequalities on the principal levels and the hull of Theorem 5 on the others gives z_{DW}.*

Here P_{rp} is the set P_p of block r and $T^*_{r\ell}$ its dominating set. The corollary follows from Theorems 5 and 6 and Proposition 4. Every straddle-free level is principal: a level that is straddle-free for some admissible capacity is straddle-free for u^* (Proposition EC.2), hence satisfies (i), (ii) or (iii) of Theorem 4; under (i) the ℓ heaviest items dominate, under (ii) all feasible sets have the same prefix counts, and under (iii) a count vector is determined by its first entry. Principal levels need not be straddle-free, however: in Example 1, the feasible pairs $\{1,2\}$ and $\{1,3\}$ are dominated by $\{1,3\}$, so the level is principal, and its straddle is removed by the single inequality $x_3 + x_2 \le 1$, which is the level cover. With three or more weights a level need not be principal: for weights $(1,2,2,3)$, $u = 4$ and $\ell = 2$, the pairs of weights $\{1,3\}$ and $\{2,2\}$ have incomparable prefix counts. In an exhaustive census of small blocks (Section EC.7), 92% of the full-dimensional levels with three distinct weights are principal, and 64% with seven.

6.2 Facets of a Level and Strengthened Covers

On a non-principal level, X_ℓ is not a matroid, and we characterize the rank-type covers $x(H) \le h$ that define facets of $\mathrm{conv}(X_\ell)$. Throughout, $n = |K|$, $1 \le \ell \le n-1$ and $O = K \setminus H$. The lightest ℓ-set with j items of H weighs $\sigma_j(H) + \sigma_{\ell-j}(O)$, so $x(H) \le h$ is valid if and only if this exceeds u for every feasible $j > h$. The dimension of $\mathrm{conv}(X_\ell)$ is $n-1$ if and only if every item belongs to some, and no item to every, feasible ℓ-set (Lemma EC.2 in Section EC.3).

Theorem 7. *Suppose* $\dim \mathrm{conv}(X_\ell) = n-1$, *and let* $x(H) \le h$ *be valid with* $1 \le h \le \min\{\ell - 1, |H|-1\}$ *and* $\ell - h \le |O|$. *Let* a, b *be a heaviest and a lightest item of* H, *and* a', b' *a heaviest and a lightest item of* O. *Then* $x(H) \le h$ *defines a facet of* $\mathrm{conv}(X_\ell)$ *if and only if*

$$\textit{(H1)}\quad d_a + \sigma_{h-1}(H \setminus \{a\}) + \sigma_{\ell-h}(O) \le u, \qquad \textit{(H2)}\quad \sigma_h(H \setminus \{b\}) + \sigma_{\ell-h}(O) \le u,$$
$$\textit{(O1)}\quad \sigma_h(H) + d_{a'} + \sigma_{\ell-h-1}(O \setminus \{a'\}) \le u, \qquad \textit{(O2)}\quad \sigma_h(H) + \sigma_{\ell-h}(O \setminus \{b'\}) \le u \text{ and } \ell - h < |O|,$$

where (O2) is required only when $|O| \ge 2$.

Conditions (H1) and (O1) say that the heaviest item of each side lies in some feasible set on the face, and (H2) and (O2) that the lightest item of each side can be avoided. Because feasible sets are closed under exchanging an item for a lighter one on the same side, testing the extreme items suffices; the proof is in Section EC.3. Theorem 7 characterizes facets within the family of rank-type covers only. This family does not describe every level: on the non-principal full-dimensional levels with three distinct weights in the census of Section EC.7, it suffices on 56%, and the remaining facets have integer coefficients up to 4, such as $x_2 + x_3 + x_4 + 2x_5 \le 2$ for weights $(1,2,2,2,4)$, $u = 5$ and $\ell = 2$. For a level cover, $H = C$ is an infeasible ℓ-set and $h = \ell - 1$, and the conditions reduce to two.

Corollary 7. *Suppose* $\dim \mathrm{conv}(X_\ell) = n-1$ *and* $\ell \ge 2$, *and let* C *be an infeasible* ℓ*-set. The level cover* $x(C) \le \ell - 1$ *defines a facet if and only if* $d(C) - \min_{c \in C} d_c + \min_{o \notin C} d_o \le u$ *and* $d(C) - \max_{c \in C} d_c + \max_{o \notin C} d_o \le u$. *At a straddle vertex* v^{Aij}, *the only violated level cover is that of* $C = A \cup \{i\}$.

When a condition fails, the face satisfies an additional equation, and for two of the four conditions this yields a stronger cover.

Corollary 8. *Let* $x(H) \le h$ *be valid with* h *in the range of Theorem 7. (a) If (O1) fails,* $x(H \cup \{a'\}) \le h$ *is valid. (b) If (H2) fails,* $x(H \setminus \{b\}) \le h - 1$ *is valid. Each new inequality implies* $x(H) \le h$ *on* $[0,1]^K$ *and is at least as violated at every point of* $[0,1]^K$. *Applying (a) and (b) while possible terminates after at most* $2n$ *steps with a valid cover that leaves the range or satisfies (H2) and (O1).*

Proof. (a) A feasible set containing a' with h items of H weighs at least the left side of (O1), which exceeds u. (b) A feasible set with h items of $H \setminus \{b\}$ has $\ell - h$ items of O and weighs at least the left side of (H2). Domination and violation follow from $x(H \cup \{a'\}) - h = x(H) - h + x_{a'}$ and $x(H \setminus \{b\}) - (h-1) = x(H) - h + (1 - x_b)$. □

6.3 Three Weights

A second mechanism describes further levels exactly. If the count changes by at most one along every edge of an integral polytope, slicing it by integer values of the count creates no new vertices. Section EC.4 characterizes the pairs of coefficients for which every knapsack polytope has this unit-edge property (Theorem EC.2) and deduces that, for three distinct weights one of whose gaps divides the other, every level is the cardinality slice of a two-coefficient knapsack polytope (Theorem EC.3), so the description of Weismantel (1996) transfers to it. For weights in arithmetic progression, a single family of covers suffices and is separated by sorting (Theorem EC.4); this case is equivalent to the case $p = 2$ of Dahl and Foldnes (2003) once Theorem EC.3 is established, and what it adds is a constructive proof and cardinality-aware right-hand sides. With these results, explicit inequality families in the space of (CI) attain z_{DW} on every block whose levels are straddle-free, principal, or have three weights with dividing gaps.

6.4 Consequences for Formulation Choice

The results of Sections 3–6 combine into a procedure that uses only the data of each resource and requires no solution of the model.

(1) *Test every resource for straddles* with Proposition 2, after fixing items at bound (Corollary EC.1). For two-valued resources under the original capacity, Corollary 4 gives the test. If every resource is straddle-free and the coupling is through aggregated variables, $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$ (Corollary 1): the Dantzig–Wolfe bound is available from one linear program whose size does not depend on the capacities, and no column generation or separation is needed.

(2) *Use exact level capacities* whenever the original capacity creates straddles that u^* avoids. They never weaken the relaxation (Proposition EC.2(a)), and on two-valued blocks they are given in closed form by Proposition 2(c).

(3) *On levels with a straddle that are principal*, add the at most $q - 1$ inequalities of Corollary 6. They are computed from the data before optimization, in $O(q\,\ell\,|K|)$ time per level.

(4) *On the remaining levels with a straddle*, the gap to z_{DW} can be closed only by inequalities valid for $\mathrm{conv}(X_\ell)$. Strengthened covers (Corollary 8) are inexpensive to generate and close part of it; exact separation over $\mathrm{conv}(X_\ell)$ closes all of it (Theorem 5) but calls an optimization oracle over the level in every round, and Section 7.4 shows that this becomes the dominant computational cost as instances grow. Where most levels fall in this case, as on the standard GAP benchmarks (Section 7.7), column generation or the objective-specific block cuts of Chen et al. (2025) are the natural alternatives for obtaining z_{DW}.

(5) *To go beyond* z_{DW}, no inequality on a single resource suffices (Theorem 5); inequalities or index relations that couple several resources are required.

The procedure also indicates when the larger size of (CI) is justified. Relative to the compact formulation, (CI) is strongest when resources serve few items whose weights take few distinct values, since

most levels are then straddle-free or principal, and weakest when many items with heterogeneous weights compete for a resource and the light items have the lower costs (Proposition EC.6).

7 Computational Study

The experiments address six questions, each associated with a theoretical result: (E1) blocks that Theorem 4 certifies as straddle-free attain z_{DW} (Corollary 1); (E2) on other blocks, the distance to z_{DW} grows with the share of levels with a straddle and depends on whether the straddles are favorable for the reduced costs; (E3) exact level capacities and fixing items at bound enlarge the class of exact blocks (Proposition EC.2, Corollary EC.1); (E4) covers strengthened by Corollary 8 are facet-defining and close more of the gap than level covers; (E5) adding the integer hulls of the levels with straddles to LP(CI) attains z_{DW} (Theorem 5), and the principal-level inequalities suffice whenever every level with a straddle is principal (Corollary 6); and (E6) on two-valued blocks, the size of (CI) and the recognition of exactness do not depend on the magnitude of the capacities. Section 7.7 applies the recognition results and the three bounds to the standard GAP benchmark instances, where the recognition test shows, without solving any model, that no agent is straddle-free, and (CI) remains well below z_{DW}. The study is not a comparison of solution methods; Section 7.8 relates it to the branch-and-cut studies that use the formulation.

Setup. All linear and integer programs are solved with HiGHS through SciPy 1.18. Optimal values z^* of the instances in Tables 3–7 are proven with a relative gap tolerance of 10^{-6}; since all costs are integers and $z^* \le 4742$, they are exact. DW bounds are computed by column generation with exact knapsack pricing until no column has reduced cost below -10^{-7}, and on every instance $z_{\mathrm{LP}}(\mathrm{C}) \le z_{\mathrm{LP}}(\mathrm{CI}) \le z_{\mathrm{DW}} \le z^*$ holds up to a relative 10^{-6}. For each bound we report the share of $z^* - z_{\mathrm{LP}}(\mathrm{C})$ that it closes, averaged over the instances with $z^* > z_{\mathrm{LP}}(\mathrm{C})$; this excludes 2 of the 10 instances of GAP class W12, on which the compact relaxation is already optimal, and no instance of any other class. The strengthened compact relaxation C^+ adds the cardinality bound $|x_r| \le L_r y_r$ and lowers each capacity to the largest attainable load; for GAP, by Corollary 5(ii), it is at least as strong as load indexing over the attainable loads. Straddles are detected with the strict inequalities of Definition 1. Section EC.7 gives the tolerances, the column generation and the instance generators.

Instances. GAP instances have $m = 5$ agents and $n = 30$ jobs, ten per class. Types C, D and E follow the generation rules of the standard benchmark sets (Chu and Beasley, 1997; Laguna et al., 1995): type C has weights uniform on $\{5, \dots, 25\}$; type D has weights uniform on $\{1, \dots, 100\}$ and costs that decrease with the weight; type E has exponentially distributed weights and costs inversely proportional to them; in all three, capacities are 80% of the average load. Class W12 has weights in $\{1, 2\}$, and classes Ss have weights uniform on $\{10, \dots, 10 + s\}$ and capacities equal to the average load. Network design instances have 10 nodes, 30 to 46 arcs and 12 commodities, ten per class, with arc capacities of about 2.5 times the mean demand, so that an arc carries only a

few commodities. Demands are uniform on $\{1,\dots,20\}$ in class H, lie in $\{5,7,8\}$ in class T3 and in $\{5,8\}$ in class T2, and are uniform on $\{10,\dots,10+s\}$ in classes Ss. The network relaxations use exact level capacities and contain no multi-arc inequalities, so they isolate the arc-level effect.

7.1 Exactness and the Loss Caused by Straddles

Table 3: Generalized assignment ($m = 5$, $n = 30$; ten instances per class). Straddle share: share of (agent, level) pairs with a straddle (%). Gap: $(z^* - z_{\mathrm{LP}}(\mathrm{C}))/z^*$ (%). Columns C$^+$ to DW: share of $z^* - z_{\mathrm{LP}}(\mathrm{C})$ closed (%); CI+LC adds level covers separated exactly. Certified: whether Theorem 4 certifies $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$. Last column: largest relative difference between z_{DW} and $z_{\mathrm{LP}}(\mathrm{CI})$ over the class. Class W12 averages over the 8 instances with $z^* > z_{\mathrm{LP}}(\mathrm{C})$.

Class	Straddle share (%)	Gap	C$^+$	CI	CI+LC	DW	Certified	$\max \lvert z_{\mathrm{DW}} - z_{\mathrm{LP}}(\mathrm{CI})\rvert / z_{\mathrm{DW}}$
C	71	3.49	0.0	21.8	42.1	88.6	no	$> 10^{-3}$
D	81	1.33	0.0	4.9	12.3	78.3	no	$> 10^{-3}$
E	91	1.90	0.0	22.3	30.7	74.1	no	$> 10^{-3}$
W12	0	0.69	0.0	81.7	81.7	81.7	yes	8×10^{-15}
S1	0	2.77	100.0	100.0	100.0	100.0	yes	1×10^{-13}
S2	17	2.92	87.6	87.6	93.7	99.3	no	$> 10^{-3}$
S4	20	3.08	12.1	75.0	89.8	99.7	no	$> 10^{-3}$
S8	41	2.73	0.0	60.2	78.0	97.3	no	$> 10^{-3}$
S16	51	2.12	0.0	57.2	77.4	96.0	no	$> 10^{-3}$

Table 4: Unsplittable network design (10 nodes, 12 commodities; ten instances per class). L: average number of levels per arc. Other columns as in Table 3.

Class	L	Straddle share (%)	Gap	C$^+$	CI	DW	Certified	$\max \lvert z_{\mathrm{DW}} - z_{\mathrm{LP}}(\mathrm{CI})\rvert / z_{\mathrm{DW}}$
H	4.7	74	8.30	0.0	18.4	94.4	no	$> 10^{-3}$
T3	3.0	23	9.12	0.2	56.3	56.9	no	$> 10^{-3}$
T2	3.0	0	8.77	8.1	62.0	62.0	yes	2×10^{-14}
S1	2.0	0	7.80	65.0	65.0	65.0	yes	8×10^{-16}
S2	2.0	0	7.67	65.1	65.1	65.1	yes	6×10^{-16}
S4	2.3	3	7.00	44.1	63.4	63.7	no	$> 10^{-3}$
S8	3.0	40	6.20	0.1	62.0	65.8	no	$> 10^{-3}$
S16	3.2	60	6.54	0.6	42.4	75.4	no	$> 10^{-3}$

(E1) Exactness. Tables 3 and 4 confirm Corollary 1 on every class. In GAP classes W12 and S1 and network classes T2 and S1, every block has at most two distinct weights, and in network class S2 every level satisfies condition (i); on all of these, $z_{\mathrm{LP}}(\mathrm{CI})$ and z_{DW} coincide to within round-off, while on every other class the largest relative difference exceeds 10^{-3}. The two-valued classes are also where the cardinality relaxation is most useful relative to the compact one: on GAP class W12 it closes 81.7% of the gap and C$^+$ none, and on network class T2 it closes 62.0% against 8.1%.

(E2) The loss caused by straddles. GAP types C, D and E have many distinct weights and capacities that are not extreme, so by Corollary 3 they contain straddles: 71–91% of the levels

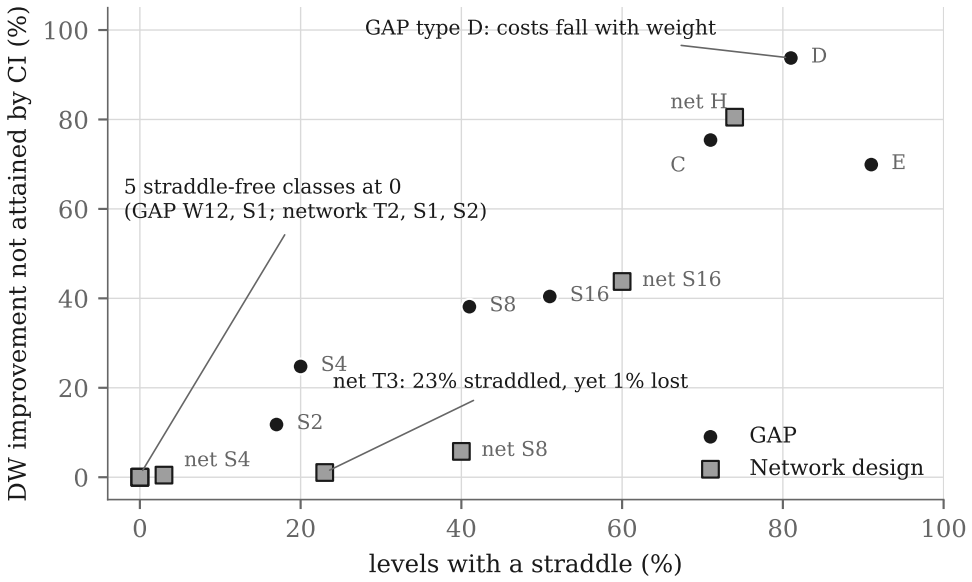


Figure 3: Bound loss against straddle frequency, one point per class (Tables 3 and 4). Horizontal axis: share of (block, level) pairs with a straddle in the class. Vertical axis: share of the DW improvement that CI fails to attain, $1 - \text{CI/DW}$. Circles are GAP classes and squares network classes; the five classes at zero are GAP W12 and S1 and network T2, S1 and S2. Network T3 (23% of levels with a straddle, about 1% lost) and network S8 (40%, 6% lost) lie below the trend, and fixing items at bound removes most of T3's straddles and many of S8's (Table 5). GAP type D, whose costs decrease with the weight, loses the most. Both observations are consistent with the a posteriori bound of Proposition EC.6.

have one, and (CI) closes 5–22% of the gap, against 74–89% for z_{DW}. Along the S series, the share closed by (CI) falls from 100% to 57% for GAP, and the distance to z_{DW} grows from zero to 33 percentage points for network design, as the share of levels with a straddle increases. The number of straddles alone does not determine the loss, however. The loss is smallest relative to z_{DW} on network class T3, where 23% of the levels have a straddle but (CI) is within one percentage point of z_{DW}, and largest on GAP type D, whose costs decrease with the weight. Both observations agree with an a posteriori bound (Proposition EC.6 in Section EC.6): a straddle vertex replaces part of a light item by a heavy one, and it lowers the bound only when the heavy item has the lower reduced cost. That bound is valid on every instance and vanishes exactly on the straddle-free classes, but it is loose, exceeding the observed gap by a median factor between 2.4 and 8.8, and it requires the DW multipliers; the a priori certificate of exactness is straddle-freeness itself. Figure 3 summarizes both observations.

7.2 Level Capacities and Fixing Items at Bound

(E3) Table 5 shows that the choice of level capacity can have a large effect on the bound. On the two-valued GAP class W58, where the weights differ by three, every level is straddle-free with $\bar{u} = u^*$, and $z_{\text{LP}}(\text{CI}) = z_{\text{DW}} = z^*$. With the original capacity, which the GAP study uses, 30.6% of the levels have a straddle and (CI) closes only 23.8% of the gap. Proposition EC.2(c) identifies

Table 5: Level capacities and fixing items at bound (ten instances per class). Straddle share: share (%) of levels with a straddle for the original capacity u, for u^*, and for u^* after fixing items at bound. Other columns: share (%) of $z^* - z_{\mathrm{LP}}(\mathrm{C})$ closed. Class W58 has GAP weights in $\{5, 8\}$ and capacities as in type C.

	Straddle share (%)			CI, $\bar{u} = u$		CI, $\bar{u} = u^*$		
Class	u	u^*	fixed	plain	fixed	plain	fixed	DW
GAP C	70.8	70.6	69.7	21.8	21.8	21.8	21.8	88.6
GAP D	82.1	81.4	79.7	4.9	4.9	4.9	4.9	78.3
GAP E	91.6	91.0	89.1	22.3	22.3	22.3	22.3	74.1
GAP S8	41.2	41.2	40.9	60.2	60.2	60.2	60.2	97.3
GAP S16	51.6	51.3	50.6	57.2	57.2	57.2	57.2	96.0
GAP W58	30.6	0.0	0.0	23.8	93.6	100.0	100.0	100.0
Network H	–	74.3	67.0	–	–	18.4	21.6	94.4
Network T3	–	23.3	3.3	–	–	56.3	56.9	56.9
Network S4	–	3.3	0.0	–	–	63.4	63.7	63.7
Network S8	–	40.0	26.7	–	–	62.0	63.9	65.8
Network S16	–	60.0	56.7	–	–	42.4	43.6	75.4

these levels exactly: a two-valued level is straddle-free under $\bar{u} = u$ if and only if its capacity does not bind or $u - u^*_\ell$ is a multiple of $\beta - \alpha$, and by Corollary 4, since 3 does not divide 5, no agent on which two or more levels bind is straddle-free. On two-valued blocks the exact capacities are given in closed form by Proposition 2(c), so this improvement requires no additional computation. Fixing items at bound removes few straddles in GAP and leaves the bound unchanged, but in network design it makes class S4 straddle-free, so that it attains z_{DW} by Corollary EC.1, and on class T3 it removes most straddles and matches z_{DW} to the reported precision. On class T3, whose demands $8 > 7 > 5$ have gaps 1 and 2, Theorem EC.3(a) applies: the integer hull of every level is the cardinality slice of a two-coefficient knapsack polytope with coefficients 3 and 2, and Section 7.3 closes the remaining gap with level-hull inequalities. Both refinements are polynomial except for the exact capacities of blocks with straddles, which require a dynamic program over count and load; both are advisable whenever (CI) is used.

7.3 Level Hulls Restore the Dantzig–Wolfe Bound

(E5) On the network instances of Table 4, regenerated deterministically with the same generator and seeds, which are listed in the code archive, and checked against the recorded values of $z_{\mathrm{LP}}(\mathrm{C})$, we add two families of inequalities to LP(CI) on the levels that contain a straddle. CI+P adds, once, the inequalities $\sum_{k \in P_p} x_{k\ell} \le C_p(T^*)\, y_\ell$ of Theorem 6 on every level with a straddle that Proposition 4 recognizes as principal. CI+H separates over $\mathrm{conv}(X_\ell)$ exactly: in each round, for every arc and every level with a straddle and $y_\ell > 0$, a linear program over the feasible ℓ-sets returns a most violated inequality $\alpha^\top x_{\cdot\ell} \le \beta\, y_\ell$ with $\|\alpha\|_\infty \le 1$, until no inequality is violated by more than 10^{-7}.

On every instance of every class, the bound of CI+H equals z_{DW} to within 1.5×10^{-14} in relative

Table 6: Level-hull inequalities on the network instances of Table 4 (ten instances per class). Levels with a straddle: levels with a straddle among all levels, summed over the class; within an instance all arcs have the same demands and capacity, so each level is counted once per instance; Principal: levels with a straddle that are principal. Columns CI to DW: share (%) of $z^* - z_{\mathrm{LP}}(\mathrm{C})$ closed. Maximum difference: largest relative difference between the bound of CI+H and z_{DW}. Cuts: average number of hull inequalities per instance. Rounds: largest number of separation rounds.

Class	Levels with a straddle	Principal	CI	CI+P	CI+H	DW	Maximum difference	Cuts	Rounds
H	35/47	4	18.4	18.4	94.4	94.4	4×10^{-16}	54.5	13
T3	7/30	7	56.3	56.9	56.9	56.9	3×10^{-15}	0.9	4
T2	0/30	0	62.0	62.0	62.0	62.0	2×10^{-14}	0.0	0
S4	1/23	1	63.4	63.7	63.7	63.7	4×10^{-15}	0.1	1
S8	12/30	8	62.0	64.4	65.8	65.8	1×10^{-16}	1.5	6
S16	19/32	2	42.4	42.8	75.4	75.4	8×10^{-15}	7.8	4

terms, as Theorem 5 states (Table 6). On class H the hull inequalities raise the share of the gap closed from 18.4% to 94.4%, and on class S16 from 42.4% to 75.4%, with at most 13 separation rounds. By Corollary 6, CI+P attains z_{DW} on the 36 instances in which every level with a straddle is principal, including all ten instances of class T3, and it does on all 36; on the other 24 instances it does on only 3. On class T3, the compact principal-level inequalities thus close the whole gap between (CI) and z_{DW} without separation, whereas on class H, where only 4 of the 35 levels with a straddle are principal, they do not improve the bound, and the general hull inequalities are needed.

7.4 Level-Hull Inequalities on Harder Instances

The instances of Table 6 are small. To measure how much of the root gap the level hulls close in the setting where Theorem 4 implies the largest effect, we generated harder network instances with tight arc capacities and high fixed costs: the capacity is $u = \kappa$ times the mean demand with $\kappa \in \{1.5, 2\}$, so that an arc carries at most 2 to 7 commodities, and the fixed costs are multiplied by 3 or 10, which gives root gaps of 11.6–23.2%. There are 13 instances, 11 with 12 nodes and 20 commodities and 2 with 20 nodes and 40 commodities, and all are reported in Table EC.6 (Section EC.9). The exact level-hull inequalities were separated at the root of (CI) on the levels with a straddle.

On all 13 instances, the level hulls close 74.3–99.1% of the root gap (mean 89.3%), against 8.1–78.8% (mean 34.5%) for (CI) alone. These means include every root separation; separation converged on all 13 instances, but not always within the time limit of the integer solve. On 12 instances it took at most 17 s. On the largest class-H instance it took 1,046 s, more than the 300 s limit, so its root bound, which closes 92.8% of the gap, is reported for completeness, and no branch-and-bound run followed. Over the other 12 instances the mean is 89.0%. The gain is largest where (CI) is weakest: on class H, whose levels have many distinct weights, (CI) closes 8.1–16.8% and the level hulls 82.9–99.0%. Section EC.9 also reports branch-and-bound runs with the open-source solver HiGHS, including the level windows and reduction of the network design study. The two

families differ in what they need. The principal-level inequalities are computed from the data before optimization, whereas exact level-hull separation solves an optimization problem over the feasible sets of each level with a straddle in every round, here by a dynamic program over count and load. On the largest class-H instance, exact separation therefore dominates the computational effort, and we report this instance as an indication of the limits of exact separation rather than as an outlier; for larger instances, column generation or objective-specific block cuts (Chen et al., 2025) are the alternatives, while the recognition test still identifies the levels on which no separation is needed. We make no claim about solution times, for which we refer to the applied studies of Section 7.8.

7.5 Strengthened Covers

(E4) Starting from LP(CI) on the GAP classes with straddles, four cutting-plane loops add at most one cut per agent and level per round, for at most 30 rounds. Loop LC adds the most violated level cover; LC-facet adds it only when Corollary 7 certifies a facet; Reduced shrinks each cover $x(C) \le \ell - 1$ to a subset H while $x(H) \le |H| - 1$ remains valid, which gives a dominating cut; and Strengthened then applies Corollary 8.

Table 7: Cut families on the GAP classes with straddles ($m = 5$, $n = 30$; ten instances per class): share (%) of $z^* - z_{\mathrm{LP}}(\mathrm{C})$ closed, and number of cuts certified facet-defining by Theorem 7 among the cuts added.

Class	CI	LC	LC-facet	Reduced	Strengthened	DW	Facets	
							Reduced	Strengthened
C	21.8	42.1	21.8	47.9	66.5	88.6	2/258	233/284
D	4.9	12.3	4.9	15.4	23.7	78.3	7/352	440/463
E	22.3	30.7	22.3	30.9	36.5	74.1	0/184	8/224
S2	87.6	93.7	87.6	94.9	99.3	99.3	0/27	23/23
S4	75.0	89.8	75.0	91.1	97.5	99.7	0/49	45/46
S8	60.2	78.0	60.2	83.7	91.2	97.3	2/83	70/76
S16	57.2	77.4	57.2	81.4	87.7	96.0	9/103	99/102

Only 4 of the 992 level covers separated by loop LC define facets, and LC-facet does not improve on (CI). Corollary 7 explains why: at a straddle vertex the only violated level cover is that of $A \cup \{i\}$, and its two exchange conditions usually fail. After strengthening, 918 of the 1,218 cuts are facet-defining, and the share of the gap closed rises on every class (Table 7), by 18.6 percentage points on type C and 8.3 on type D, matching z_{DW} on class S2. On type E, 133 of the 224 strengthened covers fail condition (H1), whose failure yields a valid inequality $x(H) + x_a \le h$ with a coefficient of two; it is not a rank-type cover, and lifting it is the natural next step.

7.6 Small Blocks and Capacity Scaling

An exhaustive exact-arithmetic census of all blocks with up to seven items and weights up to eight (Table EC.2 in Section EC.7) agreed with Theorems 6 and 7, Corollary 7 and Lemma EC.2 on every level. Most small levels are principal: 92% of the full-dimensional levels with three distinct

weights, falling to 64% with seven. These frequencies describe the enumerated grid only; on blocks with tens of distinct weights, as in the GAP benchmarks, principal levels are rare (Section 7.7).

(E6) When the weights and capacities of two-valued GAP instances with $m = 10$ and $n = 100$ are scaled by factors up to 1000, (CI) keeps its 10,100 variables and equals z_{DW}, while a load-based dynamic program grows from 5×10^4 to 5×10^7 states (Section EC.8). Since pricing for two-valued blocks is also polynomial, by enumerating the number of heavy items and sorting, this experiment demonstrates the capacity-independent size of (CI) and its recognition procedure, not a computational advantage over column generation. Used directly as an integer program without valid inequalities or reduced level sets, the unreduced formulation is not competitive in solution time with the compact formulation (Section EC.8).

7.7 The Standard GAP Benchmarks

On the standard benchmarks, the recognition test shows that no agent of any instance is straddle-free, so Corollary 1 provides no guarantee, and on the instances with a positive gap (CI) closes 26.1% of the gap on average against 80.3% for z_{DW}. The recognition results apply before any model is built, so they can be checked on public instances. We applied Theorem 4 with Proposition 2, Corollary EC.1 and Proposition 4 to every agent and level of the 117 GAP benchmark instances of OR-Library and Yagiura et al. (2006): the classical sets gap1–gap12, the instances of types A–E with up to 200 jobs, and the instances of types C, D and E with 400 to 1,600 jobs (Section EC.7). These tests are polynomial, and no exact level capacity is computed. Table 8 reports the results.

Table 8: Straddle census of the GAP benchmark instances under exact level capacities. Levels: (agent, level) pairs. Columns 5–7: share (%) of levels that are straddle-free, straddle-free after fixing items at bound, and straddle-free or principal. Exact agents: agents that are straddle-free at every level. Distinct weights: median over instances of the median number of distinct weights per agent.

Set	Instances	Agents	Levels	Straddle-free	Fixed	or principal	Exact agents	Distinct weights
gap1–12 (type C)	60	460	3,289	26.5	27.4	28.5	0	16
Type A, $n \leq 200$	6	70	2,294	28.5	28.5	29.3	0	21
Type B, $n \leq 200$	6	70	1,688	25.7	25.9	25.9	0	21
Type C, $n \leq 200$	6	70	1,612	24.9	24.9	25.0	0	21
Type D, $n \leq 200$	6	70	2,539	13.7	13.8	13.9	0	75
Type E, $n \leq 200$	6	70	3,240	4.5	4.5	4.7	0	33
Type C, $n \geq 400$	9	315	19,332	21.1	21.2	21.2	0	21
Type D, $n \geq 400$	9	315	45,578	7.5	7.5	7.6	0	100
Type E, $n \geq 400$	9	315	50,596	2.3	2.3	2.4	0	51

No agent of any benchmark instance is straddle-free, and only 8.9% of the 130,168 levels are. Of these 11,531 levels, 11,191 satisfy condition (i), so the capacity does not bind there, and the remaining 340 satisfy condition (ii); every other level has a straddle. This is consistent with Theorem 4, since the agents have 16 to 100 distinct weights. Fixing items at bound and the principal-level inequalities change these shares by less than two percentage points. On these benchmarks, therefore,

the theory gives no guarantee that LP(CI) attains z_{DW}, although a straddle lowers the bound only for suitable costs (Theorem 3). It also implies that the inequalities needed to close the gap are general level-hull inequalities rather than the principal-level inequalities.

We also computed the three bounds on the instances of types A–E with 5, 10 or 20 agents and 100 or 200 jobs, using the optimal value, or the best-known value where optimality has not been proven, as z^*. Table 9 reports all 30 instances.

Table 9: Bounds on the GAP benchmark instances of types A–E with 100 or 200 jobs, under exact level capacities. Included: instances with $z^* > z_{\mathrm{LP}}(\mathrm{C})$, over which the shares are averaged. Gap: $(z^* - z_{\mathrm{LP}}(\mathrm{C}))/z^*$ (%), averaged over all instances. Columns CI to DW: share (%) of $z^* - z_{\mathrm{LP}}(\mathrm{C})$ closed. CI=DW: instances with $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$ to a relative 10^{-6}. The value z^* is a best-known value, not a proven optimum, for d10100, d10200, d20100, d20200, e05200, e10100, e10200 and e20200; for these the shares are measured against the best-known value.

Set	Instances	Included	Gap	CI	CI+P	DW	CI=DW
Type A	6	5	0.05	52.5	52.5	100.0	3
Type B	6	6	0.48	20.7	20.7	81.3	0
Type C	6	6	0.75	18.6	18.6	84.5	0
Type D	6	6	0.27	3.8	3.8	62.4	0
Type E	6	6	0.28	39.3	39.3	76.4	0

The bounds are consistent with the census (Table 9). Over the 29 instances with a positive gap, (CI) closes 26.1% of the gap on average and z_{DW} closes 80.3%. On these instances the principal-level inequalities never changed the bound. The loss is largest on type D, where (CI) closes 3.8% against 62.4%; type D is the class whose costs decrease with the weight, which is the unfavorable case identified by Proposition EC.6. On three type-A instances, (CI) attains z_{DW} although every agent has straddles: a straddle lowers the bound only for suitable costs. On every instance, $z_{\mathrm{LP}}(\mathrm{C}) \le z_{\mathrm{LP}}(\mathrm{CI}) \le z_{\mathrm{DW}} \le z^*$ holds to within 10^{-12} in relative terms. These results do not contradict the branch-and-cut studies discussed next, which add cardinality-based valid inequalities and branch on the level variables. They show that, on these benchmarks, the root relaxation of (CI) alone stays well below z_{DW}, and Theorem 5 identifies what closes the remaining gap.

7.8 Relation to the Branch-and-Cut Studies

Cardinality indexing has been used successfully inside branch-and-cut, with reduced level sets and cardinality-based valid inequalities, for generalized assignment (Murthy and Ransbotham, 2019), concentrator location (Di Francesco et al., 2021), capacitated Steiner arborescence (Contu et al., 2026) and unsplittable network design (Anonymous, 2026a); those studies report stronger root bounds and smaller search trees than the compact formulations. Their results are not reproduced here and are not evidence for the present theorems, but the theory clarifies their scope. Inequalities on single resources cannot raise the bound above z_{DW} (Corollary 1 and Theorem 5), and on the GAP benchmarks the root relaxation of (CI) stays well below z_{DW} (Section 7.7). Bounds beyond z_{DW} therefore require inequalities that couple several resources, and the single-resource results imply

that the largest arc-level gains arise when resources serve few items with few distinct weights. The level covers of Section 3 are the single-level cases of the cardinality-based cover inequalities used in these studies, so the facet conditions of Corollary 7 and the strengthening of Corollary 8 apply to those single-level cases, and the procedure of Section 6.4 indicates, resource by resource, where such inequalities can close the gap to z_{DW} and where they cannot. In the GAP, concentrator location and Steiner arborescence formulations, which use the original capacity, Corollary 4 shows that replacing u by u^* can make two-valued resources exact without additional computation.

8 Conclusions

Cardinality-indexed formulations are exact on a resource precisely when it has no straddle, because the count changes by one along every edge of the unit cube and no fractional vertex of one level is absorbed by the others; indexing by the load lacks this property and is exact only for equal weights. Under exact level capacities, straddle-freeness means at most two distinct weights or an extreme capacity at every level, which can be tested in polynomial time before the model is built, and which guarantees the Dantzig–Wolfe bound under coupling through aggregated variables by a linear program whose size does not depend on the capacities. When straddles occur, the Dantzig–Wolfe bound is recovered exactly by the integer hulls of the levels. These hulls are nested-matroid polytopes on principal levels, recognized in polynomial time, and cardinality slices of two-coefficient knapsack polytopes for three weights whose gaps divide each other (Section EC.4). On the remaining levels, the facet-defining rank-type covers are characterized by four exchange conditions, which also yield a strengthening procedure. In the experiments, these conditions identify where the formulation attains the Dantzig–Wolfe bound and where it does not, and strengthened covers are mostly facet-defining.

The results have limits that bound their interpretation. They concern root relaxations, not solution times, and exactness does not make the integer problem easy. The necessity of straddle-freeness holds when the coupling can complete a straddle, which we establish for GAP-type assignment coupling but do not characterize for network design. The facet characterization covers rank-type covers only, and these do not describe every level. On the standard GAP benchmarks, where agents have many distinct weights, the formulation alone stays far below the Dantzig–Wolfe bound, and exact separation over the level hulls becomes expensive as instances grow. Finally, single-resource inequalities cannot go beyond the Dantzig–Wolfe bound.

Four questions remain open. The first is the description of levels with three weights whose gaps do not divide one another, or with four or more weights. The second is the lifting of the non-rank-type inequalities that arise when condition (H1) fails. The third is the structure of inequalities that couple several resources, which is where any improvement beyond the Dantzig–Wolfe bound must come from; the exactness of relaxations whose level variables are coupled across resources is studied in Anonymous (2026b). The fourth is computational: how best to exploit the root-bound gains of Section 7.4 inside branch-and-cut.

Code and data. The code, instance generators and seeds, raw results, exact-arithmetic verification scripts, and scripts that regenerate every table are submitted with the manuscript as a compressed archive with a README file, and will be released in a public repository upon acceptance.

References

4ti2 team (2018) 4ti2: A software package for algebraic, geometric and combinatorial problems on linear spaces. Software, accessed November 2026, https://4ti2.github.io.

Anonymous (2026a) A cardinality-indexed extended formulation with valid inequalities for unsplittable multicommodity network design. Working paper; author details omitted for double-anonymous review.

Anonymous (2026b) When are coupled cardinality-indexed relaxations exact? Working paper; author details omitted for double-anonymous review.

Ardila F (2003) The Catalan matroid. *J. Combin. Theory Ser. A* 104(1):49–62.

Atamtürk A, Rajan D (2002) On splittable and unsplittable flow capacitated network design arc-set polyhedra. *Math. Programming* 92(2):315–333.

Balas E (1975) Facets of the knapsack polytope. *Math. Programming* 8(1):146–164.

Balas E (1998) Disjunctive programming: Properties of the convex hull of feasible points. *Discrete Appl. Math.* 89(1–3):3–44.

Bruns W, Ichim B, Söger C, von der Ohe U (2024) Normaliz: Algorithms for rational cones and affine monoids. Software, accessed November 2026, https://www.normaliz.uni-osnabrueck.de.

Caprara A, Kellerer H, Pferschy U, Pisinger D (2000) Approximation algorithms for knapsack problems with cardinality constraints. *European J. Oper. Res.* 123(2):333–345.

Chen R, Zhu H (2025) The complexity of recognizing facets for the knapsack polytope. *Math. Oper. Res.*, ePub ahead of print, https://doi.org/10.1287/moor.2024.0481.

Chen R, Günlük O, Lodi A (2025) Recovering Dantzig–Wolfe bounds by cutting planes. *Oper. Res.* 73(2):1128–1142.

Chen L, Chen WK, Yang MM, Dai YH (2021) An exact separation algorithm for unsplittable flow capacitated network design arc-set polyhedron. *J. Global Optim.* 81(3):659–689.

Chouman M, Crainic TG, Gendron B (2017) Commodity representations and cut-set-based inequalities for multicommodity capacitated fixed-charge network design. *Transportation Sci.* 51(2):650–667.

Chu PC, Beasley JE (1997) A genetic algorithm for the generalised assignment problem. *Comput. Oper. Res.* 24(1):17–23.

Conforti M, Cornuéjols G, Zambelli G (2014) *Integer Programming* (Springer, Cham).

Contu F, Di Francesco M, Gorgone E, Murthy I (2026) An extended formulation with valid inequalities for the capacitated Steiner arborescence problem. *Networks* 88(2):250–274.

Correia I, Gouveia L, Saldanha-da-Gama F (2010) Discretized formulations for capacitated location problems with modular distribution costs. *European J. Oper. Res.* 204(2):237–244.

Crainic TG, Gendron B (2021) Exact methods for fixed-charge network design. Crainic TG, Gendreau M, Gendron B, eds. *Network Design with Applications to Transportation and Logistics* (Springer, Cham), 29–89.

Dahl G, Foldnes N (2003) Complete description of a class of knapsack polytopes. *Oper. Res. Lett.* 31(5):335–340.

de Farias IR Jr, Nemhauser GL (2003) A polyhedral study of the cardinality constrained knapsack problem. *Math. Programming* 96(3):439–467.

de Lima VL, Alves C, Clautiaux F, Iori M, Valério de Carvalho JM (2022) Arc flow formulations based on dynamic programming: Theoretical foundations and applications. *European J. Oper. Res.* 296(1):3–21.

De Loera JA, Hemmecke R, Köppe M (2013) *Algebraic and Geometric Ideas in the Theory of Discrete Optimization* (SIAM, Philadelphia).

Di Francesco M, Gaudioso M, Gorgone E, Murthy I (2021) A new extended formulation with valid inequalities for the capacitated concentrator location problem. *European J. Oper. Res.* 289(3):975–986.

Diaconis P, Graham RL, Sturmfels B (1996) Primitive partition identities. Miklós D, Sós VT, Szőnyi T, eds. *Combinatorics, Paul Erdős is Eighty*, Bolyai Society Mathematical Studies, vol. 2 (János Bolyai Mathematical Society, Budapest), 173–192.

Frangioni A, Gendron B (2009) 0–1 reformulations of the multicommodity capacitated network design problem. *Discrete Appl. Math.* 157(6):1229–1241.

Gale D (1968) Optimal assignments in an ordered set: An application of matroid theory. *J. Combin. Theory* 4(2):176–180.

Geist D, Rodin EY (1992) Adjacency of the 0–1 knapsack problem. *Comput. Oper. Res.* 19(8):797–800.

Geoffrion AM (1974) Lagrangean relaxation for integer programming. *Math. Programming Stud.* 2:82–114.

Goemans MX, Rothvoss T (2020) Polynomiality for bin packing with a constant number of item types. *J. ACM* 67(6), https://doi.org/10.1145/3421750.

Gouveia L (1995) A $2n$ constraint formulation for the capacitated minimal spanning tree problem. *Oper. Res.* 43(1):130–141.

Gouveia L (1998) Using variable redefinition for computing lower bounds for minimum spanning and Steiner trees with hop constraints. *INFORMS J. Comput.* 10(2):180–188.

Gouveia L, Saldanha-da-Gama F (2006) On the capacitated concentrator location problem: A reformulation by discretization. *Comput. Oper. Res.* 33(5):1242–1258.

Graver JE (1975) On the foundations of linear and integer linear programming I. *Math. Programming* 9(1):207–226.

Greening LM, Dey SS, Erera AL (2026) Strengthening dual bounds for multicommodity capacitated network design with unsplittable flow constraints. *European J. Oper. Res.* 334(2):496–508.

Gu Z, Nemhauser GL, Savelsbergh MWP (2000) Sequence independent lifting in mixed integer programming. *J. Combin. Optim.* 4(1):109–129.

Hampe S (2017) The intersection ring of matroids. *J. Combin. Theory Ser. B* 122:578–614.

Herrmann S, Joswig M (2008) Splitting polytopes. *Münster J. Math.* 1:109–142.

Hojny C, Gally T, Habeck O, Lüthen H, Matter F, Pfetsch ME, Schmitt A (2020) Knapsack polytopes: A survey. *Ann. Oper. Res.* 292(1):469–517.

Joswig M, Schröter B (2017) Matroids from hypersimplex splits. *J. Combin. Theory Ser. A* 151:254–284.

Köppe M, Louveaux Q, Weismantel R (2008) Intermediate integer programming representations using value disjunctions. *Discrete Optim.* 5(2):293–313.

Laguna M, Kelly JP, González-Velarde JL, Glover F (1995) Tabu search for the multilevel generalized assignment problem. *European J. Oper. Res.* 82(1):176–189.

Martin RK (1987) Generating alternative mixed-integer programming models using variable redefinition. *Oper. Res.* 35(6):820–831.

Maurras JF, Stephan R (2011) On the cardinality constrained matroid polytope. *Networks* 57(3):240–246.

Murthy I, Ransbotham S (2019) A new extended formulation of the generalized assignment problem and some associated valid inequalities. *Discrete Appl. Math.* 271:119–143.

Raack C, Koster AMCA, Orlowski S, Wessäly R (2011) On cut-based inequalities for capacitated network design polyhedra. *Networks* 57(2):141–156.

Savelsbergh M (1997) A branch-and-price algorithm for the generalized assignment problem. *Oper. Res.* 45(6):831–841.

Schrijver A (2003) *Combinatorial Optimization: Polyhedra and Efficiency* (Springer, Berlin).

Uchoa E, Fukasawa R, Lysgaard J, Pessoa A, Poggi de Aragão M, Andrade D (2008) Robust branch-cut-and-price for the capacitated minimum spanning tree problem over a large extended formulation. *Math. Programming* 112(2):443–472.

Valério de Carvalho JM (1999) Exact solution of bin-packing problems using column generation and branch-and-bound. *Ann. Oper. Res.* 86:629–659.

Weismantel R (1996) Hilbert bases and the facets of special knapsack polytopes. *Math. Oper. Res.* 21(4):886–904.

Wentges P (1997) Weighted Dantzig–Wolfe decomposition for linear mixed-integer programming. *Internat. Trans. Oper. Res.* 4(2):151–162.

Wong RT (1984) A dual ascent approach for Steiner tree problems on a directed graph. *Math. Programming* 28(3):271–287.

Yagiura M, Ibaraki T, Glover F (2006) A path relinking approach with ejection chains for the generalized assignment problem. *European J. Oper. Res.* 169(2):548–569.

Yamín D, van Hoeve WJ, Ralphs TK (2026) Dantzig–Wolfe and arc-flow reformulations: A systematic comparison. Preprint, arXiv:2602.22589.

Zeng B, Richard JPP (2011) A polyhedral study on 0–1 knapsack problems with disjoint cardinality constraints: Facet-defining inequalities by sequential lifting. *Discrete Optim.* 8(2):277–301.

E-Companion

Straddles and the Dantzig–Wolfe Bound of Cardinality-Indexed Knapsack Formulations

Md Shahrukh Anjum[1,*], Jitamitra Desai[2], and Ishwar Murthy[2]

[1]Amrut Mody School of Management, Ahmedabad University, Ahmedabad, India

[2]Decision Sciences Area, Indian Institute of Management Bangalore, Bengaluru, India

*Corresponding author: shahrukh.anjum@ahduni.edu.in

Results, sections, equations and tables without the prefix EC refer to the main text.

EC.1 Proofs and Further Results for Sections 2–5

EC.1.1 Disjunctive Cuts

Proposition EC.1. *Let $D_1 \vee D_2$ be the setup split $\{y \le 0\} \vee \{y \ge 1\}$ or a cardinality split $\{|x| \le (s-1)y\} \vee \{|x| \ge sy\}$ with $s \in \mathbb{Z}$. Then $Q \subseteq \operatorname{conv}\big((B^C \cap D_1) \cup (B^C \cap D_2)\big)$, so every disjunctive cut derived from B^C and such a disjunction is implied by Q. A disjunction $\{y_\ell \le 0\} \vee \{y_\ell \ge 1\}$ on a level indicator yields no cut for B.*

Proof. Each term of the disjunction in Theorem 1, intersected with B^C, lies in D_1 or D_2: $(\mathbf{0}, 0)$ satisfies $y \le 0$, and $|x| \le (s-1)y$ when $s \ge 1$; a term with $y = 1$ and $|x| = \ell$ satisfies $y \ge 1$ and lies in the first or second part of a cardinality split according to whether $\ell \le s-1$; for $s \le 0$ the split is trivial. Hence Q, the convex hull of the terms, lies in the convex hull of the two parts. Every vertex of B has $y_\ell \in \{0, 1\}$ by Theorem 1, so $B = \operatorname{conv}\big((B \cap \{y_\ell = 0\}) \cup (B \cap \{y_\ell = 1\})\big)$. □

Disjunctive cuts on level indicators can therefore strengthen (CI) only through Γ; disjunctions on assignment variables are not coarsenings of the cardinality disjunction and can strengthen a single block.

EC.1.2 Proof of Proposition 2

(a) An ℓ-set of weight τ_ℓ is feasible if and only if $\tau_\ell \le u$, and $u^*_\ell \le \tau_\ell$. (b) Since $\ell \in \Lambda$, $\sigma_\ell \le u$, and (ii) fails if and only if some feasible ℓ-set S has $d(S) > \sigma_\ell$. Every item outside L_ℓ is at least as heavy as every item of L_ℓ. Pairing the items of $L_\ell \setminus S$ with those of $S \setminus L_\ell$, every pair (k, o) has $d_o \ge d_k$, and $d(S) - \sigma_\ell$ is the sum of the differences, so some pair has $d_o > d_k$, and $(L_\ell \setminus \{k\}) \cup \{o\}$ has weight in $(\sigma_\ell, d(S)]$. Hence $\sigma_\ell + g_\ell$ is the smallest weight of an ℓ-set heavier than σ_ℓ. (c) An ℓ-set with p items of weight β has weight $\alpha\ell + (\beta - \alpha)p$, increasing in p; such a set exists if and only if $\max\{0, \ell - n_\alpha\} \le p \le \min\{n_\beta, \ell\}$ and fits if and only if $p \le \lfloor (u - \alpha\ell)/(\beta - \alpha) \rfloor$; the lightest ℓ-set fits, so the largest admissible p is p_ℓ. For the running time, sort the weights once; each level takes $O(|K|)$ time, and Corollary 3 combines the results. □

EC.1.3 General Level Capacities

Proposition EC.2. *Let $\bar{u}_\ell$ be admissible and $\delta_\ell = \bar{u}_\ell - u^*_\ell$. (a) If $\bar{u}'_\ell \le \bar{u}_\ell$ are admissible and level ℓ is straddle-free for $\bar{u}_\ell$, it is straddle-free for $\bar{u}'_\ell$; in particular, one of (i)–(iii) of Theorem 4 is necessary. (b) Under (i), level ℓ is straddle-free for every admissible $\bar{u}_\ell$. (c) If the weights take two values $\alpha < \beta$, level ℓ is straddle-free if and only if (i) holds or δ_ℓ is a multiple of $\beta - \alpha$. (d) Under (ii), level ℓ need not be straddle-free.*

Proof. (a) $P_\ell(\bar{u}'_\ell) \subseteq P_\ell(\bar{u}_\ell)$ and both have the integer points X_ℓ. If $P_\ell(\bar{u}_\ell)$ is integral (Theorem 2(a)), then $\mathrm{conv}(X_\ell) \subseteq P_\ell(\bar{u}'_\ell) \subseteq P_\ell(\bar{u}_\ell) = \mathrm{conv}(X_\ell)$. Taking $\bar{u}'_\ell = u^*_\ell$ and applying Theorem 4 gives the second statement. (b) Under (i), $\bar{u}_\ell - d(A) \ge \tau_\ell - d(A) \ge \max_{k \notin A} d_k$. (c) Write $\Delta = \beta - \alpha$, let n_β and n_α be the numbers of heavy and light items, and let $s = \bar{u}_\ell - \alpha\ell$. By Proposition 2(c), $u^*_\ell = \alpha\ell + \Delta p^*$ for an integer p^*, so $\Delta \mid s$ if and only if $\Delta \mid \delta_\ell$. An $(\ell-1)$-set A with q heavy items has residual $\bar{u}_\ell - d(A) = \alpha + s - q\Delta$, and with two weights a straddle requires a heavy and a light item outside A and $\alpha < \alpha + s - q\Delta < \beta$, that is, $q = \lfloor s/\Delta \rfloor$ and $\Delta \nmid s$. If $\Delta \mid s$, no residual lies strictly between α and β. Under (i) the level is straddle-free by (b). Suppose that $\Delta \nmid s$ and (i) fails, and let $q = \lfloor s/\Delta \rfloor$. Since (i) fails, $u < \tau_\ell = \alpha\ell + \Delta \min\{n_\beta, \ell\}$, so $s \le u - \alpha\ell < \Delta \min\{n_\beta, \ell\}$ and $q \le \min\{n_\beta, \ell\} - 1$. Some feasible ℓ-set exists and contains at least $\ell - n_\alpha$ heavy items, so $s \ge u^*_\ell - \alpha\ell \ge \Delta(\ell - n_\alpha)$ and $q \ge \ell - n_\alpha$. Hence an $(\ell-1)$-set with q heavy and $\ell - 1 - q$ light items exists and leaves at least one item of each weight outside, and it is the first component of a straddle. (d) For weights $(2, 3, 5)$ and $u = 6$, $u^*_2 = 5 = \sigma_2$, but $\bar{u}_2 = 6$ gives the straddle with A of weight 2 and residual $4 \in (3, 5)$. □

For example, with weights $(1, 1, 4)$ and $u = 4$, $u^*_2 = 2$, and $\bar{u}_2 = 4$ gives the straddle with A one item of weight 1 and residual $3 \in (1, 4)$, since $\delta_2 = 2$ is not a multiple of 3. Consecutive integer weights $\{w, w+1\}$ are straddle-free for every admissible capacity by (c). For the original capacity, (c) gives the test of Corollary 4 in Section 4.

Proof of Corollary 4. With $\bar{u}_\ell = u$, $\delta_\ell = u - u^*_\ell$, and $\Delta \mid \delta_\ell$ if and only if $\Delta \mid (u - \alpha\ell)$ by the proof of Proposition EC.2(c); condition (i) holds exactly when $\tau_\ell \le u$. Since τ_ℓ increases with ℓ, the levels with $\tau_\ell > u$ are consecutive, and if ℓ and $\ell + 1$ are among them, Δ divides $u - \alpha\ell$ and $u - \alpha(\ell+1)$ if and only if it divides α and u. □

Both statements were checked by enumeration: Proposition EC.2(c) on all 805,050 levels and admissible capacities with $\alpha \le 6$, $\beta \le 9$ and at most five items of each weight, and Corollary 4 on all 33,950 blocks with $\alpha \le 7$, $\beta \le 11$ and at most five items of each weight, with no exception. Class W58 of Table 5 has $\Delta = 3 \nmid 5$, so under $\bar{u} = u$ no agent on which two or more levels bind is straddle-free; Section 7.2 reports the effect on the bound.

EC.1.4 Fixing Items at Bound

The conditions refer to all items, including items that belong to no feasible ℓ-set, or to every one. Let Z_ℓ and O_ℓ be these two sets; $k \in Z_\ell$ if and only if $d_k + \sigma_{\ell-1}(K \setminus \{k\}) > u$, and $k \in O_\ell$ if and

only if $\sigma_\ell(K \setminus \{k\}) > u$, so both are found in $O(|K| \log |K|)$ time.

Corollary EC.1. *Let $\bar{u} = u^*$, and let $\tilde{Q}$ be the cardinality relaxation after fixing $x_{k\ell} = 0$ for $k \in Z_\ell$ and $x_{k\ell} = y_\ell$ for $k \in O_\ell$ at every level. Then $\tilde{Q} = \mathrm{conv}(F^\circ)$ if and only if, for every ℓ, the block on the items $K \setminus (Z_\ell \cup O_\ell)$ with capacity $u - d(O_\ell)$ satisfies (i), (ii) or (iii) at level $\ell - |O_\ell|$.*

For weights $(1, 1, 2, 2, 10)$, $u = 10$ and $\ell = 2$, the heavy item belongs to no feasible pair and creates straddles, but the reduced level has the weights $1, 1, 2, 2$ only and is exact. The corollary explains the network results of Table 5, where fixing makes class S4 straddle-free; on the GAP benchmarks its effect is small (Table 8).

Proof. By Lemma EC.2, the equations $x_k = 0$ ($k \in Z_\ell$) and $x_k = 1$ ($k \in O_\ell$) hold on $\mathrm{conv}(X_\ell)$, so fixing is valid. Identify the reduced level with its projection onto $K' = K \setminus (Z_\ell \cup O_\ell)$: it is the level polytope at level $\ell' = \ell - |O_\ell|$, with level capacity $u^*_\ell - d(O_\ell)$, of the block on K' with capacity $u - d(O_\ell)$. The feasible ℓ'-sets of this block are the sets $S \setminus O_\ell$ with $S \in X_\ell$, so its exact level capacity is $u^*_\ell - d(O_\ell)$, and Theorems 2(a) and 4 apply. The proof of Theorem 2(b) uses only the bounds and the inclusion of the reduced levels in the original ones, so it applies verbatim. □

EC.1.5 Hypersimplex Splits

For every right-hand side at once, classical tools characterize integrality: for $1 \le \ell \le |K| - 1$, $\{\chi \in [0,1]^K : |\chi| = \ell,\ d(\chi) \le \omega\}$ is integral for every integer ω if and only if $\max_k d_k - \min_k d_k \le 1$; for two weights, total unimodularity gives integrality whenever the right-hand side is congruent to $\alpha\ell$ modulo $\beta - \alpha$. These criteria give the sufficiency in Theorem 4 but not its necessity, which concerns u^*_ℓ, a value always attained by a feasible set. Under exact capacities, necessity can be derived from polyhedral geometry. Suppose (i) and (ii) fail, so $2 \le \ell \le |K| - 1$ and $\{d(\chi) = u^*_\ell\}$ meets the relative interior of $\Delta = \{\chi \in [0,1]^K : |\chi| = \ell\}$. If $\ell = |K| - 1$, Δ is a simplex, and a hyperplane through its relative interior crosses an edge, so there is a straddle. Otherwise, if level ℓ has no straddle, the hyperplane crosses no edge of Δ in its relative interior and defines a split of the hypersimplex, and by the classification of hypersimplex splits of Herrmann and Joswig (2008) (see also Joswig and Schröter 2017) it coincides on the affine hull with $\{\sum_{i \in S} x_i = \mu\}$, so d takes two values. This route is nonconstructive; the proof of Theorem 4 returns a straddle and a violated level cover.

EC.1.6 Hardness of Two-Valued GAP

Proposition EC.3. *For GAP instances in which every agent sees at most two distinct weights, deciding feasibility is strongly NP-complete already for weights in $\{1, 3\}$ and capacities 3, and minimizing cost is strongly NP-hard on instances known to be feasible; on all such instances $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$ with $\bar{u} = u^*$.*

A constant number of item sizes makes bin packing polynomial (Goemans and Rothvoss, 2020); two weights per agent do not suffice for GAP.

Proof. Reduce from three-dimensional matching: given disjoint sets W, X, Y of size q and triples $T_1, \dots, T_m \subseteq W \times X \times Y$ with $m \geq q$, decide whether q triples cover $W \cup X \cup Y$. Create one agent of capacity 3 per triple; one job of weight 1 per element, eligible exactly for the agents whose triples contain it; and $m - q$ dummy jobs of weight 3, eligible for all agents. If q triples form a cover, assign their elements to their agents and one dummy job to each other agent. Conversely, in a feasible assignment an agent with a dummy job has no other job, so the $m - q$ dummy jobs occupy $m - q$ distinct agents, the remaining q agents receive all $3q$ element jobs, at most three each and hence exactly three, and by eligibility these are the elements of their triples, which therefore form a cover. Weights and capacities are bounded by 3, so feasibility is strongly NP-complete. For optimization, give all jobs cost 0 on the triple agents and add an agent with capacity $m + 2q$, on which every job has weight 1 and cost 1; the instance is feasible, and its optimal value is 0 if and only if a cover exists. Every agent sees at most two weights, so Theorem 4(iii) and Corollary 1 give $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$, with u^* from Proposition 2(c). □

EC.1.7 Proof of Proposition 3

As in Theorem 1, $Q^{\mathrm{load}}_{\mathcal{V}} = \mathrm{conv}(\{(\mathbf{0}, 0)\} \cup \bigcup_{q \in \mathcal{V}} H_q \times \{1\})$, which grows with $\mathcal{V}$. Every integral vertex of $P_\ell(u^*)$ is a feasible load and every straddle vertex has weight u^*_ℓ, an attainable load (Theorem 2(a)), so $P_\ell(u^*) \subseteq \mathrm{conv}(\bigcup_{q \in \mathcal{V}^*} H_q)$ and $Q(u^*) \subseteq Q^{\mathrm{load}}_{\mathcal{V}^*}$. The first inclusion and its equality condition are Theorem 2(b) with Corollary 3. Points of $H_q \times \{1\}$ with $q \in \mathcal{V}^*$ lie in B^C. If all weights are equal, $Q^{\mathrm{load}}_{\mathcal{V}^*} = \mathrm{conv}(F^\circ) = Q(u^*)$ by Corollary 2(b) and Theorem 4(iii). Otherwise, the point $v = ((d_{\min}/d_k)e_k, 1)$ is a vertex of $Q^{\mathrm{load}}_{\mathcal{V}^*}$ by the proof of Corollary 2(b); every point of $Q(u^*)$ satisfies $|x| \geq y$, because $|x| = \ell \geq 1 = y$ on $P_\ell \times \{1\}$, whereas $|v_x| < 1$. Finally, every point of $Q^{\mathrm{load}}_{\mathcal{V}^*}$ with $y = 1$ has $d(x) \geq \min \mathcal{V}^* > 0$, so $(\mathbf{0}, 1) \notin Q^{\mathrm{load}}_{\mathcal{V}^*}$. For strictness of $Q(u^*) \subseteq Q^{\mathrm{load}}_{\mathcal{V}}$ on a straddle-free block, take weights $(1, 2)$ and $u = 2$: the point $x = (1, \frac{1}{2})$, $y = 1$ lies in $H_2 \times \{1\}$ but violates $|x| = y$, which holds on $Q(u^*)$ since $L = 1$. □

EC.1.8 GAP With an Outside Agent

Consider GAP with agents $r \in R$ and an outside agent o that can serve every job at once at zero fixed cost. For each $r \in R$ let $\mathcal{R}_r$ be a polytope with $(\mathbf{0}, 0) \in \mathcal{R}_r \subseteq [0,1]^{K_r} \times [0,1]$, and let $z_{\mathcal{R}}(c, f) = \min\{\sum_r (c_r^\top x_r + f_r y_r) + c_o^\top x_o : \sum_{r : k \in K_r} x_{rk} + x_{ok} = 1 \ \forall k,\ (x_r, y_r) \in \mathcal{R}_r,\ x_o \in [0,1]^K\}$. The outside agent's capacity does not bind, and its block relaxations all project onto $[0,1]^K$, so with $\mathcal{R}_r$ equal to B^C_r, Q^{load}_r, Q_r or $\mathrm{conv}(F^\circ_r)$, $z_{\mathcal{R}}$ is $z_{\mathrm{LP}}(\mathrm{C})$, z_{load}, $z_{\mathrm{LP}}(\mathrm{CI})$ or z_{DW}. Write $S^\uparrow = S + (\{\mathbf{0}\} \times \mathbb{R}_{\geq 0})$ and proj_x for the projection onto x.

Theorem EC.1. *(a) $z_{\mathcal{R}} = z_{\mathcal{R}'}$ for all (c, f) if and only if $\mathcal{R}_r = \mathcal{R}'_r$ for all r. (b) $z_{\mathcal{R}} = z_{\mathcal{R}'}$ for all c and $f \geq 0$ if and only if $\mathcal{R}^\uparrow_r = \mathcal{R}'^\uparrow_r$ for all r. (c) $z_{\mathcal{R}}(c, 0) = z_{\mathcal{R}'}(c, 0)$ for all c if and only if $\mathrm{proj}_x \mathcal{R}_r = \mathrm{proj}_x \mathcal{R}'_r$ for all r. In (b) and (c), c may be restricted to be nonnegative.*

Proof. The projection of the feasible region onto (x_r, y_r) is $\mathcal{R}_r$: set the other agents to $(\mathbf{0}, 0)$ and let the outside agent take the remaining jobs. The "if" directions follow, in (b) because for $f \geq 0$

the minimum is unchanged when $\mathcal{R}_r$ is replaced by $\mathcal{R}_r^{\uparrow} \cap \{y \le 1\}$. Conversely, a point of one set outside the other is separated by a cost on block r alone, with all other costs 0; the objective then depends only on block r, and its minima over the two regions differ. In (b), $\mathcal{R}_r^{\uparrow}$ recedes along $(\mathbf{0}, 1)$, so the separating cost has $f_r \ge 0$; in (c) the separation is in the space of x_r. Adding a constant to the cost of every job at every agent shifts both values equally. □

Proposition EC.4. *For $\omega \in \mathbb{Z}$ let $\Pi(\omega) = \{\chi \in [0,1]^K : d(\chi) \le \omega\}$. For every block,* $\operatorname{proj}_x Q^{\text{load}}_{\mathcal{V}^*} = \Pi(\hat{u})$, $\operatorname{proj}_x B^C = \Pi(u)$, $\operatorname{proj}_x Q(u^*) \subseteq \Pi(\hat{u})$, *and* $\operatorname{proj}_x \operatorname{conv}(F^\circ) = \operatorname{conv}(\Pi(u) \cap \{0,1\}^K)$. *Moreover, $\Pi(\hat{u})$ is integral if and only if all weights are equal or $u \ge d(K)$.*

Proof. $\operatorname{proj}_x Q^{\text{load}}_{\mathcal{V}^*} = \operatorname{conv}(\{\mathbf{0}\} \cup \bigcup_{q \in \mathcal{V}^*} H_q) \subseteq \Pi(\hat{u})$. Conversely, every vertex of $\Pi(\hat{u})$ is a cube vertex $\mathbf{1}_S$ with $d(S) \le \hat{u}$, hence $\mathbf{0}$ or a feasible load, or a point of a cube edge with $d = \hat{u}$, which lies in $H_{\hat{u}}$. The other projections are immediate, using $d \le u^*_\ell \le \hat{u}$ on $P_\ell(u^*)$. $\Pi(\hat{u})$ has a fractional vertex if and only if some S and $k \notin S$ satisfy $d(S) < \hat{u} < d(S) + d_k$. If all weights equal α, $\hat{u}$ and all loads are multiples of α; if $u \ge d(K)$, $\hat{u} = d(K)$. Otherwise let S^* be a feasible load of weight $\hat{u} < d(K)$, $j \in S^*$ and $k \notin S^*$: if $d_k > d_j$, $S = S^* \setminus \{j\}$ and k give a fractional vertex; if $d_k < d_j$, $S = (S^* \setminus \{j\}) \cup \{k\}$ and j do; if neither occurs for any such pair, all weights are equal. □

Proof of Corollary 5. The inequalities follow from Proposition 3, and the statements from Theorem EC.1(c) and Proposition EC.4. (i) If agent r has a straddle, its vertex projects outside $\operatorname{proj}_x \operatorname{conv}(F_r^\circ)$ by the proof of Theorem 3. (ii) $\operatorname{proj}_x Q^{\text{load}}_{\mathcal{V}^*_r}$ is the projection of the compact relaxation with capacity $\hat{u}_r$, and when neither the coupling nor the objective involves y, both bounds are minima over the same set of x, with or without the outside agent. (iii) If $\Pi_r(\hat{u}_r)$ is integral, all projections coincide. If it has a fractional vertex w, then w has one fractional coordinate, whereas every vertex of $\operatorname{proj}_x Q_r = \operatorname{conv}(\{\mathbf{0}\} \cup \bigcup_\ell P_\ell)$ has none or two (Theorem 2(a)), so $w \notin \operatorname{proj}_x Q_r$. □

With nonnegative fixed costs, as in network design, Theorem EC.1(b) compares upward closures, and the difference in Proposition 3 can be invisible: for an uncapacitated block with weights $(1, 2)$, both upward closures equal $\{(x, y) : x \in [0,1]^2,\ y \ge \max_k x_k\}$.

EC.2 Network Formulations

In network models the blocks are arcs and the items are commodities or terminals. By (1), cardinality indexing acts only through the arc relaxations Q_r; inequalities on single arcs cannot raise the bound above z_{DW}, so further strength requires structure that couples several arcs, such as load conservation in capacity-indexed formulations (Gouveia, 1995; Uchoa et al., 2008), the node copies of Contu et al. (2026), or multi-arc inequalities (Raack et al., 2011; Chouman et al., 2017).

Corollary EC.2. *If $u_r \ge d_r(K_r)$, then $Q_r = \operatorname{conv}(F_r^\circ) = \{(x, y) : 0 \le x \le y\mathbf{1},\ y \le 1,\ |x| \ge y\}$. For the uncapacitated Steiner arborescence problem in its multicommodity-flow formulation (Wong, 1984), the cardinality-indexed bound equals the multicommodity-flow bound.*

Proof. Every level satisfies (i), so $Q_r = \mathrm{conv}(F_r^\circ)$ by Proposition EC.2(b) and Theorem 2; every nonempty subset is a feasible load, and the convex hull of the nonzero cube vertices is $\{z \in [0,1]^{K_r} : |z| \geq 1\}$. In the multicommodity-flow formulation, y appears only in $x_{ka} \leq y_a$ and in the objective with nonnegative costs, so some optimal LP solution has $y_a = \max_k x_{ka} \leq |x_a|$. □

Lemma EC.1. *Consider the multicommodity-flow formulation of capacitated Steiner arborescence: one commodity per terminal, sent from the root; flow conservation; $x_a^k \leq y_a$; capacities on the arcs entering Steiner nodes; and, for every Steiner node j, $\sum_i y_{ij} \leq 1$ and $\sum_i y_{ij} \leq \sum_t y_{jt}$. If all arc costs are nonnegative, Assumption 1 holds.*

Proof. Close an installed empty arc (i_1, j_1). Flow conservation and capacities are unaffected, and the constraints at j_1 remain satisfied. Only node balancing at i_1 can fail, if i_1 is a Steiner node left with no installed outgoing arc but entered by an installed arc (i_2, i_1); that arc is empty, since a commodity on it would have to leave i_1. Close it and repeat at i_2. Each step closes an installed empty arc with nonnegative cost, so the procedure terminates with a feasible solution of no larger cost. □

The node-copy formulation. Contu et al. (2026) split each Steiner node j into copies $j(1), \ldots, j(S_j)$; an arc entering $j(s)$ carries exactly s commodities, an arc from $i(s_1)$ to $j(s_2)$ exists only if $s_1 \geq s_2$, and flow is conserved at every copy. Arcs into copies of j carry $\sum_k d_k z^k_{i(s_1)j(s)} \leq C_j y_{i(s_1)j(s)}$ and $\sum_k z^k_{i(s_1)j(s)} = s\, y_{i(s_1)j(s)}$, the bounds $z_a^k \leq y_a$ hold, and in-degree constraints bound the sum of all arcs into the copies of j by one. Summing over s_1 for fixed s, the aggregated variables $x_{k(i,j)s} = \sum_{s_1} z^k_{i(s_1)j(s)}$ and $y_{(i,j)s} = \sum_{s_1} y_{i(s_1)j(s)}$ of an arc (i,j) into a Steiner node satisfy the constraints of $B_{(i,j)}$ with $\bar{u}_{(i,j)s} = C_j$, so they lie in $Q_{(i,j)}$, and in $\mathrm{conv}(F^\circ_{(i,j)})$ if the arc is straddle-free. This does not extend to the arcs into terminals: on an arc (i,t), which can carry only commodity t, the node-copy formulation requires only $z^t_{it} \leq y_{it}$, whereas LP(CI) implies $x_{t(i,t)} = y_{(i,t)}$ by Corollary EC.2. Our results therefore give no comparison between the bounds of the two formulations.

EC.3 Facets of a Level

Throughout, $X_\ell \neq \emptyset$, $n = |K|$ and $1 \leq \ell \leq n-1$. Feasible ℓ-sets satisfy *exchange closure*: if $S \in X_\ell$, $k \in S$, $o \notin S$ and $d_o \leq d_k$, then $(S \setminus \{k\}) \cup \{o\} \in X_\ell$.

Lemma EC.2. $\dim \mathrm{conv}(X_\ell) = n - 1$ *if and only if every item belongs to some feasible ℓ-set and no item belongs to every one.*

Proof. An item in no feasible set gives $x_k = 0$, and one in every feasible set gives $x_k = 1$. Conversely, index the items with $d_1 \leq \cdots \leq d_n$; the vectors $e_{k+1} - e_k$ span $\{x : |x| = 0\}$, and $e_{k+1} - e_k$ is a difference of two points of X_ℓ whenever some $S \in X_\ell$ contains $k+1$ but not k, by exchange closure. By hypothesis there are $S, T \in X_\ell$ with $k+1 \in S$ and $k \notin T$. If $k \notin S$ or $k+1 \in T$, we are done.

Table EC.1: What the theory establishes for the four cardinality-indexed formulations. "Straddle-free" refers to the level capacities in use.

Formulation	Coupling; Assumption 1	$\bar{u}$	Linear-programming bound	Integer hull of a level
GAP (Murthy and Ransbotham, 2019)	assignment equations; holds	u	z_{DW} if every agent is straddle-free (Cor. 1); necessary with an outside agent (Thm. 3, Cor. 5); u^* is never weaker	Thms. 5–EC.4; hardness Prop. EC.3
Concentrator location (Di Francesco et al., 2021)	assignment equations; holds	u	as for GAP	as for GAP
Network design (Anonymous, 2026a)	flow conservation; holds	u^*	z_{DW} if every arc is straddle-free; dominates load indexing under the same coupling (Prop. 3)	as for GAP
Steiner arborescence (Contu et al., 2026)	flow, in-degree, node balancing; holds (Lemma EC.1)	u	arcs into Steiner nodes project into Q; no comparison of bounds	as for GAP, on arcs into Steiner nodes

If T contains $j > k+1$, then $(T \setminus \{j\}) \cup \{k+1\}$ works. Otherwise $T \subseteq \{1, \dots, k-1\}$, some $\hat{k} \in T$ lies outside S, and $(S \setminus \{k\}) \cup \{\hat{k}\}$ works. □

Proof of Theorem 7. Let $\mathcal{F} = \{S \in X_\ell : |S \cap H| = h\}$. Exchanging an item of $S \in \mathcal{F}$ for a lighter or equally heavy item on the same side of (H, O) preserves membership in $\mathcal{F}$. The left side of (H1) is the weight of the lightest ℓ-set with exactly h items of H that contains a, so (H1) holds if and only if a lies in some member of $\mathcal{F}$, and by exchange closure within H, if and only if every item of H does. Likewise (H2) holds if and only if no item of H lies in every member; (O1) and (O2) are the analogues for O, and when $|O| = 1$ the single item of O lies in every member. Differences of points of $\mathcal{F}$ lie in $L = \{x : x(H) = 0,\ x(O) = 0\}$, of dimension $n-2$. Suppose all four conditions hold, and let k, k' be consecutive items of H in nondecreasing weight order. There are $S, T \in \mathcal{F}$ with $k' \in S$ and $k \notin T$; if $k \notin S$ or $k' \in T$ we have a member containing k' but not k. Otherwise, if $T \cap H$ contains an item j after k', then $(T \setminus \{j\}) \cup \{k'\}$ is such a member; if not, every item of $T \cap H$ precedes k, some $\hat{k} \in T \cap H$ lies outside S, and $(S \setminus \{k\}) \cup \{\hat{k}\}$ is such a member. Exchanging k' for k in it shows $e_{k'} - e_k$ is a difference of points of the face. The same holds within O when $|O| \geq 2$, and the resulting $n-2$ vectors span L, so the face is a facet. Conversely, if a condition fails, the face satisfies $x_a = 0$, $x_b = 1$, or the analogous equation on O, independent of $|x| = \ell$ and $x(H) = h$, so its dimension is at most $n-3$; and if $\ell - h = |O| \geq 2$, every item of O lies in every member. □

Proof of Corollary 7. With $H = C$ and $h = \ell - 1$, (H2) and (O1) are the stated inequalities; (H1) reads $d(C) - d_{c_2} + \min_O d_o \leq u$ with c_2 a heaviest item of $C \setminus \{a\}$ and is implied by (H2), and (O2) is implied by (O1) in the same way. At v^{Aij}, an ℓ-set C has $v^{Aij}(C) > \ell - 1$ only if $A \subseteq C$,

so $C \in \{A \cup \{i\}, A \cup \{j\}\}$, and $A \cup \{j\}$ is feasible. □

If (H1) or (O2) fails, the face satisfies $x_a = 0$ or $x_{b'} = 1$, and the valid strengthenings $x(H)+x_a \leq h$ and $x(H)-x_{b'} \leq h-1$ are not rank-type covers; they can be lifted sequentially or by superadditive functions (Gu et al., 2000; Zeng and Richard, 2011). Theorem 7 does not follow from facet results for the knapsack polytope (Balas, 1975; Chen and Zhu, 2025): that polytope is down-monotone, whereas $\text{conv}(X_\ell)$ is not, and adding the cardinality equation to it does not in general give $\text{conv}(X_\ell)$. The standard route of covers on a restriction followed by lifting (Zeng and Richard, 2011) does not apply directly, because with $|x| = \ell$ and $|C| = \ell$, fixing $x_O = 0$ leaves no feasible point.

EC.4 Three Weights: Levels as Cardinality Slices

This section contains the three-weight results summarized in Section 6.3, with their proofs.

The second geometric mechanism is integrality under slicing.

Lemma EC.3. *Let P be a polytope with integral vertices and ϕ an integral linear form with $|\phi(p) - \phi(q)| \leq 1$ on every edge $[p, q]$ of P. For all integers $\alpha' \leq \beta'$, $P \cap \{\alpha' \leq \phi \leq \beta'\}$ is integral.*

Proof. Let v be a vertex of the section that is not a vertex of P, and let F be the smallest face of P containing v, so v lies in the relative interior of F and $\dim F \geq 1$. At v, the only constraints of the section beyond those defining F are the two slab constraints, so some $\phi = \gamma$ with $\gamma \in \{\alpha', \beta'\}$ is tight, ϕ is not constant on F, and v is a vertex of $F \cap \{\phi = \gamma\}$, which has dimension $\dim F - 1$. Hence $\dim F = 1$, F is an edge $[p, q]$, and $\phi(p) < \gamma < \phi(q)$, which is impossible for integers that differ by at most one. □

For $a \in \mathbb{Z}^N_{>0}$ and an integer κ, let $Y(a, \kappa) = \{y \in \{0,1\}^N : a^\top y \leq \kappa\}$. The polytope $\text{conv}\, Y(a, \kappa)$ has *unit edges* if $||p| - |q|| \leq 1$ on every edge $[p, q]$.

Theorem EC.2. *Let the entries of a take two values $\alpha < \beta$, with $\gamma = \gcd(\alpha, \beta)$. (a) If $\beta - \alpha = \gamma$, then $\text{conv}\, Y(a, \kappa)$ has unit edges for every κ and every number of items of each weight. (b) If $\beta - \alpha > \gamma$, then for α/γ items of weight β, β/γ items of weight α and $\kappa = \alpha\beta/\gamma$, the set of all heavy items and the set of all light items span an edge whose cardinalities differ by $(\beta - \alpha)/\gamma \geq 2$.*

The proof of (a) shows, by four exchange cases, that two feasible sets whose cardinalities differ by at least two have the same midpoint as two other feasible sets, so they are not adjacent by the criterion of Geist and Rodin (1992); for (b), the costs $1/\alpha$ and $1/\beta$ expose the edge; the proofs follow below. With three weights $w_1 > w_2 > w_3$, gaps $g_1 = w_1 - w_2$ and $g_2 = w_2 - w_3$, and class counts $h(S), b(S), c(S)$, (2) reads $d(S) = w_3\ell + (g_1 + g_2)h(S) + g_2 b(S) = w_1\ell - g_1 b(S) - (g_1 + g_2)c(S)$, so eliminating the lightest or the heaviest class turns a level into a two-coefficient knapsack cut by the count.

Theorem EC.3. *Let the weights of a block take three distinct values, and let $X_\ell \neq \emptyset$.*

(a) If $g_1 \mid g_2$, then with $Y = \{y \in \{0,1\}^{K_1 \cup K_2} : (g_1 + g_2)y(K_1) + g_2 y(K_2) \leq u - w_3\ell\}$, $\text{conv}(X_\ell) = \{x \in [0,1]^K : x|_{K_1 \cup K_2} \in \text{conv}(Y),\ |x| = \ell\}$.

(b) If $g_2 \mid g_1$, the analogous statement holds with $K_2 \cup K_3$, complemented variables, and the coefficients g_1 and $g_1 + g_2$.

(c) If neither gap divides the other, both right-hand sides strictly contain $\operatorname{conv}(X_\ell)$ *for suitable class sizes and capacities.*

Proof of (a). The integer points of the right-hand side are X_ℓ. With $g_2 = kg_1$, the coefficients of Y are $(k+1)g_1$ and kg_1, consecutive multiples of g_1, so $\operatorname{conv}(Y)$ has unit edges by Theorem EC.2(a). The edges of $\operatorname{conv}(Y) \times [0,1]^{K_3}$ are products of an edge of one factor with a vertex of the other, so the count changes by at most one along each, and Lemma EC.3 gives integrality. □

Parts (b) and (c) are proved below. By Theorem EC.3, the complete description of two-coefficient knapsack polytopes by Weismantel (1996) transfers to every three-weight level with dividing gaps; separation over these levels is polynomial, as over any level with a fixed number of distinct weights (Section 6). For weights in arithmetic progression, $g_1 = g_2 = g$, the weight of an ℓ-set is $w_2\ell + g(h(S) - c(S))$, and a single family of covers suffices. Let $\mu^* = \max_{S \in X_\ell}(h(S) - c(S))$ and $\rho(t) = \max_{S \in X_\ell}(h(S) + \min\{b(S), t\})$ for $0 \le t \le |K_2|$.

Theorem EC.4. *If the weights of a block take three distinct values in arithmetic progression and $X_\ell \neq \emptyset$, then*

$$\operatorname{conv}(X_\ell) = \left\{x \in [0,1]^K : \ |x| = \ell,\ x(K_1) - x(K_3) \le \mu^*,\ x(K_1) + x(B) \le \rho(|B|) \text{ for all } B \subseteq K_2\right\},$$

where $B = \emptyset$ gives $x(K_1) \le \rho(0)$, and the covers can be separated in $O(n \log n)$ time once ρ has been computed in $O(\ell^2 |K_2|)$ time.

The proof, given below, is constructive: given x, it writes x as a convex combination of feasible sets with at most three count vectors, using a Minkowski-sum description of mixtures of hypersimplices (Schrijver, 2003, Thm. 44.6), or returns a violated cover. Given Theorem EC.3, Theorem EC.4 is equivalent to the case $p = 2$ of Dahl and Foldnes (2003); its own contributions are the constructive proof, right-hand sides that account for the cardinality, and the separation by sorting. Theorems 5, EC.3 and EC.4 together give explicit inequality families, in the space of (CI), that attain z_{DW} on every block whose levels are straddle-free, principal, or have three weights with dividing gaps.

Theorem EC.4 does not extend to four weights in arithmetic progression, for which facets with coefficient two occur. For more than two coefficients, the unit-edge property of Theorem EC.2 is characterized by a finite test on the Graver basis of the row vector of coefficients followed by a one; pairwise consecutive multiples are necessary, sufficient for three coefficients, and not sufficient for four, and every level whose weight differences pass the test is a cardinality slice of a knapsack polytope with one class fewer (Theorems EC.5 and EC.6 in Section EC.5).

Proof of Theorem EC.2. Scale so that $\gamma = 1$. (a) Here $\beta = \alpha + 1$. Two points $p, q \in Y$ are non-adjacent if $p + q = p' + q'$ for some other $p', q' \in Y$ (Geist and Rodin, 1992). Let $P, Q \in Y$ with

$|P| \geq |Q|+2$, $C = P \cap Q$, $D_P = P \setminus Q$, $D_Q = Q \setminus P$ and $r = \kappa - a(C)$, so $a(D_P), a(D_Q) \leq r$ and $|D_P| \geq |D_Q|+2$. It suffices to split $D_P \cup D_Q$ into two parts of weight at most r other than $\{D_P, D_Q\}$. If D_P and D_Q contain items of equal weight, exchange them. If $D_Q = \emptyset$, split D_P into two nonempty parts. Otherwise all items of D_P have one weight and those of D_Q the other. If D_P is heavy and D_Q light, move one heavy item to D_Q: the new part weighs $|D_Q|\alpha+\beta \leq (|D_P|-2)\alpha+\beta < |D_P|\beta \leq r$. If D_P is light and D_Q heavy with $s = |D_Q|$, then $r \geq (s+2)\alpha$; if $s \leq \alpha$, move one light item to D_Q, giving weight $s(\alpha+1)+\alpha \leq r$; if $s \geq \alpha+1$, then $|D_P| \geq \alpha+3 > \beta$, and exchanging β light items with α heavy ones preserves both weights. (b) Give heavy items cost $1/\alpha$ and light items $1/\beta$. A feasible set with t heavy and j light items has value $(t\beta+j\alpha)/(\alpha\beta) \leq 1$, with equality only if $t\beta+j\alpha = \alpha\beta$, which with $\gcd(\alpha,\beta) = 1$ and $t \leq \alpha$ forces $(t,j) \in \{(\alpha,0),(0,\beta)\}$; the optimal face is the edge between the two sets. □

Theorem EC.2 is not an instance of $M^\natural$-convexity: for weights $\{1,2\}$ and capacity 2, the sets $\{a\}$ of weight 2 and $\{i,j\}$ of weight 1 span an edge with direction $e_a - e_i - e_j$. Weismantel (1996) describes the facets of two-weight knapsack polytopes through the Hilbert basis of a cone of exchange vectors; Theorem EC.2 concerns their edges.

Proof of Theorem EC.3(b) and (c). (b) With $y' = \mathbf{1} - x$ on $K_2 \cup K_3$, an ℓ-set is feasible if and only if its complemented restriction lies in $Y' = \{y' : g_1 y'(K_2) + (g_1+g_2)y'(K_3) \leq g_1|K_2| + (g_1+g_2)|K_3| - w_1\ell + u\}$. With $g_1 = mg_2$, the coefficients are consecutive multiples of g_2, so $\mathrm{conv}(Y')$ has unit edges, the form $x(K_1) - |y'|$ changes by at most one along every edge of $[0,1]^{K_1} \times \mathrm{conv}(Y')$, and Lemma EC.3 applies. (c) If $g_1 \nmid g_2$, let $\gamma = \gcd(g_1,g_2)$, $\alpha = g_2$, $\beta = g_1+g_2$; take $|K_1| = \alpha/\gamma$, $|K_2| = \beta/\gamma$, $|K_3| = 1$, $\ell = |K_1|+1$ and $u = \alpha\beta/\gamma + w_3\ell$. By Theorem EC.2(b), $[\mathbf{1}_{K_1}, \mathbf{1}_{K_2}]$ is an edge of $\mathrm{conv}(Y)$ with $|K_1| < \ell < |K_2|$, and $\{|x| = \ell\}$ cuts the edge $[\mathbf{1}_{K_1}, \mathbf{1}_{K_2}] \times \{0\}$ of the product in its relative interior, creating a fractional vertex. The other case is symmetric. For weights $(6,4,1)$ with class sizes $(3,5,1)$, $\ell = 4$ and $u = 19$, exact vertex enumeration finds fractional vertices of both lifts. □

Corollary EC.3. *For every block with three distinct weights one of whose gaps divides the other,* $\mathrm{conv}(X_\ell)$ *is described by the bounds,* $|x| = \ell$*, and the inequalities of Weismantel (1996) for* $\mathrm{conv}(Y)$ *or* $\mathrm{conv}(Y')$*, after dividing their coefficients by* g_1 *or* g_2*; every facet other than a bound is induced by a facet of the two-coefficient polytope. Separation over* $\mathrm{conv}(X_\ell)$ *is polynomial in this case, because optimization over* X_ℓ *is polynomial for a fixed number of distinct weights.*

For $g_2 = kg_1$ with $k \geq 2$, the reduced knapsack has coefficients $(k+1,k)$ and can have facets with coefficients larger than one; the arithmetic progression is the case $(2,1)$, which Theorem EC.4 describes.

Proof of Theorem EC.4. For $0 \leq k \leq N$ let $\Delta(k,N) = \{y \in [0,1]^N : \sum_i y_i = k\}$, and let $s_y(t)$ be the sum of the t largest entries of y.

Lemma EC.4. *Let κ be a random variable with distribution π on $\{0,\dots,N\}$. Then $y \in \sum_k \pi_k \Delta(k,N)$ if and only if $s_y(N) = \mathbb{E}\kappa$ and $s_y(t) \le \mathbb{E}\min\{\kappa, t\}$ for all t. If π is supported on two consecutive integers, $\sum_k \pi_k \Delta(k,N) = \{y \in [0,1]^N : \sum_i y_i = \mathbb{E}\kappa\}$.*

Proof. $f_k(T) = \min\{k, |T|\}$ is nondecreasing and submodular, and $\Delta(k,N)$ is the face of its polymatroid maximizing $y(\{1,\dots,N\})$. The polymatroid of a sum of such functions is the Minkowski sum of their polymatroids (Schrijver, 2003, Thm. 44.6), and the maximizing face of a Minkowski sum is the sum of the maximizing faces, so $\sum_k \pi_k \Delta(k,N)$ is the set of y with $y(T) \le \mathbb{E}\min\{\kappa, |T|\}$ for all T, with equality for the ground set. If π is supported on $\{k_0, k_0+1\}$, the right-hand side is t for $t \le k_0$ and $\mathbb{E}\kappa$ otherwise, so only the bounds and the sum remain. □

Validity. For $S \in X_\ell$, the knapsack inequality is $h(S) - c(S) \le \mu^*$, and $|S \cap (K_1 \cup B)| = h(S) + |S \cap B| \le h(S) + \min\{b(S), |B|\} \le \rho(|B|)$. We show that every x in the right-hand side lies in $\mathrm{conv}(X_\ell)$; integrality of the system follows. Write $n_p = |K_p|$ and $x^{(p)}$ for the restriction of x to K_p.

Step 1: reduction to count distributions. On the level, $d(S) = w_2\ell + g\,(h(S) - c(S))$, so an ℓ-set is feasible if and only if $h(S) - c(S) \le \mu^*$; indeed some feasible set attains $h - c = \mu^*$, and every set with $h - c \le \mu^*$ is no heavier. Hence the set F of count vectors of feasible sets consists of the integer vectors (h,b,c) with $0 \le h \le n_1$, $0 \le b \le n_2$, $0 \le c \le n_3$, $h+b+c = \ell$ and $h - c \le \mu^*$, and X_ℓ contains every ℓ-set whose count vector lies in F. Consequently $\Delta(h,n_1) \times \Delta(b,n_2) \times \Delta(c,n_3) \subseteq \mathrm{conv}(X_\ell)$ for $(h,b,c) \in F$. Suppose that a random vector $(\mathsf{h},\mathsf{b},\mathsf{c})$ with values in F satisfies $x^{(1)} \in \sum_k \Pr[\mathsf{h} = k]\Delta(k,n_1)$, $x^{(2)} \in \sum_k \Pr[\mathsf{b} = k]\Delta(k,n_2)$ and $x^{(3)} \in \sum_k \Pr[\mathsf{c} = k]\Delta(k,n_3)$. Write $x^{(1)} = \sum_k \Pr[\mathsf{h} = k] z^1_k$ with $z^1_k \in \Delta(k,n_1)$, and similarly z^2_k, z^3_k. Then $x = \sum_{(h,b,c)\in F} \Pr[(\mathsf{h},\mathsf{b},\mathsf{c}) = (h,b,c)]\,(z^1_h, z^2_b, z^3_c)$, because summing over two of the counts leaves the marginal of the third, and every term lies in $\mathrm{conv}(X_\ell)$.

Step 2: the two outer classes. Let $\eta = x(K_1)$, $\sigma = x(K_2)$ and $\lambda = x(K_3)$, so $\eta + \sigma + \lambda = \ell$. Write $\eta = \eta_0 + \varphi$ and $\lambda = \lambda_0 + \psi$ with integers η_0, λ_0 and $\varphi, \psi \in [0,1)$, and put $\nu = \ell - \eta_0 - \lambda_0$ and $\theta = \mu^* - \eta_0 + \lambda_0$. The knapsack inequality $\eta - \lambda \le \mu^*$ gives $\theta \ge \varphi - \psi > -1$, so the integer θ is nonnegative. Set $\mathsf{h} = \eta_0 + \epsilon$, $\mathsf{c} = \lambda_0 + \delta$ and $\mathsf{b} = \nu - \epsilon - \delta$, where $(\epsilon,\delta) \in \{0,1\}^2$ is a random vector with $\Pr[\epsilon = 1] = \varphi$ and $\Pr[\delta = 1] = \psi$, chosen below. Then $\mathbb{E}\mathsf{h} = \eta$, $\mathbb{E}\mathsf{c} = \lambda$ and $\mathbb{E}\mathsf{b} = \sigma$. If $\varphi > 0$ then $\eta_0 < \eta \le n_1$, so $\mathsf{h} \in \{0,\dots,n_1\}$; likewise $\mathsf{c} \in \{0,\dots,n_3\}$. Since h and c take at most two consecutive values, Lemma EC.4 gives the conditions of Step 1 for $x^{(1)}$ and $x^{(3)}$. The condition $\mathsf{h} - \mathsf{c} \le \mu^*$ reads $\epsilon - \delta \le \theta$. It remains to choose the joint law of (ϵ,δ) so that $\epsilon - \delta \le \theta$, $\mathsf{b} \in \{0,\dots,n_2\}$, and $x^{(2)}$ satisfies the conditions of Lemma EC.4 for b.

Step 3, case $\theta \ge 1$ or $\varphi = 0$. If $\varphi + \psi \le 1$, let (ϵ,δ) equal $(1,0)$, $(0,1)$ and $(0,0)$ with probabilities φ, ψ and $1 - \varphi - \psi$; if $\varphi + \psi > 1$, let it equal $(1,1)$, $(1,0)$ and $(0,1)$ with probabilities $\varphi + \psi - 1$, $1 - \psi$ and $1 - \varphi$. Both laws have the required marginals, and $\epsilon + \delta \in \{\lfloor \varphi + \psi \rfloor, \lceil \varphi + \psi \rceil\}$. If $\theta \ge 1$, then $\epsilon - \delta \le 1 \le \theta$; if $\varphi = 0$, then $\epsilon = 0$. The variable b takes the single value σ, if $\varphi + \psi$ is an integer, or the two values $\lfloor \sigma \rfloor$ and $\lceil \sigma \rceil$ otherwise; since $0 \le \sigma \le n_2$, these lie in $\{0,\dots,n_2\}$, and Lemma EC.4 applies.

Step 4, case $\theta = 0$ and $\varphi > 0$. Then $\psi \geq \varphi > 0$. Let (ϵ, δ) equal $(1,1)$, $(0,1)$ and $(0,0)$ with probabilities φ, $\psi - \varphi$ and $1 - \psi$. The marginals are correct, $\epsilon \leq \delta$ gives $\epsilon - \delta \leq 0 = \theta$, and b equals $\nu - 2$, $\nu - 1$ and ν with these probabilities.

(4a) *Range of $t^* = \nu - 1$.* From $\sigma = \nu - \varphi - \psi \geq 0$ and $\varphi > 0$ we get $\nu > 0$, so $t^* \geq 0$; from $\sigma \leq n_2$ and $\varphi + \psi < 2$ we get $\nu < n_2 + 2$, so $t^* \leq n_2$.

(4b) $\rho(t^*) \leq \eta_0 + t^*$. Let $(h, b, c) \in F$. Since $\theta = 0$, $h - c \leq \mu^* = \eta_0 - \lambda_0$, and substituting $c = \ell - h - b$ gives $2h + b \leq 2\eta_0 + \nu$. If $b \geq t^*$, then $2h \leq 2\eta_0 + 1$, so $h \leq \eta_0$ and $h + \min\{b, t^*\} \leq \eta_0 + t^*$. If $b \leq t^* - 1$, then $h + b \leq \eta_0 + (\nu + b)/2 \leq \eta_0 + (\nu + t^* - 1)/2 = \eta_0 + t^*$.

(4c) *The cover.* Let $B \subseteq K_2$ consist of t^* items with the largest values of x. The cover $x(K_1) + x(B) \leq \rho(t^*)$ and (4b) give $\eta + s_{x^{(2)}}(t^*) \leq \eta_0 + t^*$, that is,

$$s_{x^{(2)}}(t^*) \leq t^* - \varphi. \tag{EC.1}$$

If $t^* = 0$, (EC.1) reads $0 \leq -\varphi$, contradicting $\varphi > 0$. If $t^* = n_2$, it reads $\sigma \leq t^* - \varphi$, whereas $\sigma = t^* + 1 - \varphi - \psi > t^* - \varphi$ because $\psi < 1$. Hence $1 \leq t^* \leq n_2 - 1$, that is, $2 \leq \nu \leq n_2$, and the values $\nu - 2, \nu - 1, \nu$ of b lie in $\{0, \ldots, n_2\}$.

(4d) *Lemma EC.4 for b.* For $t \leq \nu - 2$, $\mathbb{E}\min\{\mathsf{b}, t\} = t \geq s_{x^{(2)}}(t)$ because $x \leq \mathbf{1}$. For $t \geq \nu$, $\mathbb{E}\min\{\mathsf{b}, t\} = \sigma \geq s_{x^{(2)}}(t)$ because $x \geq 0$, with equality at $t = n_2$. For $t = t^*$, $\mathbb{E}\min\{\mathsf{b}, t^*\} = (\nu - 2)\varphi + (\nu - 1)(1 - \varphi) = t^* - \varphi$, and (EC.1) applies. These are all integers $0 \leq t \leq n_2$, so the conditions hold.

Steps 3 and 4 exhaust the cases, because $\theta \geq 0$. □

Lemma EC.4 needs no separate nonnegativity condition: $s_y(N-1) \leq \mathbb{E}\min\{\kappa, N-1\} \leq \mathbb{E}\kappa = s_y(N)$ forces the smallest entry of y to be nonnegative, and $s_y(1) \leq 1$ bounds the largest by one. The construction was also executed in exact rational arithmetic on random points of the right-hand side, including 1,070 points in the case of Step 4, and every intermediate claim held.

The proof writes x as a combination with at most three count vectors or returns the violated cover with B the $\nu - 1$ largest entries of $x^{(2)}$; sorting separates the whole family in $O(n \log n)$ time once ρ is computed from the $O(\ell^2)$ count pairs. With unequal gaps, the two-for-two exchange changes the weight, and facets with larger coefficients appear, as $x_2 + x_3 + x_4 + 2x_5 \leq 2$ for weights $(1, 2, 2, 2, 4)$, $u = 5$ and $\ell = 2$.

Corollary EC.4. *In the setting of Theorem EC.4 with $\dim \mathrm{conv}(X_\ell) = n - 1$, for each t either every cover $x(K_1) + x(B) \leq \rho(t)$ with $|B| = t$ defines a facet or none does, as decided by Theorem 7; the facets other than the bounds and the knapsack inequality are these covers for at most $|K_2| + 1$ values of t, found in $O(n)$ time once ρ is known.*

Corollary EC.5. *For $Y = \{y \in \{0,1\}^{N_1 \cup N_2} : y(N_1) + 2y(N_2) \leq \kappa\}$ with $N_1, N_2 \neq \emptyset$, $\kappa^* = \max_Y(y(N_1) + 2y(N_2))$ and $\rho_0(t) = \max_Y(y(N_2) + \min\{y(N_1), t\})$, $\mathrm{conv}(Y) = \{y \in [0,1]^N : y(N_1) + 2y(N_2) \leq \kappa^*,\ y(N_2) + y(S) \leq \rho_0(|S|)\ \forall S \subseteq N_1\}$.*

Proof. Embed Y as a level of a block with weights $3, 2, 1$: $K_1 = N_2$, $K_2 = N_1$, $|N|$ items of weight 1, $\ell = |N|$ and $u = \kappa + \ell$. Then $\mathrm{conv}(Y)$ is the projection of $\mathrm{conv}(X_\ell)$ onto N, and projecting the

description of Theorem EC.4 by substituting $x(K_3) = \ell - |y|$ gives the system, with $\mu^* = \kappa^* - \ell$ and $\rho = \rho_0$. □

This agrees with the case $p = 2$ of Dahl and Foldnes (2003) and reproves it without total dual integrality.

EC.5 More Than Two Coefficients

Call a finite set W of positive integers *unit-edged* if every knapsack polytope with coefficients in W, for every number of items of each coefficient and every capacity, has unit edges. For $u, v \in \mathbb{Z}^m$ write $u \sqsubseteq v$ if $u_i v_i \geq 0$ and $|u_i| \leq |v_i|$ for all i; the *Graver basis* of an integer matrix is the set of $\sqsubseteq$-minimal nonzero integer vectors in its kernel (Graver, 1975; De Loera et al., 2013). For $1 \times n$ matrices, Graver bases are the primitive partition identities studied by Diaconis et al. (1996).

Proposition EC.5. *(a) W is unit-edged if and only if every cardinality section* $\operatorname{conv} Y \cap \{|y| = \lambda\}$ *of every knapsack polytope with coefficients in W is integral. (b) W is unit-edged if, for all count vectors $p, q \in \mathbb{Z}^W_{\geq 0}$ with disjoint supports and $|p| \geq |q| + 2$, some $x \leq p + q$ with $x \notin \{p, q\}$ has weight between $d(q)$ and $d(p)$. (c) If W is unit-edged, every two of its elements are consecutive multiples of their greatest common divisor.*

Proof. (a) One direction is Lemma EC.3; conversely, an edge with $|P| \geq |Q| + 2$ meets an intermediate section in a fractional vertex. (b) On an edge $[\mathbf{1}_P, \mathbf{1}_Q]$, $P \setminus Q$ and $Q \setminus P$ share no weight, since an exchange would give another pair with the same midpoint; a vector x as stated splits $P \triangle Q$ into two other parts of weight at most $\max\{d(p), d(q)\}$, giving another pair of feasible sets with the same midpoint. (c) Fixing the items of the other coefficients at 0 gives a face to which Theorem EC.2(b) applies. □

Theorem EC.5. *Let $(w\ 1)$ be the row vector of the elements of W followed by a 1. (a) W is unit-edged if and only if every Graver element (z, ζ) of $(w\ 1)$ satisfies $|\sum_\delta z_\delta| \leq 1$. (b) If the weights of a block take $q \geq 2$ values $w_1 > \cdots > w_q$ and $\{w_p - w_q : p < q\}$ is unit-edged, then* $\operatorname{conv}(X_\ell) = \{x \in [0,1]^K : x|_{K \setminus K_q} \in \operatorname{conv}(Y),\ |x| = \ell\}$ *with* $Y = \{y \in \{0,1\}^{K \setminus K_q} : \sum_{p<q}(w_p - w_q) y(K_p) \leq u - w_q \ell\}$.

Proof. (a) If every Graver element satisfies the bound, let p, q be as in Proposition EC.5(b) with $T = d(p) - d(q)$. The kernel vector $(p - q, -T)$ is a sum of Graver elements $\sqsubseteq$ it, none with $z = 0$, whose counts add to $|p| - |q| \geq 2$ and are at most one each, so there are at least two; one of them, $(v, -\tau)$ with $v = a - b$, $0 \leq a \leq p$, $0 \leq b \leq q$ and τ between 0 and T, gives $x = p - a + b$ as in Proposition EC.5(b). Conversely, let $(z, -\tau)$ be a Graver element with $\sum z \geq 2$, $p = z^+$, $q = z^-$. Take $p_\delta + q_\delta$ items of each weight and capacity $\max\{d(p), d(q)\}$; the sets P, Q of the classes in the supports of p, q are the only feasible sets with these count vectors. If $[\mathbf{1}_P, \mathbf{1}_Q]$ were not an edge, the midpoint $\frac{1}{2}\mathbf{1}$ would be a convex combination $\lambda_P \mathbf{1}_P + \lambda_Q \mathbf{1}_Q + \sum_X \mu_X \mathbf{1}_X$ of feasible sets in which some $X \notin \{P, Q\}$ has $\mu_X > 0$. Since P and Q partition the items, comparing coordinates gives $\lambda_P, \lambda_Q \leq \frac{1}{2}$, so the sets $X \notin \{P, Q\}$ have average weight $\big((\frac{1}{2} -$

$\lambda_P)d(p) + (\frac{1}{2} - \lambda_Q)d(q)\big)/(1 - \lambda_P - \lambda_Q) \geq \min\{d(p), d(q)\}$, and one of them, with count vector x, has $\min\{d(p), d(q)\} \leq d(x) \leq \max\{d(p), d(q)\}$. With $a = p - x|_{\operatorname{supp} p}$ and $b = x|_{\operatorname{supp} q}$, we have $a - b \sqsubseteq z$, $a - b \notin \{0, z\}$ because $X \notin \{P, Q\}$, and $d(a) - d(b) = d(p) - d(x)$ lies between 0 and $d(p) - d(q)$; so $(a - b, -(d(a) - d(b)))$ is a kernel element $\sqsubseteq (z, -\tau)$ other than 0 and $(z, -\tau)$, contradicting minimality. So $[\mathbf{1}_P, \mathbf{1}_Q]$ is an edge with a count change of at least two. (b) On the level, $d(x) = w_q \ell + \sum_{p<q}(w_p - w_q)x(K_p)$, and Lemma EC.3 applies to $\operatorname{conv}(Y) \times [0,1]^{K_q}$. □

Graver bases are finite and computable (4ti2 team, 2018), so part (a) is a finite test, though not a polynomial one. The pairwise condition is not sufficient: $\{30, 36, 40, 45\}$ satisfies it, but $(9, -4, -2, -1, -1)$ is a Graver element of (30 36 40 45 1); nine items of weight 30 and the seven others, with capacity 270, span an edge with a count change of two.

Theorem EC.6. *A set of three positive integers is unit-edged if and only if every two of its elements are consecutive multiples of their greatest common divisor.*

Proof. Necessity is Proposition EC.5(c). For sufficiency, let W have three elements, every two of which are consecutive multiples of their greatest common divisor. Multiplying all coefficients and the capacity by a common factor does not change the knapsack polytope, so we may assume $W = \{a, b, c\}$ with $a < b < c$ and $\gcd(a, b, c) = 1$. Put

$$e = b - a, \qquad g = c - a, \qquad h = c - b = g - e.$$

The pairwise condition says $\gcd(a, b) = e$, $\gcd(a, c) = g$ and $\gcd(b, c) = h$, so $e \mid a$, $g \mid a$ and $h \mid b$, and hence $e \mid b$, $g \mid c$ and $h \mid c$. A common divisor of e and g divides a, $b = a + e$ and $c = a + g$, so $\gcd(e, g) = 1$, and $\gcd(g, h) = \gcd(g, e) = 1$. Therefore $eg \mid a$ and $gh \mid c$; write $a = meg$ with $m \geq 1$. For $u \in \mathbb{Z}^3$, indexed by (a, b, c), let $|u| = u_a + u_b + u_c$ and $d(u) = au_a + bu_b + cu_c$. Since $a = c - g$ and $b = c - h$,

$$d(u) = c\,|u| - g\,u_a - h\,u_b. \tag{EC.T1}$$

Reduction. By Theorem EC.5(a), it suffices to show that no Graver element (z, ζ) of $(a\ b\ c\ 1)$ has $|z| \geq 2$; the case $|z| \leq -2$ follows by negation. Here $\zeta = -d(z)$. A vector (u, η) in the kernel with $(u, \eta) \sqsubseteq (z, -d(z))$ and $(u, \eta) \notin \{0, (z, -d(z))\}$ has $\eta = -d(u)$, $u \notin \{0, z\}$, and $d(u)$ between 0 and $d(z)$, both included; conversely, every such u gives such a vector. Call $u \in \mathbb{Z}^3$ a *witness* for z if $u \sqsubseteq z$, $u \notin \{0, z\}$, and $d(u)$ lies between 0 and $d(z)$. Then $(z, -d(z))$ is a Graver element exactly when z has no witness, and we show that every z with $N := |z| \geq 2$ has one. Let $\Delta = d(z)$. Every vector u below has $|u| = 1 < N$, so $u \neq z$ holds automatically, and $u \neq 0$ is clear. Moreover 0 lies between 0 and Δ whatever the sign of Δ, so every $u \sqsubseteq z$ with $d(u) = 0$ is a witness.

The cases below are summarized in the following table. Each row lists the sign pattern of z in the order (a, b, c), the witness, its weight, and the inequalities that give $u \sqsubseteq z$; in Cases 2–5 the single item e_a (in Case 2, e_b) is tried first, and the listed witnesses are used only when it fails.

Case	Signs	Witness u	$d(u)$	Why $u \sqsubseteq z$
–	one sign	e_δ	δ	$\Delta \geq \delta$
–	two nonzero entries	Thms. EC.2(a), EC.5(a)		
1	$(-,+,+)$, $(-,-,+)$	e_{δ^+}	δ^+	δ^+ the lightest positive class; $\Delta \geq 2\delta^+$
2	$(-,+,-)$	$(0,\, c/h,\, -b/h)$	0	$\gamma > b/h$, $\beta \geq \gamma + 2$
3	$(+,-,+)$	$(b/e,\, -a/e,\, 0)$	0	$\alpha > b/e + 1$, $\beta > a/e$
4(i)	$(+,-,-)$	$(1 + gi + e(m-i),\, -gi,\, -e(m-i))$	0	$gi \leq \beta$, $e(m-i) \leq \gamma$, and $\alpha \geq \beta + \gamma + 2$
4(ii)		$(1+g,\, -g,\, 0)$ or $(1+e,\, 0,\, -e)$	$a - eg$	$\beta \geq g$ or $\gamma \geq e$
4(iii)		$(1+\beta,\, -\beta,\, 0)$, here $m = 1$	$a - e\beta$	$\alpha \geq \beta + 2$
5	$(+,+,-)$	$(hj,\, c/h - gj,\, -(c/h - ej - 1))$	0	$hj \leq \alpha$, $c/h - gj \leq \beta$; then $r_j < \gamma$
5, residual		$(\rho,\, 0,\, -(\rho - 1))$, here $N = 2$	$c - g\rho$	$\rho \leq \alpha$, $\rho - 1 \leq \alpha + \beta - 2 = \gamma$

Every witness other than a single item has $|u| = 1$.

If z has a single nonzero entry, then $z = Ne_\delta$ and e_δ is a witness. If z has exactly two nonzero entries, with coefficients δ, δ', then $(z, -\Delta)$, restricted to these coordinates and the last one, would be a Graver element of $(\delta\ \delta'\ 1)$, whose elements change the count by at most one by Theorems EC.2(a) and EC.5(a). So let all three entries be nonzero. If they all have the same sign, $|z| \geq 2$ makes them positive, and e_a is a witness. Six sign patterns remain.

Case 1: every positive entry belongs to a coefficient larger than every negative one, that is, the patterns $(-,+,+)$ and $(-,-,+)$. Let δ^+ be the smallest coefficient with a positive entry, and let $p = z^+$ and $q = z^-$, so that $|p| \geq |q| + 2$. Then $\Delta \geq \delta^+|p| - \delta^+|q| \geq 2\delta^+$, and e_{δ^+} is a witness.

Case 2: $z = (-\alpha, \beta, -\gamma)$ with $\alpha, \beta, \gamma \geq 1$ and $\beta \geq \alpha + \gamma + 2$. If $\Delta \geq b$, then e_b is a witness. Otherwise $b > \Delta \geq (\alpha + \gamma + 2)b - \alpha a - \gamma c = 2b + \alpha e - \gamma h$, so $\gamma h > b$, that is, $\gamma > b/h$; and $\beta \geq \gamma + 2 > c/h = b/h + 1$. Hence $u = (0,\, c/h,\, -b/h) \sqsubseteq z$, with $|u| = 1$ and $d(u) = (bc - cb)/h = 0$, so u is a witness.

Case 3: $z = (\alpha, -\beta, \gamma)$ with $\alpha, \beta, \gamma \geq 1$ and $\beta \leq \alpha + \gamma - 2$. If $\Delta \geq a$, then e_a is a witness. Otherwise $\beta b > (\alpha - 1)a + \gamma c$, and $\beta \leq \alpha + \gamma - 2$ gives $(\alpha + \gamma - 2)b > (\alpha - 1)a + \gamma c$, that is, $(\alpha - 1)e > \gamma h + b > b$. Hence $\alpha > b/e + 1$, and then $\beta b > (\alpha - 1)a > (b/e)\, a$ gives $\beta > a/e$. So $u = (b/e,\, -a/e,\, 0) \sqsubseteq z$, with $|u| = 1$ and $d(u) = 0$, is a witness.

Case 4: $z = (\alpha, -\beta, -\gamma)$ with $\alpha, \beta, \gamma \geq 1$ and $\alpha \geq \beta + \gamma + 2$. If $\Delta \geq a$, then e_a is a witness, so let $\Delta < a$. For $0 \leq \beta' \leq \beta$ and $0 \leq \gamma' \leq \gamma$, the vector $u = (1 + \beta' + \gamma', -\beta', -\gamma')$ satisfies $u \sqsubseteq z$, since $1 + \beta' + \gamma' \leq \alpha$, and $d(u) = a - e\beta' - g\gamma'$. Put $R = e\beta + g\gamma$; then $\Delta \geq (\beta + \gamma + 2)a - \beta b - \gamma c = 2a - R$.

(i) If some $i \in \{0, \ldots, m\}$ has $gi \leq \beta$ and $e(m - i) \leq \gamma$, then $(\beta', \gamma') = (gi, e(m - i))$ gives $d(u) = a - meg = 0$, a witness.

(ii) Otherwise let $i_0 = \lfloor \beta/g \rfloor$. Then $i_0 < m$, since otherwise $i = m$ would qualify, and $e(m - i_0) > \gamma$, since i_0 does not qualify. Hence $R < eg(i_0 + 1) + ge(m - i_0) = a + eg \leq 2a$, so $\Delta \geq 2a - R > a - eg \geq 0$. If $\beta \geq g$, take $(\beta', \gamma') = (g, 0)$; if $\gamma \geq e$, take $(0, e)$. In both cases $d(u) = a - eg$, and $0 \leq a - eg < \Delta$.

(iii) Otherwise $\beta < g$ and $\gamma < e$, so $R < 2eg$, while $\Delta < a$ and $\Delta \geq 2a - R$ give $R > a = meg$.

Hence $m = 1$ and $a = eg$. Take $(\beta', \gamma') = (\beta, 0)$: then $d(u) = a - e\beta > 0$, and $\Delta \geq 2a - R = (a - e\beta) + (a - g\gamma) > a - e\beta$, because $g\gamma < ge = a$.

Case 5: $z = (\alpha, \beta, -\gamma)$ with $\alpha, \beta, \gamma \geq 1$ and $\gamma = \alpha + \beta - N$. If $\Delta \geq a$, then e_a is a witness, so let $\Delta < a$. By (EC.T1), $\Delta = cN - F$ with $F = g\alpha + h\beta$, so

$$F > c(N-1) + g \geq c + g. \tag{EC.T2}$$

Let $J = c/(gh)$. For $j = 0, \ldots, J$, put $s_j = hj$, $t_j = c/h - gj \geq 0$ and $r_j = s_j + t_j - 1$. Then $s_j + t_j \geq 1$, since $t_j \geq g$ for $j < J$ and $s_J = c/g$, so $r_j \geq 0$; moreover $gs_j + ht_j = c$, and $u_j = (s_j, t_j, -r_j)$ has $|u_j| = 1$ and, by (EC.T1), $d(u_j) = 0$. We claim that u_j is a witness whenever $s_j \leq \alpha$ and $t_j \leq \beta$; it remains to check $r_j \leq \gamma$. From $d(u_j) = 0$, $r_j c = s_j a + t_j b$, and from $\Delta = \alpha a + \beta b - \gamma c < a$, $\gamma c > (\alpha - 1)a + \beta b$. Since $gs_j + ht_j = c < F$, we have $(s_j, t_j) \neq (\alpha, \beta)$, so $s_j \leq \alpha - 1$ or $t_j \leq \beta - 1$. In either case $s_j a + t_j b \leq (\alpha - 1)a + \beta b$, using $a < b$ in the second. Hence $r_j c < \gamma c$.

If $\alpha \geq c/g = s_J$, then $j = J$ qualifies, since $t_J = 0$. Otherwise let $j_0 = \lfloor \alpha/h \rfloor < J$ and $\rho = \alpha - hj_0 \in \{0, \ldots, h-1\}$. If $t_{j_0} \leq \beta$, then j_0 qualifies. Otherwise $c/h - gj_0 > \beta$, that is, $c > h\beta + g(\alpha - \rho)$, so $F < c + g\rho$. With (EC.T2) this forces $\rho \geq 2$, and also $F < c + gh \leq 2c$, so $N = 2$ by (EC.T2), and $\Delta = 2c - F > c - g\rho$. Take $u = (\rho, 0, -(\rho - 1))$. It satisfies $u \sqsubseteq z$, because $\rho \leq \alpha$ and $\rho - 1 \leq \alpha + \beta - 2 = \gamma$. It has $|u| = 1$ and, by (EC.T1), $d(u) = c - g\rho$, which lies strictly between 0 and Δ because $g\rho < gh \leq c$. So u is a witness.

The sign patterns with three nonzero entries and both signs present are $(-,+,+)$ and $(-,-,+)$ (Case 1), $(-,+,-)$ (Case 2), $(+,-,+)$ (Case 3), $(+,-,-)$ (Case 4) and $(+,+,-)$ (Case 5). This proves the theorem. □

The case analysis was run as an explicit procedure on all 143 admissible sets with $c \leq 100$ and every $z \in [-20, 20]^3$ with three nonzero entries and $|z| \geq 2$, and independently the Graver bases (computed with 4ti2) of all 2,342 admissible sets with $c \leq 400$ change the count by at most one (Table EC.3).

EC.6 An A Posteriori Bound on the Loss

For a block with costs $\bar{c}$, let $\Delta(\bar{c}) = \max\{t_{Aij}(\bar{c}_j - \bar{c}_i)^+ : (A, i, j) \text{ a straddle}\}$, with $\Delta = 0$ on straddle-free blocks. Give an integral vertex of Q the value 0 and a vertex $(v^{Aij}, 1)$ the value $t_{Aij}(\bar{c}_j - \bar{c}_i)^+$, and let $\Delta^*(\bar{c}, \bar{f})$ be the smallest value of a vertex of Q minimizing $\bar{c}^\top x + \bar{f} y$.

Proposition EC.6. *Suppose Γ involves only (x, y), and let λ^* be an optimal multiplier of the Lagrangian dual of Γ over $\prod_r \mathrm{conv}(F_r^\circ)$, with reduced costs $\bar{c}_r$ and $\bar{f}_r$. Then $0 \leq z_{\mathrm{DW}} - z_{\mathrm{LP}}(\mathrm{CI}) \leq \sum_r \Delta_r^*(\bar{c}_r, \bar{f}_r) \leq \sum_r \Delta_r(\bar{c}_r)$. In particular, $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$ if, in every block with a straddle, $\bar{c}_{ri} \geq \bar{c}_{rj}$ whenever $d_{ri} > d_{rj}$.*

Proof. Let ϕ_S be the minimum of $\bar{c}^\top x + \bar{f} y$ over S. If the chosen minimizing vertex of Q is integral, $\phi_Q = \phi_{\mathrm{conv}(F^\circ)}$. If it is $(v^{Aij}, 1)$, then $A \cup \{j\}$ is a feasible load and $\bar{c}^\top v^{Aij} = \bar{c}(A \cup \{j\}) + t_{Aij}(\bar{c}_i - \bar{c}_j) \geq$

$\bar{c}(A\cup\{j\})-t_{Aij}(\bar{c}_j-\bar{c}_i)^+$. Hence $\phi_{\mathrm{conv}(F^\circ)}-\phi_Q \le \Delta^* \le \Delta$. The Lagrangian functions of LP(CI) and of the DW relaxation differ only in these block terms, so $z_{\mathrm{LP}}(\mathrm{CI}) \ge \mathcal{L}_Q(\lambda^*) \ge z_{\mathrm{DW}} - \sum_r \Delta_r^*$. □

The bound requires the DW multipliers, so it explains the loss a posteriori but cannot certify it in advance; in the GAP experiments, $\sum_r \Delta_r^*$ exceeds the observed gap by a median factor between 2.4 and 8.8, and $\sum_r \Delta_r$ by a factor between about 16 and 128.

EC.7 Computational Details

Solvers and tolerances. All programs are solved in Python with HiGHS through SciPy 1.18 with default tolerances unless stated. Optimal values are obtained from the compact integer formulation with a relative gap tolerance of 10^{-6}; all costs are integers and $z^* \le 4742$ in Tables 3–7, so every reported z^* is proven optimal. Every integer program terminated with status optimal within its time limit.

Column generation. The restricted master contains the coupling constraints, one convexity row per block, and artificial variables with cost 10^5 (GAP) or 10^6 (network design). Columns are nonempty feasible loads of single blocks, priced by a knapsack dynamic program over $\{0,\dots,u_r\}$, until no column has reduced cost below -10^{-7}. All runs converged, network instances within 30 iterations, and the artificial variables summed to at most 10^{-9} in every final master. Exact level capacities are computed by a dynamic program over count and load, or in closed form on two-valued blocks (Proposition 2(c)).

GAP instances. With $U[a,b]$ the uniform distribution on $\{a,\dots,b\}$ and $U(0,1)$ the continuous uniform distribution: type C has $a_{ij}\sim U[5,25]$ and $c_{ij}\sim U[10,50]$; type D has $a_{ij}\sim U[1,100]$ and $c_{ij}=111-a_{ij}+U[-10,10]$; type E has $a_{ij}=\lfloor 1-10\ln U(0,1)\rfloor$ and $c_{ij}=\lfloor 1000/a_{ij}-10\,U(0,1)\rfloor$; in all three, $b_i=\lfloor 0.8\sum_j a_{ij}/m\rfloor$, and for type E additionally $b_i\ge\max_j a_{ij}$. Class W12 has $a_{ij}\in\{1,2\}$, $c_{ij}\sim U[10,50]$ and b_i as in type C; class W58 has $a_{ij}\in\{5,8\}$ and otherwise the same rules; classes Ss have $a_{ij}\sim U[10,10+s]$, $c_{ij}\sim U[10,50]$ and $b_i=\lfloor\sum_j a_{ij}/m\rfloor$. Seeds: 11 for C, D, E, W12 and W58, and 23 for Ss. The generator discards infeasible instances and instances not solved within 60 s; none was discarded.

Network design instances. Directed graphs with 10 nodes contain a Hamiltonian cycle in both directions and every further ordered pair as an arc with probability 0.3. There are 12 commodities with random distinct origins and destinations. All arcs have capacity $u=\max\{\mathrm{round}(2.5\bar{d}),\max_k d_k\}$ with $\bar{d}$ the mean demand; fixed costs are $U[50,150]$ and unit routing costs $U[1,5]$. The compact relaxation includes $x_{ka}\le y_a$. The DW master has flow conservation for every commodity and node and one convexity row per arc. Seed 31; two class-H instances were infeasible and were replaced. All 80 integer programs were solved to proven optimality.

Level-hull experiments. The network instances of Section 7.3 are regenerated with the original generator and seed 31, and each regenerated instance is matched to the recorded value of $z_{\mathrm{LP}}(\mathrm{C})$; all 60 instances matched. Levels with a straddle are identified by enumerating the $(\ell-1)$-sets under exact level capacities. Principal levels and their right-hand sides $C_p(T^*)$ are computed by Proposition 4. Exact separation over $\mathrm{conv}(X_\ell)$ at $\hat{x} = x_{\cdot\ell}/y_\ell$ solves $\max\{\alpha^\top \hat{x} - \beta : \alpha^\top \mathbf{1}_S \le \beta \ (S \in X_\ell),\ -\mathbf{1} \le \alpha \le \mathbf{1}\}$ over the explicitly enumerated feasible ℓ-sets (12 commodities, at most 5 levels). Every cut is added in homogenized form $\alpha^\top x_{\cdot\ell} \le \beta y_\ell$, and the loop stops when no cut is violated by more than 10^{-7}.

Benchmark instances. The GAP benchmark files are the OR-Library instances of types A–D and the type-E instances and the instances with 400 or more jobs made available by Yagiura et al. (2006), all in OR-Library format: 60 small type-C instances (the classical sets gap1–gap12), 30 instances of types A–E with 5, 10 or 20 agents and 100 or 200 jobs, and 27 instances of types C, D and E with 400 to 1,600 jobs. We use the copies distributed with the open-source package AssignmentProblems.jl, together with its table of optimal and best-known minimization values. Items heavier than an agent's capacity are excluded from that agent's block. The census uses only Theorem 4, Proposition 2, Corollary EC.1 and Proposition 4, so no exact level capacity is computed. On 3,000 random blocks with up to 8 items (9,964 levels), the implementation agreed with brute-force enumeration of straddles and of the principal property on every level, and its test after fixing items at bound agreed on all 9,985 levels of a second sample. The bound comparison computes the exact level capacities by a dynamic program over count and load. It computes z_{DW} by column generation with a warm-started master problem in HiGHS, exact knapsack pricing and the dual smoothing of Wentges (1997) with smoothing factor 0.8. The procedure stops only when no column has reduced cost below -10^{-7} at the true duals; on the 23 instances run with smoothing, the best Lagrangian bound found during pricing agreed with the final master value, and the other 7 were run without smoothing; on the 4 instances run both ways, the two values agreed exactly.

Exact verification. Table EC.3 lists the exact-arithmetic checks of the results of Section 6 and Section EC.5, computed with Normaliz (Bruns et al., 2024) through PyNormaliz, with edges determined combinatorially from facet incidences and every reported vertex certified by a rank test. An implementation of the construction in the proof of Theorem 4, with every intermediate claim checked as an assertion, returned a valid straddle on every block with at most six items and weights in $\{1,\dots,5\}$ and on 20,000 random blocks with up to nine items, exercising Step 1, Case A and the two sub-cases of Case B ($\gamma < \alpha$ and $\gamma > \beta$) 74,637, 57,824, 2,423 and 5,507 times. An independent brute-force comparison of the statement of Theorem 4 on about 8.9 million levels found no exception.

Census. The census of Table EC.2 enumerates all facets of $\mathrm{conv}(X_\ell)$ in exact arithmetic for every block with 3 to 7 items and weights in $\{1,\dots,8\}$, every capacity with $\max_k d_k \le u < d(K)$ and every level $2 \le \ell \le n-1$. Lemma EC.2 agreed with the computed dimension on all 424,369

Table EC.2: Exhaustive census of $\mathrm{conv}(X_\ell)$ for blocks with 3 to 7 items, weights in $\{1,\dots,8\}$, every capacity with $\max_k d_k \le u < d(K)$ and every level $2 \le \ell \le n-1$, computed in exact arithmetic; full-dimensional levels only. Principal: levels covered by Theorem 6. Last two columns: levels whose integer hull is described by the bounds, the knapsack inequality at u^*_ℓ, and level covers or rank-type covers, respectively.

Distinct weights	Levels	Principal	With level covers	With rank-type covers
1	1,260	100%	100%	100%
2	20,580	100%	100%	100%
3	88,354	92%	82%	97%
4	130,396	83%	76%	91%
5	70,000	75%	73%	86%
6	12,654	68%	71%	81%
7	549	64%	70%	76%

nonempty levels, and Corollary 7 with the facet status of all 677,338 level covers; Theorem 7 agreed on all 143,964 valid rank-type covers in its range for blocks with up to 6 items and weights up to 8, and on all 186,470 for up to 7 items and weights up to 6. Theorem 6 was confirmed on all 406,253 levels with at least two feasible sets by testing the basis-exchange axiom. On non-principal levels with three distinct weights, rank-type covers suffice on 56% of the full-dimensional levels, and the remaining facets have coefficients at most 4.

Table EC.3: Exact checks of the results of Section 6 and Sections EC.4–EC.5.

Claim	Instances	Result
Corollary 2	every multiset of 2–5 weights in $\{1,\dots,6\}$, every capacity: 5,075 cases, 382,953 fractional term vertices	vertices as stated in the corollary
Theorem EC.2(a)	weights $\{\alpha, \alpha+1\}$, $\alpha \le 4$, at most 9 items, every capacity: 2,520 polytopes	largest count change 1
Theorem EC.2(b)	$(\alpha,\beta) = (3,5),(2,5),(4,7),(5,7)$	count changes $2,3,3,2$
Theorem EC.3(a),(b)	three weights in $\{1,\dots,7\}$, 3–8 items, every capacity, $2 \le \ell \le n-2$: 113,413 levels with dividing gaps	lift equals $\mathrm{conv}(X_\ell)$ on all
Theorem EC.3(c)	same grid, non-dividing gaps: 15,224 levels; and the stated example	lift fails on 282; both lifts fail on the example
Theorem EC.4	full-dimensional levels with at most 10 items; integrality of the system on all 14,443 nonempty levels with at most 12 items	facets as stated; integral
Corollary EC.4	9,156 classes of covers not equivalent to a bound	facet status as stated in the corollary
Corollary EC.5	all 3,353 facets of knapsack polytopes with weights $\{1,2\}$, at most 8 items	of the stated form
Theorem EC.5(a)	admissible sets with gcd 1: 314 with 3 elements up to 200, 83 with 4 up to 120, 10 with 5 up to 120 (4ti2)	all 3-element sets unit-edged; 3 four- and 3 five-element sets not
Theorem EC.6	143 sets with $c \le 100$, 4,331,041 vectors z; Graver bases of 2,342 sets with $c \le 400$	witness found for every z; count change at most 1
Proposition EC.5(b), small grids	3, 4, 5 coefficients up to 40, 48, 60 with at most 12, 7, 4 items per coefficient	no long edge within these bounds
Proposition EC.4	all weight multisets with at most 6 items in $\{1,\dots,6\}$: 13,714 blocks	criterion exact
Corollary 5(ii)	200 random GAP instances, 2–3 agents, 4–7 jobs (HiGHS)	largest relative difference 2.4×10^{-15}

EC.8 Larger Instances and Capacity Scaling

These experiments use GAP instances with $m = 10$ and $n = 100$ from types C and D and from a two-valued class W58-L with weights in $\{5, 8\}$ (seed 101; these are different instances from class W58 of Table 5), and network design instances with 12 nodes, 20 commodities and 56–62 arcs from classes H, T3 and T2 (seed 7), five per class. Both (C) and (CI), with exact level capacities and no additional inequalities, are solved as integer programs by HiGHS with a relative gap tolerance of 10^{-6} and time limits of 120 s (GAP) and 200 s (network design). For GAP type D, where no instance was solved, z^* is the best value found by either formulation. Table EC.4 reports the results.

Table EC.4: Larger instances (five per class). Closed: share (%) of $z^* - z_{\mathrm{LP}}(\mathrm{C})$ closed. LP time: average time (s) for LP(C), LP(CI) and column generation (CG). Branch and bound: instances solved, and shifted geometric means (shift 1) of time (s) and nodes.

		Closed		LP time			B&B, (C)			B&B, (CI)		
Problem	Class	CI	DW	C	CI	CG	solved	time	nodes	solved	time	nodes
GAP	C	21.2	87.3	0.01	2.33	8.2	5/5	2.6	42	5/5	66.5	483
	D	2.4	41.7	0.01	16.51	9.6	0/5	120.0	9072	0/5	120.2	1
	W58-L	100.0	100.0	0.01	0.78	10.2	5/5	0.3	1	5/5	1.1	1
Network	H	8.5	85.7	0.04	0.90	1.8	5/5	3.8	2	2/5	87.4	40
	T3	52.3	56.7	0.05	0.25	1.1	5/5	17.3	41	5/5	27.7	24
	T2	63.3	63.3	0.05	0.22	1.0	5/5	15.9	27	5/5	17.1	7

On the two-valued classes W58-L and T2, $z_{\mathrm{LP}}(\mathrm{CI}) = z_{\mathrm{DW}}$ to a relative 10^{-15}, as Theorem 4(iii) implies, and on W58-L both equal z^*. Used directly as an integer program, (CI) is not faster than (C): it has 10 to 28 times as many variables for GAP and 3 to 7 times as many for network design, and the solver closes (C) at or near the root on GAP types C and W58-L and network class H. On network classes T2 and T3, (CI) needs fewer nodes, 7 against 27 and 24 against 41, at comparable or greater times. On GAP type D, neither formulation finishes, and the single node reported for (CI) means that the solver was still processing the root node when the time limit was reached. This experiment uses the unreduced formulation without the valid inequalities, reduced level sets and separation routines of the branch-and-cut studies discussed in Section 7.8.

Table EC.5: Capacity scaling on two-valued GAP instances ($m = 10$, $n = 100$, weights in $\{5s, 8s\}$, capacities multiplied by s; seed 202, averages over three instances). DP states: $\sum_i (u_i + 1)n$.

s	Capacity	Variables of (CI)	LP(CI) (s)	CG (s)	Pricing (s)	DP states	$\lvert z_{\mathrm{DW}} - z_{\mathrm{LP}}(\mathrm{CI})\rvert$
1	52	10,100	0.86	10.4	0.5	5.2×10^4	3×10^{-12}
10	519	10,100	0.82	10.9	0.6	5.2×10^5	3×10^{-12}
100	5,190	10,100	0.72	11.6	1.5	5.2×10^6	3×10^{-12}
1000	51,896	10,100	0.88	21.6	11.6	5.2×10^7	3×10^{-12}

Table EC.5 reports the capacity-scaling experiment. (CI) is built with the closed form of Proposition 2(c). Column generation needs the same number of iterations at every scale, and

only its dynamic-programming pricing time grows. For two-valued blocks, pricing can be done in polynomial time by enumerating the number of heavy items and sorting, which would remove this growth; the experiment shows that the size of (CI) and the recognition of exactness are independent of the capacities, not that (CI) dominates column generation.

EC.9 Cut-and-Branch Experiments

Instances. The generator is that of Section EC.7 with two changes: the arc capacity is $u = \max\{\text{round}(\kappa\bar{d}), \max_k d_k\}$ with $\kappa \in \{1.5, 2\}$ instead of 2.5, and the fixed costs are multiplied by a factor $f \in \{3, 10\}$. The seed is 7 for every configuration, and instances that are infeasible as integer programs are replaced, as in Section EC.7. Each configuration in the first part of Table EC.6 was planned with two instances. The screening was stopped after the 11 instances shown, when the pattern was clear, so this part is an exploratory screen and its class means are descriptive rather than representative estimates; no generated instance is omitted.

Method. (C) and (CI) are as in Section 7, with exact level capacities. For CI+H, the exact level-hull inequalities of Theorem 5 are separated on the levels of LP(CI) that contain a straddle, as in Section 7.3. Separation uses row generation with a dynamic program over count and load as the optimization oracle, and stops when no inequality is violated by more than 10^{-6} or after 200 rounds. All root inequalities are kept, and the model is then solved by HiGHS 1.15.1 through highspy, with default settings, a relative gap tolerance of 10^{-6}, and a time limit equal to the limit minus the separation time. Times for CI+H include separation. On the largest class-H instance, separation converged after 57 rounds but took 1,046 s, more than the 300 s limit, so no branching took place; its CI+H entry in Table EC.6 is TL without a bracket. On every other instance separation took at most 17 s. The scripts are `sweep.py` and `cut_and_branch.py`, whose default options reproduce this setting, and `cut_and_branch_windows.py`; the raw results are in `results/sweep_*.jsonl`, `results/scale_*.jsonl` and `results/cbw_*.jsonl`. Times depend on the machine.

Level windows and reduction. On the two instances with 20 nodes and 40 commodities, we also ran the pipeline of the network design study.

- Each arc keeps the exact levels within one of the count estimated from LP(C); the levels below and above this window are each merged into one block with a count range and the largest exact capacity of the range. This preserves the integer optimum.
- Exact level-hull inequalities are separated on the window levels with a budget of 25% of the time limit, stopping early when the bound improves by less than a relative 10^{-5} over three rounds.
- Every arc that carries no inequality is returned to compact form, but only if the root bound is unchanged.

Table EC.6: Cut-and-branch with exact level-hull inequalities on unsplittable network design (all instances generated). Size: nodes/commodities. κ: capacity multiplier; f: fixed-cost factor; L: number of levels per arc. Gap: $(z^* - z_{\mathrm{LP}}(\mathrm{C}))/z^*$ (%), with z^* the best integer value found. Root: share (%) of the gap closed by LP(CI) and by CI+H at the root, before any solver cut. Last three columns: time in seconds and nodes in parentheses when solved; TL when the limit (120 s for size 12/20, 300 s for 20/40) was reached, with the share of the gap closed by the final bound in brackets; TL without brackets: root separation alone exceeded the limit, so no branch-and-bound run took place.

						Root		Branch-and-bound		
Class	Size	κ	f	L	Gap	CI	CI+H	C	CI	CI+H
H	12/20	1.5	3	6	11.6	9.1	89.8	5.1 (1)	23.8 (1)	TL [96.1]
H	12/20	1.5	3	5	17.1	16.8	99.0	2.1 (1)	TL [93.7]	28.4 (1)
H	12/20	1.5	10	6	13.8	8.1	82.9	6.3 (1)	TL [80.3]	TL [83.9]
S16	12/20	1.5	3	2	20.4	47.7	91.5	4.7 (1)	9.4 (1)	5.8 (0)
S16	12/20	1.5	3	2	19.7	53.4	89.6	4.6 (1)	10.5 (1)	3.0 (0)
S16	12/20	1.5	10	2	23.2	48.8	91.9	4.3 (1)	11.4 (1)	6.7 (1)
S16	12/20	1.5	10	2	23.1	50.9	87.4	9.8 (5)	9.5 (1)	6.1 (1)
S16	12/20	2	3	3	16.8	27.8	78.9	8.5 (3)	93.7 (466)	24.4 (1)
S16	12/20	2	3	3	17.9	31.5	89.5	10.9 (3)	17.8 (1)	14.2 (1)
S16	12/20	2	10	3	20.3	29.0	74.3	16.6 (7)	TL [77.7]	46.3 (31)
S16	12/20	2	10	3	20.0	35.5	93.7	8.9 (1)	21.0 (1)	28.1 (1)
H	20/40	1.5	10	7	17.8	10.9	92.8	TL [96.1]	TL [43.0]	TL
S16	20/40	1.5	10	2	22.1	78.8	99.1	35.4 (1)	TL [94.6]	57.6 (0)

The reduction was rejected on both instances because every arc carried an inequality, so the windowed model was solved. Table EC.7 reports CIW, the windowed model without cuts, and CIW+H, the windowed model with the root inequalities. On the class-H instance the separation budget stopped the root loop after 12 rounds, at 62.5% of the gap.

Table EC.7: Level windows and the bound-preserving reduction on the two instances of size 20/40 in Table EC.6 (time limit 300 s). Root: share (%) of the gap closed at the root. Levels kept: arcs that keep level variables after the reduction. Other columns as in Table EC.6.

		Root			Branch-and-bound		
Class	Size	CIW	CIW+H	Levels kept	C	CIW	CIW+H
H	20/40	10.9	62.5	142/142	TL [96.1]	TL [52.4]	TL [81.7]
S16	20/40	78.8	99.1	142/142	35.4 (1)	96.2 (1)	89.5 (1)