

Matteo Fischetti*
Andrea Lodi**
Andrea Tramontani***

On the Separation of Disjunctive Cuts

the date of receipt and acceptance should be inserted later

Abstract. Disjunctive cuts for Mixed-Integer Linear Programs have been introduced by Egon Balas in the late 70's, and successfully exploited in practice since the late 90's. In this paper we investigate the main ingredients of a disjunctive cut separation procedure, and analyze their impact on the quality of the root-node bound for a set of instances taken from MIPLIB library. We compare alternative normalization conditions, and try to better understand their role. In particular we point out that the constraints that become redundant (because of the disjunction used) can produce over-weak cuts, and analyze this property with respect to the normalization used. Finally, we introduce a new normalization condition and analyze its theoretical properties and computational behavior. Along the paper, we make use of a number of small numerical examples to illustrate some basic (and often misinterpreted) disjunctive programming features.

1. Introduction

We consider the Mixed-Integer Linear Program (MIP)

$$\min\{cx : Ax \geq b, x_j \text{ integer for all } j \in J\} \quad (1)$$

with bounds on x (if any) included in $Ax \geq b$, where A is a given $m \times n$ matrix and $J \subseteq \{1, \dots, n\}$. For technical reasons, we assume w.l.o.g. that the system $Ax \geq b$ implies (or contains explicitly) the trivial inequality $0x \geq -1$, in the sense that this latter inequality can be obtained as a nonnegative combination of the rows of $Ax \geq b$ ¹.

Let x^* denote an optimal solution of the continuous relaxation $\min\{cx : x \in P\}$ where

$$P := \{x \in \mathbb{R}^n : Ax \geq b\}. \quad (2)$$

Address(es) of author(s) should be given

* DEL, University of Padova, via Gradenigo 6A - 35131 Padova, Italy. Supported in part by the EU project ARRIVAL (contract n. FP6-021235-2) and by MiUR, Italy (PRIN 2006 project "Models and algorithms for robust network optimization"). e-mail: matteo.fischetti@unipd.it

** DEIS, University of Bologna, viale Risorgimento 2, 40136 Bologna, Italy. Supported in part by the EU project ARRIVAL (contract n. FP6-021235-2). andrea.lodi@unibo.it

*** DEIS, University of Bologna, viale Risorgimento 2, 40136 Bologna, Italy. andrea.tramontani@unibo.it

¹ For problems with at least one bounded variable, the trivial inequality can always be obtained by adding the bound constraints on a single variable, say $x_j \geq LB_j$ and $-x_j \geq -UB_j$, and dividing the resulting inequality by $UB_j - LB_j > 0$.

We are given a disjunction of the form

$$\pi x \leq \pi_0 \quad \text{OR} \quad \pi x \geq \pi_0 + 1 \quad (3)$$

such that (π, π_0) is integer, $\pi_j = 0, \forall j \notin J$ and $\pi x^* - \pi_0 = \eta^*$, with $\eta^* \in]0, 1[$.

In this paper we are interested in deriving the “strongest” (in some sense to be discussed later) disjunctive cut $\gamma x \geq \gamma_0$ violated by x^* , according to the classical approach of Balas [2]. (Disjunctive cuts which can be derived by imposing a single disjunction as (3) on a polyhedron P are also known as *split* cuts; see Cook, Kannan and Schrijver [14].) To this end, let us denote by P_0 (respectively, P_1) the polyhedron obtained from P by imposing the additional restriction $\pi x \leq \pi_0$ (resp., $\pi x \geq \pi_0 + 1$). By Farkas lemma, the validity of $\gamma x \geq \gamma_0$ for P_0 and for P_1 , and hence for $\text{conv}(P_0 \cup P_1)$, can always be certified by means of nonnegative multipliers (u, u_0, v, v_0) associated with the inequalities defining P_0 and P_1 according to the following scheme:

$$\begin{array}{ll} P_0 & P_1 \\ (u) \quad Ax \geq b & (v) \quad Ax \geq b \\ (u_0) \quad -\pi x \geq -\pi_0 & (v_0) \quad \pi x \geq \pi_0 + 1 \end{array}$$

A most-violated disjunctive cut can therefore be found by solving the following *Cut Generating Linear Program* (CGLP) that determines the Farkas multipliers so as to maximize the violation of the resulting cut with respect to the given point x^* :

$$\text{(CGLP)} \quad \min \quad \gamma x^* - \gamma_0 \quad (4)$$

$$\gamma = uA - u_0\pi \quad (5)$$

$$\gamma = vA + v_0\pi \quad (6)$$

$$\gamma_0 = ub - u_0\pi_0 \quad (7)$$

$$\gamma_0 = vb + v_0(\pi_0 + 1). \quad (8)$$

$$u, v, u_0, v_0 \geq 0 \quad (9)$$

Note that, according to Farkas lemma, the two equations (7) and (8) defining γ_0 should be relaxed into \leq inequalities. However it is not difficult to see that, due to the (possibly implicit) presence of the trivial inequality $0x \geq -1$, one can always require that equality holds in both cases.

By construction, any feasible CGLP solution with negative objective function value corresponds to a violated disjunctive cut. However, as stated, the feasible CGLP set is a cone and needs to be truncated so as to produce a bounded LP in case a violated cut exists. This crucial step will be addressed in the next session.

Usually, the CGLP is projected on the support of x^* . Given a variable x_k restricted to be nonnegative and such that $x_k^* = 0^2$, it is well known that one

² Of course, variables with nonzero lower bound can be shifted, while variables at the upper bound can be complemented.

can project x_k away, not considering explicitly the CGLP constraints associated with γ_k . The cut coefficient γ_k is derived afterwards by solving the trivial lifting problem

$$\min\{\gamma_k : \gamma_k = uA_k - u_0\pi_k = vA_k + v_0\pi_k, \quad u, v \geq 0\}, \quad (10)$$

where all the Farkas multipliers are fixed, except those related to the bound constraint $x_k \geq 0$ in P_0 and P_1 .

In practice, disjunction (3) typically involves only one integer variable—for 0-1 ILPs, it reads $x_j \leq 0$ OR $x_j \geq 1$, with x_j^* fractional. As such, the disjunctive cut only exploits the integrality requirement on a single variable and can therefore be improved easily by an *a posteriori cut strengthening* procedure as the one proposed by Balas and Jeroslow [5].

Recently, Balas and Perregaard [7] developed an elegant and efficient way of solving the CGLP by making pivot operations in the “natural” tableau involving the original x variables only, which represents a crucial speed-up in the implementation of the method.

In this paper we investigate computationally the main ingredients of a disjunctive cut separation procedure, and analyze their impact on the overall performance at the root node of the branching tree. To be more specific, we consider a testbed of MIPs taken from MIPLIB library [10]. For each instance, we solve the root-node LP relaxation and generate 10 rounds of disjunctive cuts computed according to alternative strategies. In each round, a violated disjunctive cut is generated for each fractional LP components x_j^* , by exploiting the disjunction $x_j \leq \lfloor x_j^* \rfloor$ OR $x_j \geq \lfloor x_j^* \rfloor + 1$. In order to limit possible side effects, no *a posteriori* cut strengthening procedure is applied.

The paper is organized as follows. In Section 2 we compare some classical normalization conditions used to truncate the CGLP cone, and try to better understand their role. In Section 3 we characterize weak rays/vertices of the CGLP leading to dominated cuts. In Section 4 we show that using redundant constraints in the CGLP can lead to very weak cuts, and we analyze such an issue with respect to the normalization used. In Section 5 we introduce a new normalization and we analyze its theoretical properties and computational behavior.

2. The role of normalization

In order to truncate the CGLP cone one can introduce a suitable cut normalization condition expressed as a linear (in)equality. A possible normalization, called *trivial* in the sequel, is as follows:

$$u_0 + v_0 = 1. \quad (11)$$

One of the most widely-used (and effective) truncation condition, called the *Standard Normalization Condition* (SNC) in the following, reads instead:

$$\sum_{i=1}^m u_i + \sum_{i=1}^m v_i + u_0 + v_0 = 1. \quad (12)$$

This latter condition was proposed in Balas [1] and investigated by Ceria and Soares [12], and by Balas and Perregaard [6, 7].

The choice of the normalization condition turns out to be crucial for an effective selection of a “strong” disjunctive cut in that it affects heavily the ranking of the feasible CGLP solutions. To see this it is enough to observe that, since the CGLP feasible set is a cone and assuming a violated cut exists, one can always swap the role of the objective function and of the normalization condition. In other words, one could equivalently fix the objective function to a given negative value (say, -1) so as to only allow for violated cuts, and use the left-hand side of the normalization condition as the objective function to be minimized. Hence, the actual CGLP “optimal” cut depends heavily on the normalization condition.

Balas and Perregaard [7] showed that the well-known Gomory Mixed-Integer (GMI) cut [18] is a basic solution of the CGLP when either the SNC or the trivial normalization is applied. In fact, we next show that this solution is indeed optimal when the trivial normalization (11) is used. We start with a useful lemma.

Lemma 1 *Let $x^* \in P$ and let $(\gamma, \gamma_0, u, v, u_0, v_0)$ be a feasible solution of the CGLP (4)-(9). Then valid upper bounds on the cut violation can be computed as follows:*

$$\begin{aligned} \text{UB1: } & \gamma_0 - \gamma x^* \leq u_0 \eta^* \\ \text{UB2: } & \gamma_0 - \gamma x^* \leq v_0 (1 - \eta^*) \\ \text{UB3: } & \gamma_0 - \gamma x^* \leq (u_0 + v_0) (1 - \eta^*) \eta^*. \end{aligned}$$

Proof. Because of (5) and (7), $\gamma x^* - \gamma_0 = u(Ax^* - b) - u_0(\pi x^* - \pi_0) \geq -u_0 \eta^*$. Analogously, from (6) and (8) we obtain $\gamma x^* - \gamma_0 = v(Ax^* - b) + v_0(\pi x^* - \pi_0 - 1) \geq -v_0(1 - \eta^*)$. Adding up the two inequalities above weighed by $1 - \eta^*$ and η^* , respectively, one gets the claimed UB3 bound. \square

Given a vertex x^* of P and the associated basis, the next theorem shows how to compute a solution of the CGLP whose violation is equal to bound UB3 above—for any given disjunction (3). Moreover, as shown in [6, 7], for an appropriate choice of the disjunction this CGLP solution yields precisely a GMI cut associated with the optimal LP tableau. As a consequence of Lemma 1, this easily-computable cut has a violation that is optimal among the cuts with constant $u_0 + v_0$, i.e., when the trivial normalization (11) is imposed. Note however that this is not necessarily the case when a different normalization (in particular, the SNC one) is applied.

For any vector v , let operator $[v]_+$ takes the maximum between the argument and zero (componentwise); by definition, $v \equiv [v]_+ - [-v]_+$ with $[v]_+ \geq 0$ and $[-v]_+ \geq 0$.

Theorem 2 *Assume w.l.o.g. $\text{rank}(A) = n$. Given a vertex x^* of P , let system $Ax \geq b$ be partitioned into $Bx \geq b_B$ and $Nx \geq b_N$, where $Bx^* = b_B$ and B is an $n \times n$ nonsingular matrix. Let (u_B, v_B) and (u_N, v_N) denote the Farkas*

multipliers associated with the rows of B and N , respectively. For a given disjunction (3) with $\eta^* = \pi x^* - \pi_0 \in [0, 1]$, let $u_0^* = 1 - \eta^*$, $v_0^* = \eta^*$, $u_N^* = v_N^* = 0$, $u_B^* = [\pi B^{-1}]_+$ and $v_B^* = [-\pi B^{-1}]_+$, while γ^* and γ_0^* are defined through equations (5) and (7), respectively. Then $(\gamma^*, \gamma_0^*, u^*, v^*, u_0^*, v_0^*)$ is an optimal CGLP solution w.r.t. the trivial normalization(11).

Proof. We first prove feasibility. Consistency between (5) and (6) requires $u^* A - u_0^* \pi = v^* A + v_0^* \pi$, i.e., $u_B^* - v_B^* = (u_0^* + v_0^*) \pi B^{-1} = \pi B^{-1}$, a condition that follows directly from the definition of u_B^* and v_B^* . Analogously, consistency between (7) and (8) requires $(u_B^* - v_B^*) b_B = (u_0^* + v_0^*) \pi_0 + v_0^*$, i.e., $\pi B^{-1} b_B = \pi_0 + v_0^*$. This latter equation is indeed satisfied because $B^{-1} b_B = x^*$ and $v_0^* = \eta^* = \pi x^* - \pi_0$. As to optimality, we first observe that $u_0^* + v_0^* = 1$ holds by definition. Because of (5) and (7), $\gamma x^* - \gamma_0 = u^*(Ax^* - b) - u_0^*(\pi x^* - \pi_0) = u_B^*(Bx^* - b_B) + u_N^*(Nx^* - b_N) - u_0^* \eta^* = 0 + 0 - (1 - \eta^*) \eta^*$, hence the cut violation attains bound UB3 of Lemma 1. \square

The theorem above shows that, in case the trivial normalization is adopted, the CGLP can be solved in a closed form for any vertex x^* . Moreover, with this normalization, in all optimal CGLP solutions the slack constraints receive a null Farkas multiplier, i.e., only tight constraints play a role in the cut derivation—a choice that typically leads to weak cuts.

The first set of experiments we designed was aimed at evaluating the actual practical impact of different normalization conditions. In particular, we compared the SNC normalization (12) with the alternative trivial normalization (11).

The outcome of our experiments is given in Table 1. As already mentioned, we applied 10 rounds of cuts. At each round, a cut was generated from each fractional variable. No a-posteriori cut strengthening was applied. As usual, the CGLP is solved projected on the x^* support. Instances denoted as “*” are neglected in the average computations. The table reports (i) the number of separated cuts, (ii) the quality of the lower bound (i.e., percentage gap closed at the root node) and (iii) the average cardinality of $S(u, v) := \{i \in \{1, \dots, m\} : u_i + v_i > 0\}$ ($|S|$ for short), i.e., how many constraints are actually used, on average, to generate a cut.





Table 1 shows clearly that normalizations (12) and (11) yield quite different results. As a matter of fact, the dual support of cuts separated with (12) is much sparser (i.e., less constraints are used in the cut derivation) and the quality of final bound is significantly improved. To get more insights on the different behaviors of (12) and (11), for instance p0201 we provide a full picture of the main differences between the separated inequalities.

Figures 1–3 report, for each iteration, the dual bound reached after adding the cuts, the average density of the cuts (i.e., the number of nonzero coefficients), and the average cardinality of $S(u, v)$. Figure 4 reports, for each $k = 1, \dots, 10$, the number of separated cuts having “rank” k . Here, we use a relaxed definition of rank, namely we compute the rank $rnk(\gamma, \gamma_0)$ of a cut $\gamma x \geq \gamma_0$ as

$$rnk(\gamma, \gamma_0) := 1 + \max_{i \in S(u, v)} rnk(a_i, b_i),$$

Table 1. Trivial vs. SNC normalization.

Instance	Trivial normalization (GMI)			SNC normalization		
	# cuts	%gap	S	# cuts	%gap	S
bell3a	137	70.74	59.49	71	70.74	43.72
bell5	202	28.18	31.20	178	94.29	11.75
blend2	156	28.73	11.70	192	30.51	8.10
flugpl	93	15.15	7.57	92	18.36	5.85
gt2	191	98.71	14.52	196	93.46	10.28
lseu	152	32.94	14.34	196	41.33	9.17
*markshare1	68	0.00	1.00	74	0.00	1.39
mod008	104	12.09	10.40	139	17.05	12.41
p0033	103	58.33	5.72	113	67.86	4.81
p0201	574	18.58	56.03	767	93.82	13.43
rout	445	8.52	135.39	434	24.26	68.07
*stein27	235	0.00	19.74	252	0.00	6.53
vpml	255	36.95	9.03	263	55.84	5.39
vpml2	424	42.08	71.72	403	74.96	17.27
avg.	236.333	37.583	35.593	253.667	56.873	17.521

**Fig. 1.** SNC vs. GMI: dual bound for instance $p0201$.**Fig. 2.** SNC vs. GMI: avg. cut density for instance $p0201$.**Fig. 3.** SNC vs. GMI: avg. cardinality of $S(u, v)$ for instance $p0201$.**Fig. 4.** SNC vs. GMI: cut rank for instance $p0201$.

where $rnk(a_i, b_i)$ is the rank of constraint $a_i x \geq b_i$ (constraints in the original formulation are defined to be of zero rank)³.

³ Note that this way of computing the rank provides just an upper bound on the classical definition of Chvátal rank [13].

2.1. Why does SNC normalization work so well?

A careful analysis of the computational results in Table 1 and Figures 1–4 reveal a very (tricky but) important feature of the SNC scheme that improves significantly its performance. Indeed, it turns out that the use of the SNC normalization (12) enforces the following very nice properties:

1. The norm of the separated cuts tends to become smaller and smaller as a result of the small multipliers used for the newly generated cuts (that is, in turn, a consequence of having limited the multiplier sum to 1). This means that the separated cuts inserted in the LP are automatically scaled so as to have “small coefficients”. Therefore, in the subsequent iterations these cuts would need big Farkas multipliers to become relevant, a situation that is however penalized by the normalization condition itself. As a consequence, the normalization penalizes implicitly the rank of the cuts to be generated, because high-rank cuts will be “expensive” in terms of multiplier sum, hence low-rank cuts tend to be separated at each step.
2. Since low-rank cuts are preferred and since the original (rank-0) inequalities are generally sparse, the separated cuts tend to remain sparse; this is also a consequence of the fact that the SNC normalization tends to reduce the sum of the components of the Farkas multiplier vector and hence it increases the sparsity of its support, so a small number of constraints are typically used in the disjunctive cut derivation.

Trivial normalization (11), instead, takes care only of the Farkas multipliers u_0 and v_0 associated with the disjunction. Indeed, as shown in Section 2, only constraints which are tight at x^* are used in the cut derivation, thus the rank of the cuts increases very quickly, basically at each iteration. Moreover, all other constraint multipliers are not penalized, hence (i) several constraints are used in the cut derivation, thus cuts increase their density, and (ii) Farkas multipliers can assume huge values, thus the subsequent cut lifting procedure may produce very weak coefficients for the variables outside the support of x^* .

In the SNC normalization case the coefficient lifting is not an issue. Indeed, since all the constraint multipliers in the SNC normalization are penalized and each multiplier tends to be small, the coefficient lifting of the variables outside the support of x^* – to be performed afterwards – is “safe”, i.e., also the coefficients of these variables remain under control.

2.2. Nothing is perfect!

Although it produced good results in the experiments reported in Table 1, there are cases where normalization (12) may lead to very weak disjunctive cuts.

Bad Scaling. A bad feature of the SNC normalization is its dependency on the relative scaling of the constraints, in the sense that the relative size of the Farkas multipliers (whose sum is fixed to 1) depends on the relative size of the

coefficients of the corresponding constraints. Indeed, it is easy to see that the multiplication by a positive factor ϕ of the i -th constraint in the system $Ax \geq b$ implies that the corresponding u_i and v_i multipliers are divided by ϕ , which in turn is equivalent to use a coefficient $1/\phi$ (instead of 1) in the normalization condition (12). Thus, the scaled constraint is “cheaper” if one interprets the right hand side of (12) as a resource.

The following experiment clearly demonstrates this unstable behavior: we ran the CGLP code with the classical SNC normalization condition, as in Table 1, but we just multiply by 1,000 each disjunctive cut before its addition to the current LP. At first glance, one could guess that this “innocent change” would not have any impact on the overall performance, but the actual results reported in Table 2 show that this is definitely not the case.

Table 2. “Classical” SNC approach vs. “Bad scaled” SNC approach.

Instance	“Classical” SNC			“Bad scaled” SNC		
	# cuts	%gap	S	# cuts	%gap	S
bell3a	71	70.74	43.72	69	70.74	44.32
bell5	178	94.29	11.75	214	88.83	17.47
blend2	192	30.51	8.10	166	28.91	11.71
flugpl	92	18.36	5.85	90	15.40	7.40
gt2	196	93.46	10.28	184	93.42	17.22
lseu	196	41.33	9.17	137	38.58	10.88
*markshare1	74	0.00	1.39	206	0.00	14.60
mod008	139	17.05	12.41	104	3.90	10.21
p0033	113	67.86	4.81	94	57.09	6.40
p0201	767	93.82	13.43	610	49.91	45.72
rout	434	24.26	68.07	435	13.03	152.66
*stein27	252	0.00	6.53	248	0.00	22.39
vpm1	263	55.84	5.39	244	47.59	8.50
vpm2	403	74.96	17.27	420	54.39	22.27
avg.	253.667	56.873	17.521	230.583	46.816	29.563

As explained, multiplying by 1,000 the generated cuts is equivalent to dividing by 1,000 the coefficient of the corresponding Farkas multipliers u_i and v_i in the normalization condition, so we actually weaken the penalty on the choice $u_i + v_i > 0$ that leads to low-rank sparse cuts. In other words, the scaling operation interferes with the nice SNC tendency of producing low-rank cuts, and the overall performance deteriorates significantly, as shown in detail for problem p0201 in Figures 5–8. (Incidentally, the above example shows the importance of “small implementation details” when evaluating the performance of a method—two apparently equivalent implementations of precisely the same idea lead to very different outcomes.)

Bad Examples. Even for toy instances, the CGLP can have hard time in finding a good disjunctive cut. This is illustrated by the following two simple 2-dimensional cases, where the optimal CGLP solution may correspond to very weak cuts.




Fig. 5. “Classical” SNC vs. “Bad scaled” SNC: dual bound for instance *p0201*.




Fig. 6. “Classical” SNC vs. “Bad scaled” SNC: avg. cut density for instance *p0201*.




Fig. 7. “Classical” SNC vs. “Bad scaled” SNC: avg. cardinality of $S(u, v)$ for instance *p0201*.




Fig. 8. “Classical” SNC vs. “Bad scaled” SNC: cut rank for instance *p0201*.

Example 1. Consider the simple ILP whose continuous relaxation, depicted in Figure 9, has one of the constraints, namely (a5), scaled by a parameter $k > 0$:

$$\begin{aligned}
 \min \quad & -x_1 - 2x_2 \\
 \text{(a1)} \quad & 4x_1 - 4x_2 \geq -2 \\
 \text{(a2)} \quad & -2x_1 - 2x_2 \geq -3 \\
 \text{(a3)} \quad & 8x_1 - 4x_2 \geq -1 \\
 \text{(a4)} \quad & -x_1 \geq -1 \\
 \text{(a5)} \quad & -kx_2 \geq -k \\
 \text{(a6)} \quad & x_1 \geq 0 \\
 \text{(a7)} \quad & x_2 \geq 0
 \end{aligned}$$

The optimal solution of the LP relaxation is $x^* = (\frac{1}{2}, 1)$ and three cuts can be derived from disjunction $x_1 \leq 0$ OR $x_1 \geq 1$, namely:

- (c1) $2x_2 \leq 1$, corresponding to the basic solution of the CGLP (u_1, v_2, u_0, v_0) , of value $z_1 = -\frac{2}{11}$, optimal for $k \leq 8$;
- (c2) $-x_1 + 4x_2 \leq 1$, corresponding to the basic solution of the CGLP (u_3, v_2, u_0, v_0) , of value $z_2 = -\frac{1}{6}$, never optimal.
- (c3) $-x_1 + 2x_2 \leq 1$, corresponding to the basic solution of the CGLP (u_1, v_5, u_0, v_0) , of value $z_3 = -\frac{k}{4+5k}$, optimal for $k \geq 8$.

So, depending on the value of k , the optimal CGLP solution corresponds to weak cuts, either (c1) or (c3), whereas the facet-defining cut (c2) will never be selected.




Fig. 9. Example 1 depicted.

□

Example 2. For the simple ILP whose continuous relaxation is depicted in Figure 10:

$$\begin{array}{ll}
 \min & -x_1 - 2x_2 \\
 \text{(b1)} & 2x_1 - 2x_2 \geq -1 \\
 \text{(b2)} & -2x_1 - 2x_2 \geq -3 \\
 \text{(b3)} & 4x_1 + 4x_2 \geq 3 \\
 \text{(b4)} & -x_1 \geq -1 \\
 \text{(b5)} & -x_2 \geq -1 \\
 \text{(b6)} & x_1 \geq 0 \\
 \text{(b7)} & x_2 \geq 0
 \end{array}$$

the optimal solution of the continuous relaxation is again $x^* = (\frac{1}{2}, 1)$ and three cuts can be derived from disjunction $x_1 \leq 0$ OR $x_1 \geq 1$ (yielding $P_0 = \emptyset$), namely:

- (c1) $2x_2 \leq 1$, corresponding to the basic solution of the CGLP (u_1, v_2, u_0, v_0) , of value $z_1 = -\frac{1}{6}$ (optimal);
- (c2) $x_1 \geq 1$, corresponding to the basic solution of the CGLP (u_1, u_3, u_0, v_0) , of value $z_2 = -\frac{1}{22}$ (nonoptimal);
- (c3) $-x_1 + 2x_2 \leq 1$, corresponding to the basic solution of the CGLP (u_1, v_5, u_0, v_0) , of value $z_3 = -\frac{1}{10}$ (nonoptimal).

Since (c2) is not optimal for the associated CGLP, the only facet-defining cut (c2) will not be selected. □




Fig. 10. Example 2 depicted.

2.3. Comments

The examples above show clearly the following fact: even if the solution of the CGLP is a vertex, the corresponding disjunctive cut can be very weak. At first glance, this may be seen as a counter-intuitive result as one would expect that CGLP vertices correspond to facets of $\text{conv}(P_0 \cup P_1)$. This is however not the case, as discussed e.g. in Balas and Perregaard [6], since the CGLP is not defined in the “natural” reverse polar space (γ, γ_0) but in an enlarged space involving the Farkas variables explicitly. As a matter of fact, in the extended space $(\gamma, \gamma_0, u, v, u_0, v_0)$ there are several rays/vertices whose projection in the (γ, γ_0) space is nonextremal, therefore the corresponding cut can be obtained as the sum of other valid cuts and hence is dominated. By using software PORTA [15] we can get a clear picture of the situation in Example 1. In the natural polar space (γ, γ_0) , the projected CGLP cone has only 4 extreme rays that correspond to the facets of $\text{conv}(P_0 \cup P_1)$. In space $(\gamma, \gamma_0, u, v, u_0, v_0)$, instead, the CGLP cone has 117 extreme rays that correspond to 117 vertices once normalization (12) is applied. Only 6 of these vertices correspond to violated constraints, and 3 of them correspond to the cuts depicted in Figure 9. So, most CGLP vertices in the $(\gamma, \gamma_0, u, v, u_0, v_0)$ space correspond to very weak cuts, and the cut separation procedure can be in trouble in returning a facet-defining cut even in this toy example. As mentioned above, this is essentially due to the fact that the cut is separated in the extended space $(\gamma, \gamma_0, u, v, u_0, v_0)$, where a dominated cut could turn out not to be dominated in terms of the multipliers used for its generation. For instance, 3 extreme rays of the CGLP cone for Example 1 are

reported below.

	γ_1	γ_2	γ_0	u_1	u_2	u_3	u_4	u_5	u_6	u_7	v_1	v_2	v_3	v_4	v_5	v_6	v_7	u_0	v_0
(r_1)	1	-4	-1	0	0	1	0	0	0	0	0	2	0	0	0	0	0	7	5
(r_2)	-1	0	-1	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0
(r_3)	0	-4	-2	1	0	0	0	0	0	0	0	2	0	0	0	0	0	4	4

In the (γ, γ_0) space, the third constraint is clearly dominated as it is just the sum of the previous ones, but there is no way to obtain ray r_3 as conic combination of rays r_1 and r_2 in the extended space, due to the presence of the Farkas components. The above drawback is even more evident in Example 2 where $P_0 = \emptyset$, hence $x_1 \geq 1$ itself is a valid cut (c2), but not the best one for the CGLP.

3. Weak CGLP rays/vertices and dominated cuts

The examples in the previous section show that some rays/vertices of the CGLP lead to weak cuts and should not be used. In this section we formally characterize those rays/vertices which correspond to cuts which are trivially dominated by other cuts associated with solutions of the same CGLP. The first step is the following definition.

Definition 3 (*Strictly dominated cuts*) Let $\tilde{\gamma}x \geq \tilde{\gamma}_0$ be a cut valid for $\text{conv}(P_0 \cup P_1)$ but not for P . If there exists another cut $\bar{\gamma}x \geq \bar{\gamma}_0$ valid for $\text{conv}(P_0 \cup P_1)$ such that $\{x \in P : \bar{\gamma}x \geq \bar{\gamma}_0\} \subset \{x \in P : \tilde{\gamma}x \geq \tilde{\gamma}_0\}$, then the cut $\tilde{\gamma}x \geq \tilde{\gamma}_0$ is said to be strictly dominated w.r.t. P .

Note that, in the above definition, the domination of cut $\tilde{\gamma}x \geq \tilde{\gamma}_0$ only depends on a single other cut $(\bar{\gamma}x \geq \bar{\gamma}_0)$.

Lemma 4 Let $\tilde{\gamma}x \geq \tilde{\gamma}_0$ be a valid cut for $\text{conv}(P_0 \cup P_1)$ such that $\tilde{P} := \{x \in P : \tilde{\gamma}x \geq \tilde{\gamma}_0\} \subset P$, and assume \tilde{P} full dimensional. If there exists another cut $\bar{\gamma}x \geq \bar{\gamma}_0$ valid for $\text{conv}(P_0 \cup P_1)$ and such that $\tilde{\gamma} = \bar{\gamma} + \mu A$, $\tilde{\gamma}_0 = \bar{\gamma}_0 + \mu b$ for a certain $\mu \in \mathbb{R}_+^m \setminus \{0\}$, then $\tilde{\gamma}x \geq \tilde{\gamma}_0$ is strictly dominated w.r.t. P .

Proof. Define $\bar{P} := \{x \in P : \bar{\gamma}x \geq \bar{\gamma}_0\}$. By definition, $x \in P$ and $\bar{\gamma}x \geq \bar{\gamma}_0$ imply $\tilde{\gamma}x \geq \tilde{\gamma}_0$, hence $\bar{P} \subseteq \tilde{P}$. We need to show that the above inclusion is always strict. Indeed, let $\tilde{F} := \{x \in P : \tilde{\gamma}x = \tilde{\gamma}_0\}$ denote the face of \tilde{P} induced by $\tilde{\gamma}x \geq \tilde{\gamma}_0$, and consider any given $h \in \{1, \dots, m\}$ such that $\mu_h > 0$. Since \tilde{P} is full dimensional, there exists $\hat{x} \in \tilde{F}$ such that $a_h \hat{x} > b_h$ (otherwise $\tilde{\gamma}x \geq \tilde{\gamma}_0$ would be a positive multiple of $a_h x \geq b_h$, impossible since we are assuming $\tilde{P} \subset P$). Hence $\bar{\gamma} \hat{x} - \bar{\gamma}_0 = (\tilde{\gamma} \hat{x} - \tilde{\gamma}_0) - \mu(A \hat{x} - b) \leq -\mu_h(a_h \hat{x} - b_h) < 0$, i.e., $\hat{x} \in \tilde{P} \setminus \bar{P}$. \square

For any feasible solution $(\gamma, \gamma_0, u, v, u_0, v_0)$ of (5)–(9), define $S(u) := \{i \in \{1, \dots, m\} : u_i > 0\}$ and $S(v) := \{i \in \{1, \dots, m\} : v_i > 0\}$. It is not difficult to show that in any extreme ray of (5)–(9) yielding a cut nonvalid for P , both u_0 and v_0 are strictly positive, while $S(u)$ and $S(v)$ are disjoint (i.e., $S(u) \cap S(v) = \emptyset$).

This property is also inherited by the vertices of the CGLP with normalization (12) (see, Balas and Perregaard [7]). We next give a characterization of the extreme rays/vertices of the CGLP that lead to strictly dominated cuts according to Definition 3.

Theorem 5 *Assume $\text{conv}(P_0 \cup P_1)$ full dimensional. Let $(\tilde{\gamma}, \tilde{\gamma}_0, \tilde{u}, \tilde{v}, \tilde{u}_0, \tilde{v}_0)$ be an extreme ray of the CGLP cone (5)–(9) corresponding to a cut $\tilde{\gamma}x \geq \tilde{\gamma}_0$ nonvalid for P . Then $\tilde{\gamma}x \geq \tilde{\gamma}_0$ is strictly dominated w.r.t. P if and only if there exists a feasible solution $(\bar{\gamma}, \bar{\gamma}_0, \bar{u}, \bar{v}, \bar{u}_0, \bar{v}_0)$ of (5)–(9) such that $S(\bar{u}) \cap S(\bar{v}) \neq \emptyset$.*

Proof. We first prove the if condition. Given a feasible solution $(\tilde{\gamma}, \tilde{\gamma}_0, \tilde{u}, \tilde{v}, \tilde{u}_0, \tilde{v}_0)$ of (5)–(9) such that $S(\tilde{u}) \cap S(\tilde{v}) \neq \emptyset$, define $\mu = \min\{\tilde{u}, \tilde{v}\}$ (componentwise) and note that $\mu_i > 0$ for any $i \in S(\tilde{u}) \cap S(\tilde{v})$. Then, define $\bar{u} = \tilde{u} - \mu \geq 0$, $\bar{v} = \tilde{v} - \mu \geq 0$, $\bar{\gamma} = \tilde{\gamma} - \mu A$, $\bar{\gamma}_0 = \tilde{\gamma}_0 - \mu b$. Since $(\bar{\gamma}, \bar{\gamma}_0, \bar{u}, \bar{v}, \tilde{u}_0, \tilde{v}_0)$ is a feasible solution of (5)–(9), the cut $\bar{\gamma}x \geq \bar{\gamma}_0$ is valid for $\text{conv}(P_0 \cup P_1)$ and dominates $\tilde{\gamma}x \geq \tilde{\gamma}_0$ w.r.t. P from Lemma 4. Concerning the only if condition, assume $\tilde{\gamma}x \geq \tilde{\gamma}_0$ to be strictly dominated w.r.t. P by $\bar{\gamma}x \geq \bar{\gamma}_0$, and let $(\bar{\gamma}, \bar{\gamma}_0, \bar{u}, \bar{v}, \bar{u}_0, \bar{v}_0)$ be a feasible solution of (5)–(9) yielding the dominating cut. Then, there exist $\mu \in \mathbb{R}_+^m \setminus \{0\}$ and $\mu_0 > 0$ such that $\tilde{\gamma} = \mu A + \mu_0 \bar{\gamma}$, $\tilde{\gamma}_0 = \mu b + \mu_0 \bar{\gamma}_0$. Hence $(\tilde{\gamma}, \tilde{\gamma}_0, \tilde{u}, \tilde{v}, \tilde{u}_0, \tilde{v}_0)$ is a feasible solution of (5)–(9) yielding the dominated cut, where $\hat{u} = \mu + \mu_0 \bar{u}$, $\hat{v} = \mu + \mu_0 \bar{v}$, $\hat{u}_0 = \mu_0 \bar{u}_0$, $\hat{v}_0 = \mu_0 \bar{v}_0$ and $S(\hat{u}) \cap S(\hat{v}) \neq \emptyset$. \square

Corollary 6 *Let $(\gamma, \gamma_0, u, v, u_0, v_0)$ be an optimal solution of the CGLP with normalization (12), yielding a cut violated by x^* (i.e., $\gamma x^* - \gamma_0 < 0$). Then $S(u) \cap S(v) = \emptyset$.*

Note that the above corollary holds even if the CGLP cone is truncated with a more general normalization than (12), e.g., the one to be discussed in Section 5.

4. Redundancy hurts

In the attempt of finding a way to get rid of the “weak vertices” in the CGLP, we looked for more combinatorial properties that allow us to force some Farkas multipliers to zero, thus avoiding that certain “bad” constraints participate in the definition of the optimal disjunctive cut. Our starting point was the observation that a previous disjunctive cut generated for a certain disjunction, is certainly “bad” and it should not to be used in the next iteration if the same disjunction is used⁴, in the sense that we can force its Farkas multiplier to zero. However, this multiplier is nonzero precisely in the cases where the cut rank keeps increasing at each iteration.

A more careful analysis of Example 1 reveals a more general property that allows one to classify as “bad” certain constraints. Indeed, consider the role of constraint (a1) with respect to the left-branch polytope P_0 . This constraint is clearly redundant, i.e., it can be removed without affecting P_0 (note that this is not the case if the original P is considered). However, if constraint (a1)

⁴ Assuming the original set of inequalities is the same.

participates with a positive multiplier to the definition of the disjunctive cut whereas constraint (a3) does not (i.e., if $u_1 > 0$ and $u_3 = 0$), then the cut itself has to be valid for the point $x_1 = 0, x_2 = 1/2$ and cannot be “pushed” any further inside P_0 . This is precisely what happens for the weak cuts (c3) and (c1), that cannot be supporting for P_0 precisely because of the bad choice $u_1 > 0$.

The role of redundancy is formally stated as follows. Loosely speaking, a *redundant* constraint for a polyhedron is a constraint whose removal does not enlarge the polyhedron itself. More specifically, given a system of inequalities $\hat{A}x \geq \hat{b}$, with $\hat{A} \in \mathfrak{R}^{q \times n}, \hat{b} \in \mathfrak{R}^q$, and its associated polyhedron $Q = \{x \in \mathfrak{R}^n : \hat{A}x \geq \hat{b}\}$, for any $i \in \{1, \dots, q\}$ let \hat{A}_I and \hat{b}_I denote, respectively, the submatrix of \hat{A} and the subvector of \hat{b} whose rows are indexed by the index set $I = \{1, \dots, q\} \setminus \{i\}$, and define $Q_I = \{x \in \mathfrak{R}^n : \hat{A}_I x \geq \hat{b}_I\}$. Then the constraint $\hat{a}_i x \geq \hat{b}_i$ corresponding to the row i of the system $\hat{A}x \geq \hat{b}$ is *redundant* for Q if $Q_I = Q$. In particular, we say that the constraint $\hat{a}_i x \geq \hat{b}_i$ is *strictly redundant* for Q if there exists $(\lambda_I, \delta) \in \mathfrak{R}_+^q$, with $\delta > 0$, such that $\hat{a}_i = \lambda_I \hat{A}_I$ and $\hat{b}_i = \lambda_I \hat{b}_I - \delta$. Note that, for any *strictly redundant* constraint $\hat{a}_i x \geq \hat{b}_i$, if $0x \geq -1$ can be obtained as conic combination of constraints $\hat{A}_I x \geq \hat{b}_I$, then $\hat{a}_i x \geq \hat{b}_i$ can be obtained as conic combination of $\hat{A}_I x \geq \hat{b}_I$ as well.

Proposition 7 *If a constraint that is strictly redundant for P_0 (resp. P_1) is used in the cut derivation with a nonzero multiplier, then the resulting disjunctive cut is nonsupporting in P_0 (resp. P_1).*

Proof. Let $(\tilde{\gamma}, \tilde{\gamma}_0, \tilde{u}, \tilde{u}_0, \tilde{v}, \tilde{v}_0)$ be a feasible solution of the CGLP, with $\tilde{u}_i > 0$, and assume that constraint i is strictly redundant for P_0 . Then there exists $(\lambda_I, \lambda_0, \delta) \in \mathfrak{R}_+^{m+1}$, with $\delta > 0$ such that $a_i = \lambda_I A_I - \lambda_0 \pi$ and $b_i = \lambda_I b_I - \lambda_0 \pi_0 - \delta$. By using equations (5) and (7), we get

$$\begin{aligned} \tilde{\gamma} &= \tilde{u}_I A_I + \tilde{u}_i a_i - \tilde{u}_0 \pi = (\tilde{u}_I + \tilde{u}_i \lambda_I) A_I - (\tilde{u}_0 + \tilde{u}_i \lambda_0) \pi \\ \tilde{\gamma}_0 &= \tilde{u}_I b_I + \tilde{u}_i b_i - \tilde{u}_0 \pi_0 = (\tilde{u}_I + \tilde{u}_i \lambda_I) b_I - (\tilde{u}_0 + \tilde{u}_i \lambda_0) \pi_0 - \tilde{u}_i \delta \end{aligned}$$

Thus, for each $x \in P_0$ we have $\tilde{\gamma}x - \tilde{\gamma}_0 = (\tilde{u}_I + \tilde{u}_i \lambda_I)(A_I x - b_I) - (\tilde{u}_0 + \tilde{u}_i \lambda_0)(\pi x - \pi_0) + \tilde{u}_i \delta \geq \tilde{u}_i \delta > 0$, and this shows that cut $\tilde{\gamma}x \geq \tilde{\gamma}_0$ is nonsupporting in P_0 . In the same way it can be shown that if $v_h > 0$ for a constraint h strictly redundant for P_1 , then the cut $\tilde{\gamma}x \geq \tilde{\gamma}_0$ does not support P_1 . \square

By definition, a redundant constraint for P_0 or P_1 can be obtained as a conic combination of other constraints. If the sum of the multipliers in the conic combination is greater than 1, then using a redundant constraint is cheaper (with respect to normalization (12)) than using the constraints that generate it, hence a redundant constraint can in fact be preferred by the CGLP. This is formally proved by the following theorem dealing with redundancy for P_0 (the case dealing with P_1 being perfectly analogous).

Theorem 8 *Assume that constraint $a_i x \geq b_i$ is redundant for P_0 , hence it can be obtained as a conic combination of constraints $A_I x \geq b_I, -\pi x \geq -\pi_0$, with multipliers $(\lambda_I, \lambda_0) \in \mathfrak{R}_+^m$; i.e., $a_i = \lambda_I A_I - \lambda_0 \pi, b_i = \lambda_I b_I - \lambda_0 \pi_0$. Further, let $(\tilde{\gamma}, \tilde{\gamma}_0, \tilde{u}_I, \tilde{u}_i, \tilde{u}_0, \tilde{v}, \tilde{v}_0)$ be a feasible solution of the CGLP with normalization (12), corresponding to a cut $\tilde{\gamma}x \geq \tilde{\gamma}_0$ violated by x^* , and assume $\tilde{u}_i > 0$. Then*

- (i) there exist $\theta > 0$ and a feasible solution $(\tilde{\gamma}, \tilde{\gamma}_0, \tilde{u}_I, \tilde{u}_i, \tilde{u}_0, \tilde{v}, \tilde{v}_0)$ of the CGLP with normalization (12) such that $\tilde{u}_i = 0$, $\tilde{\gamma} := \bar{\gamma}/\theta$, $\tilde{\gamma}_0 = \bar{\gamma}_0/\theta$;
(ii) $\bar{\gamma}x^* - \bar{\gamma}_0 = \theta(\tilde{\gamma}x^* - \tilde{\gamma}_0)$, and $\theta > 1$ if and only if $1\lambda_I + \lambda_0 > 1$.

Proof. Since $(\bar{\gamma}, \bar{\gamma}_0, \bar{u}_I, \bar{u}_i, \bar{u}_0, \bar{v}, \bar{v}_0)$ is feasible for the CGLP with normalization (12), writing $a_i x \geq b_i$ in terms of the multipliers (λ_I, λ_0) one gets

$$\begin{aligned}\bar{\gamma} &= (\bar{u}_I + \bar{u}_i \lambda_I) A_I - (\bar{u}_0 + \bar{u}_i \lambda_0) \pi = \bar{v} A + \bar{v}_0 \pi \\ \bar{\gamma}_0 &= (\bar{u}_I + \bar{u}_i \lambda_I) b_I - (\bar{u}_0 + \bar{u}_i \lambda_0) \pi_0 = \bar{v} b + \bar{v}_0 (\pi_0 + 1)\end{aligned}$$

while from the normalization condition $1\bar{u} + 1\bar{v} + \bar{u}_0 + \bar{v}_0 = 1$ one obtains

$$1(\bar{u}_I + \bar{u}_i \lambda_I) + (\bar{u}_0 + \bar{u}_i \lambda_0) + 1\bar{v} + \bar{v}_0 = \theta := 1 + \bar{u}_i (1\lambda_I + \lambda_0 - 1).$$

Since cut $\bar{\gamma}x \geq \bar{\gamma}_0$ is violated, one must have $\bar{u}_0 + \bar{v}_0 > 0$, hence $\theta > 0$ holds. Therefore, one can define the nonnegative quantities $\tilde{u}_I := (\bar{u}_I + \bar{u}_i \lambda_I)/\theta$, $\tilde{u}_i = 0$, $\tilde{u}_0 := (\bar{u}_0 + \bar{u}_i \lambda_0)/\theta$, $\tilde{v} = \bar{v}/\theta$, $\tilde{v}_0 = \bar{v}_0/\theta$, $\tilde{\gamma} := \bar{\gamma}/\theta$, $\tilde{\gamma}_0 = \bar{\gamma}_0/\theta$, thus getting a feasible solution of the CGLP that satisfies normalization (12) and such that $\bar{\gamma}x^* - \bar{\gamma}_0 = \theta(\tilde{\gamma}x^* - \tilde{\gamma}_0)$. Moreover, being $\bar{u}_i > 0$, one has $\theta > 1$ if and only if $1\lambda + \lambda_0 > 1$. \square

The above theorem shows that redundant constraints do not introduce new cuts, but just scaled copies of already-existing cuts that may have a better objective function (violation). Loosely speaking, redundant constraints can “cheat” normalization (12), in the sense that they can create vertices of the CGLP corresponding to scaled copies of cuts that are strictly dominated but more attractive (i.e., with a better objective function value) than the dominating ones.

A very natural way to cope with redundancy is to just eliminate the redundant constraints from the CGLP, or equivalently to fix their Farkas multipliers to zero⁵. In Example 1, the CGLP without redundant constraints has only 9 extreme rays and 9 vertices (instead of 117), and only one of them corresponds to a violated constraint – namely, the facet-defining cut (c2).

At a first glance, this example seems to suggest that only *strictly redundant* (i.e., nonsupporting) constraints should be avoided in the cut generation. However, redundant constraints should be avoided even in case they are supporting, as shown by the example reported below.

Example 3. For the simple ILP

$$\begin{aligned}\min & -x_1 - 2x_2 + 10x_3 \\ \text{(a1)} & 4x_1 - 4x_2 \geq -2 \\ \text{(a2)} & -2x_1 - 2x_2 - x_3 \geq -3 \\ \text{(a3)} & 3x_1 - 2x_2 - x_3 \geq -1 \\ \text{(a4)} & x_1 \geq 0 \\ \text{(a5)} & x_2 \geq 0 \\ \text{(a6)} & x_3 \geq 0\end{aligned}$$

⁵ Of course, a more efficient approach would be to check redundancy “on the fly”, by only considering those constraints that have a nonzero Farkas multiplier in the optimal CGLP solution; the implementation of this approach is however outside the scope of the present paper, and is left to future work on the subject.

only two cuts violated by the optimal solution of the LP relaxation $x^* = (\frac{1}{2}, 1, 0)$ can be derived from disjunction $x_1 \leq 0$ OR $x_1 \geq 1$, namely:

- (c1) $2x_2 \leq 1$, corresponding to the basic solution of the CGLP $(u_1, v_2, v_6, u_0, v_0)$, of value $z_1 = -\frac{2}{13}$ (optimal);
- (c2) $2x_2 + x_3 \leq 1$, corresponding to the basic solution of the CGLP $(u_3, v_2, v_6, u_0, v_0)$, of value $z_2 = -\frac{1}{7}$ (nonoptimal).

$\text{Conv}(P_0 \cup P_1)$ has 6 vertices, namely $V_1 = (0, 0, 0)$, $V_2 = (0, \frac{1}{2}, 0)$, $V_3 = (0, 0, 1)$, $V_4 = (\frac{3}{2}, 0, 0)$, $V_5 = (1, \frac{1}{2}, 0)$, and $V_6 = (1, 0, 1)$. In the reverse polar space (γ, γ_0) , the projected CGLP cone has only 5 extreme rays that correspond to the facets of $\text{conv}(P_0 \cup P_1)$. In space $(\gamma, \gamma_0, u, v, u_0, v_0)$, instead, the CGLP cone has 33 extreme rays that correspond to 33 vertices once normalization (12) is applied. In the optimal basis, constraint (a1) (which is redundant but supporting for P_0) is used with $u_1 > 0$, and the corresponding cut (c1) supports both P_0 and P_1 in V_2 and V_5 , respectively, but it is not facet-defining. The CGLP without redundant constraints (in particular without (a1)) has only 10 extreme rays and 10 vertices (instead of 33), and only one of them corresponds to a violated constraint—namely, the facet-defining cut (c2). Note that cut (c2) dominates cut (c1) and is facet-defining since it supports $\text{conv}(P_0 \cup P_1)$ in the 3 affinely independent vertices V_2, V_3 , and V_5 . For illustration purposes, 3 extreme rays and an interior point of the CGLP cone are reported below.

	γ_1	γ_2	γ_3	γ_0	u_1	u_2	u_3	u_4	u_5	u_6	v_1	v_2	v_3	v_4	v_5	v_6	u_0	v_0
(r_1)	0	-2	-1	-1	0	0	1	0	0	0	0	1	0	0	0	0	3	2
(r_2)	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	1	0	0
(p_1)	0	-2	0	-1	0	0	1	0	0	1	0	1	0	0	0	1	3	2
(r_3)	0	-2	0	-1	1/2	0	0	0	0	0	0	1	0	0	0	1	2	2

The weak cut (c1) is strictly dominated w.r.t. P by (c2), as the interior point p_1 is just the sum of the extreme rays r_1 and r_2 , the latter corresponding to the original constraint $x_3 \geq 0$. However, the redundant constraint (a1) creates an extremal copy of the weak cut – the extreme ray r_3 – which turns out to be the optimal vertex once normalization (12) is applied.

□

The previous discussion shows that extreme rays of the CGLP cone in the extended space $(\gamma, \gamma_0, u, v, u_0, v_0)$ may be nonextremal when projected onto the (γ, γ_0) space, and redundant constraints can add to the cone several extreme-rays corresponding to very weak cuts. The SNC normalization (12) simply maps extreme rays to vertices, and creates a possibly “wrong” ranking among the vertices. Unfortunately, as far as we know no normalization equation is able of truncating the CGLP cone so as to guarantee that an optimal CGLP vertex in the extended space remains a vertex when projected in the (γ, γ_0) space. For instance, let consider normalizations of the form

$$\gamma(q - x^*) = 1, \tag{13}$$

which have been deeply investigated in Bonami [11]. Balas and Perregaard [6] proved that, if $q \in \text{conv}(P_0 \cup P_1)$, then the CGLP truncated with (13) has a finite optimum and that there *exists* an optimal vertex of the resulting polyhedron in the extended space whose projection in the natural reverse polar space (γ, γ_0) remains extremal. However, this does not imply that any optimal vertex in the extended space is a vertex in the projected space – hence even normalization (13) could not help in finding a facet-defining cut.

Example 4. (Example 3 continued)

Let consider again the simple ILP discussed in Example 3. If the corresponding CGLP cone is truncated with normalization (13), with $q = (0, 0, 0)$, the resulting polyhedron has 60 extreme rays and 20 vertices. As before, only two vertices correspond to violated cuts, namely:

- i) the basic solution $(u_1, v_2, v_6, u_0, v_0)$, of value $z_1 = -\frac{1}{2}$ (optimal), corresponding to the weak cut (c1);
- ii) the basic solution $(u_3, v_2, v_6, u_0, v_0)$, of value $z_2 = -\frac{1}{2}$ (optimal), corresponding to the facet-defining cut (c2).

So, even in this case, the separation procedure could select the weak cut (c1), since the choice of q makes (c1) and (c2) completely equivalent in terms of objective function.

□

4.1. Empirical Analysis

In our third set of experiments we eliminated redundant constraints in a trivial way (i.e., by solving LPs) before solving the CGLP. To get a clearer picture, we did not project the separation problem on the support of x^* since such a projection makes the definition of what is redundant and what is not less clear.

The results are reported in Table 3 and show that removing redundant constraints is indeed very useful. Besides an average improvement in the percentage gap closed of around 2.5%, only for two problems, namely **bel15** and **gt2**, the “Classical” SNC is slightly better than the “No redundancy” SNC version, while for some single problems the improvement is very large, up to 13% for instance **p0033**. Concerning the average cardinality of the dual support of the cut there is a slight increase in the “No redundancy” version which however does not seem relevant.

4.2. Working on the Support

Projecting the separation problem into the support of x^* has of course the advantage of dealing with a problem of smaller size. However, according to our experience the projection can enlarge the set of redundant constraints in a way that vanishes part of the positive effects related to their removal. A possible explanation of this behavior is that projection may hide the redundancy of some bound

Table 3. “Classical” SNC vs. “No redundancy” SNC with no projection.

Instance	“Classical” SNC			“No redundancy” SNC		
	# cuts	%gap	S	# cuts	%gap	S
bell3a	71	70.74	64.65	54	70.74	66.19
bell5	188	94.12	16.83	189	93.54	15.80
blend2	197	30.49	71.42	212	30.63	119.90
flugpl	93	18.34	6.45	90	18.83	6.48
gt2	218	94.13	58.11	167	93.68	63.16
lseu	171	42.46	23.86	184	45.10	30.96
*markshare1	77	0.00	55.99	77	0.00	56.00
mod008	107	15.46	304.18	107	15.48	304.19
p0033	116	57.25	8.75	126	70.32	10.99
p0201	692	92.53	23.40	757	98.31	37.44
roul	349	29.46	189.07	384	31.93	202.18
*stein27	251	0.00	7.29	249	0.00	6.46
vpm1	267	50.62	11.13	282	54.55	11.10
vpm2	390	74.73	24.23	376	76.47	22.82
avg.	238.250	55.861	66.840	244.000	58.298	74.267

constraints, hence weakening the final disjunctive cut. Indeed, consider a variable x_k restricted to being nonnegative and such that $x_k^* = 0$. If x_k is projected away with the aim of computing coefficient γ_k afterwards through (10), then we loose any control on the Farkas variables associated with the constraint $x_k \geq 0$, say $u_{i(k)}$ and $v_{i(k)}$. In fact, if it happens that constraint $x_k \geq 0$ is redundant, it is very useful to keep explicitly constraints $\gamma_k = uA_k - u_0\pi_k = vA_k + v_0\pi_k$ in the CGLP and to impose the additional requirement $u_{i(k)} = 0$ and/or $v_{i(k)} = 0$.

As the above property seems to be crucial for the variable bounds, we defined an *extended support* of x^* by avoiding projecting away any variable whose bound condition is (tight in x^* and) redundant. The results when using the extended support are reported in Table 4, where %*supp* indicates the average percentage of the x variables which are kept in the (extended) support.

Table 4. “Classical” SNC vs. “No redundancy” SNC with cuts separated projected on the support.

Instance	“Classical” SNC				“No redundancy” support				“No redundancy” ext. support			
	# cuts	%gap	%supp	S	# cuts	%gap	%supp	S	# cuts	%gap	%supp	S
bell3a	71	70.74	69.25	43.72	88	70.74	69.32	44.82	54	70.74	65.61	44.60
bell5	178	94.29	72.69	11.75	207	94.62	72.88	13.32	180	94.29	71.64	11.99
blend2	192	30.51	53.06	8.10	200	30.99	53.54	10.84	193	30.53	53.99	8.34
flugpl	92	18.36	86.11	5.85	93	18.94	86.11	5.89	93	18.86	86.29	5.95
gt2	196	93.46	18.30	10.28	191	94.13	18.14	10.58	187	93.88	20.00	13.10
lseu	196	41.33	29.44	9.17	191	40.16	27.08	12.28	178	43.45	29.41	9.08
*markshare1	74	0.00	11.94	1.39	130	0.00	13.39	2.56	77	0.00	12.59	1.69
mod008	139	17.05	4.51	12.41	136	17.70	4.42	12.17	157	19.13	5.85	14.43
p0033	113	67.86	55.76	4.81	106	70.32	55.76	5.74	146	70.29	58.84	5.89
p0201	767	93.82	45.02	13.43	873	81.59	43.43	25.83	769	100.00	48.93	13.39
roul	434	24.26	42.19	68.07	355	6.56	38.11	58.23	353	30.88	69.46	140.29
*stein27	252	0.00	93.70	6.53	252	0.00	93.70	6.68	251	0.00	93.61	7.13
vpm1	263	55.84	62.14	5.39	275	50.18	62.25	6.30	259	57.63	65.18	6.60
vpm2	403	74.96	64.74	17.27	377	75.30	65.08	18.10	373	75.84	67.15	17.71
avg.	253.667	56.873	50.268	17.521	257.667	54.269	49.677	18.675	245.167	58.793	53.529	24.281

The first part of Table 4 reports the same figures as Table 1, plus a column which gives the percentage size (w.r.t. the nominal size) of the CGLP projected

on the support. By comparing the first and second part of the table, we note that the gain shown by Table 3 due to the redundancy removal is lost here (the average gap closed of 56.873% deteriorates to 54.269%), thus confirming our intuition about the smaller precision of the redundancy test in such a case. However, the situation is totally recovered using the extended support as defined above. Indeed, the percentage value 56.873 improves to 58.793, and the average size of the support does not increase much (from 50.268% to 53.529%). The only large increase on the size of the CGLP arises for problem *rou1*, which is in fact a very instructive case: the “Classical” SNC closed 24.26% of the gap, the “No redundancy” SNC version on the support closes only 6.56%, while in the extended support the situation is totally recovered (and improved) with 30.88% gap closed. For such a particular instance the extended support size is also substantially different from the support size, namely 69.46% with respect to 42.19%. Our interpretation is the following: in order to forbid the use of some variable bounds in the derivation of the cut we have to enlarge the support considerably (half of the projected variables are re-inserted) with the overall effect of generating much stronger cuts.

Detailed results on instance *p0201* are given on Figures 11–14.




Fig. 11. “Classical” SNC vs. “No redundancy” SNC: dual bound for instance *p0201*.




Fig. 12. “Classical” SNC vs. “No redundancy” SNC: avg. cut density for instance *p0201*.




Fig. 13. “Classical” SNC vs. “No redundancy” SNC: avg. cardinality of $S(u, v)$ for instance *p0201*.




Fig. 14. “Classical” SNC vs. “No redundancy” SNC: cut rank for instance *p0201*.

5. A different normalization

As shown in the previous sections, the standard normalization has the main advantage of generating low-rank inequalities, which is in general a desirable property. As a matter of fact, it has been recently showed that rank-1 inequalities alone are able to close a large portion of the integrality gap (see, e.g., Fischetti and Lodi [17], Balas and Saxena [8], Dash, Günlük and Lodi [16]). When normalization (12) is applied, the norm of the separated cuts tends to be smaller with respect to the constraints used for their generation, and small-norm constraints are implicitly penalized by the normalization itself. Thus, high-rank constraints are selected in the cut derivation only if needed, hence generating weak cuts does not hurt the overall separation procedure in that these cuts are less likely to be used in the next iterations. However, as stated in Section 4, the standard normalization creates a ranking among the CGLP vertices which depends on the scaling of the constraints, i.e., the overall separation procedure is heavily affected by the scaling of the constraints in the original formulation. To overcome the latter drawback, one can replace the standard normalization with the following *Euclidean Normalization* (EN):

$$\sum_{i=1}^m \|a_i\|u_i + \sum_{i=1}^m \|a_i\|v_i + \|\pi\|u_0 + \|\pi\|v_0 = 1, \quad (14)$$

where $\|t\|$ denotes the Euclidean norm of vector t .

Lemma 9 *Let $\tilde{A}x \geq \tilde{b}$ be a scaled copy of system $Ax \geq b$ where, for all $i \in \{1, \dots, m\}$, $\tilde{a}_i := a_i/K_i$ and $\tilde{b}_i := b_i/K_i$, with $K_i > 0$. For any solution of the CGLP with normalization (14) corresponding to a cut $\gamma x \geq \gamma_0$, there exists a solution of the CGLP associated with the system $\tilde{A}x \geq \tilde{b}$, still with normalization (14), corresponding to the same cut.*

Proof. Let $(\gamma, \gamma_0, u, v, u_0, v_0)$ be a solution of the CGLP with normalization (14), and for all $i \in \{1, \dots, m\}$ define $\tilde{u}_i = K_i u_i$ and $\tilde{v}_i = K_i v_i$. From equations (5) we obtain

$$\gamma = uA - u_0\pi = \sum_{i=1}^m u_i a_i - u_0\pi = \sum_{i=1}^m (K_i u_i)(a_i/K_i) - u_0\pi = \tilde{u}\tilde{A} - u_0\pi.$$

Analogously, from (6) we get $\gamma = \tilde{v}\tilde{A} + v_0\pi$ and from (7)–(8) we have $\gamma_0 = \tilde{u}\tilde{b} - u_0\pi_0 = \tilde{v}\tilde{b} + v_0(\pi_0 + 1)$. Hence $(\gamma, \gamma_0, \tilde{u}, \tilde{v}, u_0, v_0)$ is a feasible solution of the CGLP associated with system $\tilde{A}x \geq \tilde{b}$ yielding the same cut as $(\gamma, \gamma_0, u, v, u_0, v_0)$. Since $\|\tilde{a}_i\|\tilde{u}_i + \|\tilde{a}_i\|\tilde{v}_i = \|a_i\|u_i + \|a_i\|v_i \forall i \in \{1, \dots, m\}$, then $(\gamma, \gamma_0, \tilde{u}, \tilde{v}, u_0, v_0)$ fulfills normalization (14) as well. \square

The above lemma shows that the CGLP with Euclidean normalization is not affected by scaling issues. Moreover, the CGLP with (14) is the same as the CGLP with the standard normalization for a system $\tilde{A}x \geq \tilde{b}$ where all the constraints have been scaled in order to have Euclidean norm equal to 1 (i.e.,

$\|\tilde{a}_i\| = 1 \forall i \in \{1, \dots, m\}$). By replacing normalization (12) with (14) we are losing the implicit penalization of high-rank inequalities hidden in the standard normalization. However, the Euclidean normalization associates penalties with the Farkas multipliers which are in some way related to the structure of the corresponding constraints.

Computational results comparing the standard and the Euclidean normalization are reported in Table 5.

Table 5. SNC normalization vs. Euclidean normalization on MIPLIB instances.

Instance	SNC normalization			Euclidean normalization		
	# cuts	%gap	S	# cuts	%gap	S
bell3a	71	70.74	43.72	71	70.74	34.95
bell5	178	94.29	11.75	213	93.01	13.78
blend2	192	30.51	8.10	182	31.80	9.04
flugpl	92	18.36	5.85	89	18.75	6.41
gt2	196	93.46	10.28	166	93.68	12.03
lseu	196	41.33	9.17	177	42.35	10.40
*markshare1	74	0.00	1.39	75	0.00	1.77
mod008	139	17.05	12.41	135	18.27	11.98
p0033	113	67.86	4.81	97	70.32	4.28
p0201	767	93.82	13.43	741	99.07	28.72
rout	434	24.26	68.07	412	7.92	71.53
*stein27	252	0.00	6.53	254	0.00	12.39
vpm1	263	55.84	5.39	256	67.17	5.59
vpm2	403	74.96	17.27	379	71.68	12.34
avg.	253.667	56.873	17.521	243.167	57.063	18.421

Table 5 shows that the proposed normalization slightly improves on average over the “Classical” SNC with the only relevant exception of instance **rout** on which the gap closed decreases from 24.26% to 7.92%. (Without such an instance, the Euclidean normalization would exhibit an average improvement of 1.7%.) The negative result for **rout** is not a total surprise since Table 4 had shown that both redundancy and projection play crucial roles for such an instance. How much one of the two is more important than the other is hard to state, but the gap closed without projection is 20.10%, i.e., considerably smaller than the 29.46% of SNC on the complete model (see Table 3). A natural question is now if the removal of redundant constraints would also help the proposed Euclidean normalization. The answer is empirically given by the computational results in Table 6, by comparing the “Classical” SNC and the Euclidean normalization after removal of redundant constraints and by working in the extended support. According to the table, removing redundant constraints helps also the Euclidean normalization as the average gap closed raised from 58.793% (“No redundancy” SNC) to 60.014% (“No redundancy” EN). Moreover, for instance **rout** the gap closed is 29.42%, i.e., very similar to the SNC one (30.88%).

Table 6. SNC vs. Euclidean normalization on MIPLIB instances. Redundant constraints removed in both versions and projection on the extended support of x^* .

Instance	“No redundancy” SNC			“No redundancy” EN		
	# cuts	%gap	S	# cuts	%gap	S
bell3a	54	70.74	44.60	50	70.74	44.21
bell5	180	94.29	11.99	194	94.07	15.39
blend2	193	30.53	8.34	181	32.83	9.99
flugpl	93	18.86	5.95	92	19.20	6.02
gt2	187	93.88	13.10	215	94.71	13.54
lseu	178	43.45	9.08	165	45.00	8.86
*markshare1	77	0.00	1.69	75	0.00	1.77
mod008	157	19.13	14.43	135	18.36	11.98
p0033	146	70.29	5.89	99	70.32	4.47
p0201	769	100.00	13.39	765	100.00	32.44
rout	353	30.88	140.29	351	29.42	193.85
*stein27	251	0.00	7.13	251	0.00	13.17
vpm1	259	57.63	6.60	272	68.91	6.44
vpm2	373	75.84	17.71	361	76.61	14.38
avg.	245.167	58.793	24.281	240.000	60.014	30.131

5.1. The Set Covering (special) case

As showed in Tables 5–6, associating with each constraint a penalty equal to the Euclidean norm of the constraint itself leads on average to a slight improvement over the “Classical” SNC. For a general-purpose MIP such a choice might have drawbacks as in the instance `rout` discussed above. However, if the problem has some special structure the improvement over SNC can be more consistent, as in the case e.g. of the well-known *Set Covering* problem. Indeed, set covering constraints have nonnegative coefficients only, and this property is known to be inherited by nontrivial valid inequalities, including the disjunctive cuts we can separate through our procedure.

We performed additional computational experiments on a test-bed of Set Covering instances taken from the OR-Library [9], and the results are reported in Table 7. The improvement in the percentage gap closed by EN with respect

Table 7. SNC normalization vs. Euclidean normalization on SCP instances.

Instance	SNC normalization			Euclidean normalization		
	# cuts	%gap	S	# cuts	%gap	S
scpnre1	904	13.67	89.22	951	17.35	93.62
scpnre2	963	9.38	95.42	997	12.51	98.14
scpnre3	923	15.14	91.41	944	18.13	92.82
scpnre4	878	13.25	85.99	897	15.70	87.82
scpnre5	889	16.84	87.77	935	21.03	91.16
scpnrf1	678	10.23	67.75	682	12.62	67.77
scpnrf2	655	9.62	65.42	689	12.90	68.50
scpnrf3	586	12.08	58.34	617	15.58	60.93
scpnrf4	664	10.21	66.35	692	12.59	68.91
scpnrf5	661	8.63	66.05	700	11.85	69.70
avg.	780.100	11.905	77.372	810.400	15.026	79.937

to SNC is quite substantial as it ranged from 2.38% to 4.19%, with an average of 3.12%.

Figures 15–18 describe the behavior of the two normalizations on the particular instance `scpnre5`. Figures 16 and 18 are particularly interesting. The




Fig. 15. SNC vs. Euclidean normalization: dual bound for instance `scpnre5`.




Fig. 16. SNC vs. Euclidean normalization: avg. cut density for instance `scpnre5`.




Fig. 17. SNC vs. Euclidean normalization: avg. cardinality of $S(u, v)$ for instance `scpnre5`.




Fig. 18. SNC vs. Euclidean normalization: cut rank for instance `scpnre5`.

considerably higher rank of the cuts generated using the Euclidean normalization with respect to those obtained through SNC (see Figure 18) does not correspond at all to denser cuts. Indeed, as shown in Figure 16, the former cuts are consistently sparser than the latter. (Cuts generated using SNC are fully dense: the number of variables of the instance is 5,000 and the number of non-zero coefficients is almost always very close to 5,000 too; see Figure 16).

References

1. E. Balas, A modified lift-and-project procedure, *Mathematical Programming* **79** (1997), 19–31.
2. E. Balas, Disjunctive programming, *Annals of Discrete Mathematics* **5** (1979), 3–51.
3. E. Balas, Disjunctive programming: properties of the convex hull of feasible points, *Discrete Applied Mathematics* **89** (1998) 3–44.

4. E. Balas, S. Ceria, G. Cornuéjols, Mixed 0–1 programming by lift-and-project in a branch-and-cut framework, *Management Science* **42** (1996), 1229–1246.
5. E. Balas, R. Jeroslow, Strengthening Cuts for Mixed Integer Programs, *European Journal of Operations Research* **4** (1980), 224–234.
6. E. Balas and M. Perregaard, Lift-and-project for mixed 0–1 programming: recent progress, *Discrete Applied Mathematics* **123** (2002), 129–154.
7. E. Balas and M. Perregaard, A precise correspondence between lift-and-project cuts, simple disjunctive cuts, and mixed integer gomory cuts for 0–1 programming, *Mathematical Programming Series B* **94** (2003), 221–245.
8. E. Balas and A. Saxena, Optimizing over the split closure, *Mathematical Programming Series A*, to appear, doi:10.1007/s10107-006-0049-5.
9. J. E. Beasley, OR–Library: a collection of test data sets for a variety of Operations Research (OR) problems, <http://people.brunel.ac.uk/~mastjjb/jeb/info.html>.
10. R. E. Bixby, S. Ceria, C. M. McZeal, M. W. P. Savelsbergh, An updated mixed integer programming library: MIPLIB 3.0, <http://www.caam.rice.edu/~bixby/miplib/miplib.html>.
11. P. Bonami, Étude et mise en oeuvre d’approches polyédriques pour la résolution de programmes en nombres entiers ou mixtes généraux, *PhD Thesis, Université de Paris 6* (2003).
12. S. Ceria, J. Soares, Disjunctive cuts for mixed 0–1 programming: duality and lifting, *GSB, Columbia University* (1997).
13. V. Chvátal, Edmonds polytopes and a hierarchy of combinatorial problems, *Discrete Mathematics* **4** (1973), 305–337.
14. W. J. Cook, R. Kannan, and A. Schrijver, Chvátal closures for mixed integer programming problems, *Mathematical Programming Series A* **47** (1990), 155–174.
15. T. Christof, A. Löbel, PORTA - Polyhedron Representation Transformation Algorithm, <http://www.zib.de/Optimization/Software/Porta/>.
16. S. Dash, O. Günlük, A. Lodi, “On the MIR closure of polyhedra”, in M. Fischetti, D.P. Williamson, Eds., *Integer Programming and Combinatorial Optimization - IPCO 2007*, Lecture Notes in Computer Science 4513, Springer-Verlag, Berlin Heidelberg, 2007, 337–351.
17. M. Fischetti and A. Lodi, Optimizing over the first Chvátal closure, *Mathematical Programming Series B* **110** (2007), 3–20.
18. R. E. Gomory, An algorithm for the mixed integer problem, RM-2597, The RAND Corporation (1960).