It's Farkas All the Way Down: A Generalized Farkas Lemma and Its Implications

Ted Ralphs[1]

[1]COR@L Lab, Department of Industrial and Systems Engineering, Lehigh University

MIP Workshop, 3 June 2024

Attributions

Many current and former students contributed in various ways to the development of this long line of research.

Current/Former Students/Postdocs

- Federico Battista
- Suresh Bolusani
- Scott DeNegre
- Samira Fallah
- Menal Gúzelsoy
- Anahita Hassanzadeh
- Shannon Kelley
- Ashutosh Mahajan
- Sahar Tahernejad
- Yu Xie

Thanks!

"It's All Just Duality"

Quote from the Internet: *Duality* is a woefully overloaded mathematical term for a relation that groups elements of a set into "dual" pairs.
Bold claim: Many (most?) duality concepts can be seen as roughly "isomorphic".

Duality Concepts

- Sets: Projection/complement, intersection/union
- Conic duality: Cones and their duals, convexity/nonconvexity
- Farkas duality: Theorems of the alternative, empty/non-empty
- Complexity: Languages and their complements (NP vs. co-NP)
- Quantifier duality: Existential versus universal quantification
- De Morgan duality: Conjunction versus disjunction
- Weyl-Minkowski duality: V representation versus H representation
- Polarity: Optimization versus separation
- Dual problems: Primal and dual problems in optimization
- Inverses: Functions and inverses, inverse optimization inverses

Theorems About Sets

- Mathematically speaking, we can think of "solving" an optimization problem with an "exact" solution method as proving a theorem about a given set.
- The solver produces not only a solution, but also a proof.
- Let $\mathcal{S} = \{x \in \mathbb{Q}^n \mid P(x)\}$, where $P : \mathbb{Q}^n \to \{TRUE, FALSE\}$.
- The simplest question we can ask is whether $\mathcal{S}$ is non-empty

$$\mathcal{S} \overset{?}{=} \emptyset.$$

- Given function f and constant K, the related question of

$$\mathcal{S}(f, K) := \{x \in \mathcal{S} \mid f(x) < K\} \overset{?}{=} \emptyset$$

is the *decision version* of the optimization problem

$$\min_{x \in \mathcal{S}} f(x) \qquad \text{(OPT)}$$

Constructing Proofs

- What do proofs of theorems about sets look like?
  - Certifying $\mathcal{S} \neq \emptyset$ is "easy": produce a point in the set.
  - Certifying $\mathcal{S} = \emptyset$ is more difficult in general.
- The difficulty of proving a set is empty is most easily seen by re-stating the theorems we are trying to prove/disprove, as follows.

$$\mathcal{S} \neq \emptyset \Leftrightarrow \exists x \in \mathcal{S}$$
$$\mathcal{S} = \emptyset \Leftrightarrow \forall x \in \mathbb{Q}^n \; x \notin \mathcal{S} \Leftrightarrow \forall x \in \mathbb{Q}^n \; x \in \bar{\mathcal{S}}$$

- The statement that a set is non-empty is *existentially quantified*, whereas the statement that a set is empty is *universally quantified*.
- Universally quantified statements are intuitively more difficult to prove than existentially quantified ones.

De Morgan Duality

- There is a duality between existential and universal quantifiers that can be seen as one of a number of generalized forms of De Morgan's Laws.

DeMorgan's Laws

- The complement of the union is the intersection of the complements.
- The complement of the intersection is the union of the complements.

- These laws can be used to equivalently formulate logical statements in different dual forms to aid in constructing proofs.

$$P(x)\ \forall x \in \mathcal{S} \Leftrightarrow \neg[\exists x \in \mathcal{S}\ \neg P(x)] \Leftrightarrow \neg \bigvee_{x \in \mathcal{S}} \neg P(x) \Leftrightarrow \bigwedge_{x \in \mathcal{S}} P(x)$$

$$\exists x \in \mathcal{S} : P(x) \Leftrightarrow \neg[\forall x \in \mathcal{S}\ P(x)] \Leftrightarrow \neg \bigwedge_{x \in \mathcal{S}} \neg P(x) \Leftrightarrow \bigvee_{x \in \mathcal{S}} P(x)$$

- Note also the duality between conjunction and disjunction.

Convexity and Nonconvexity

- Related dualities exist between between conjunction and disjunction, which are reflected in the way convex and nonconvex sets are described.
- Convex sets are described by conjunctive logic: the *intersection* of convex sets is convex.
- Nonconvex sets are described using disjunctive logic: the *union* of convex sets is nonconvex (in general).
- Proving that a point is *not* in a convex set is "easy," whereas doing the same for a nonconvex set is difficult in general.

Short Proofs of Emptiness

- The essence of the Farkas Lemma is that it allows to obtain a short proof that a convex set is empty.
- Consider again the polyhedron

$$\mathcal{P} = \{x \in \mathbb{R}^n_+ \mid Ax = b\}$$

given in standard form with $A \in \mathbb{Q}^{m \times n}$, $b \in \mathbb{Q}^m$.

Farkas Lemma

$$\mathcal{P} = \emptyset \Leftrightarrow \exists u \in \mathbb{R}^m \text{ s.t. } u^\top A \leq 0, u^\top b > 0$$

- Equivalently, $\mathcal{P} = \emptyset$ if and only if we can separate b from the convex cone

$$\begin{aligned} \mathcal{C} &= \{Ax \mid x \in \mathbb{R}^n_+\} \\ &= \{\beta \in \mathbb{R}^m \mid u^\top \beta \leq 0 \ \forall u \in \mathcal{C}^*\}, \end{aligned}$$

where $\mathcal{C}^* = \{u \in \mathbb{R}^m \mid u^\top A \leq 0\}$ (the *dual* of $\mathcal{C}$).

Example 1

$$6y_1 + 7y_2 + 5y_3 = 1/2$$
$$2y_1 - 7y_2 + y_3 = 1/2$$
$$y_1, y_2, y_3 \in \mathbb{R}_+$$

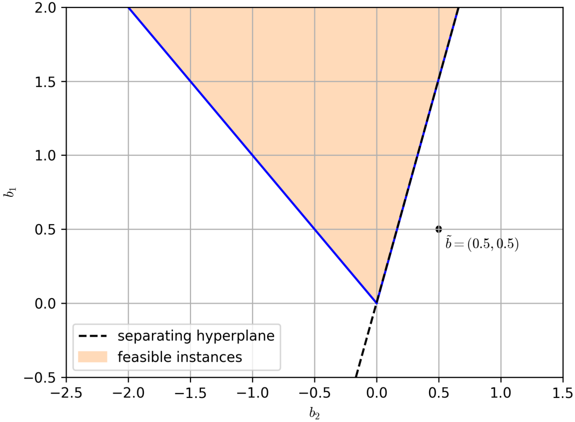

Another Interpretation

- We lift the problem into a higher dimensional space by making b a vector of variables and homogenizing.

$$\mathcal{P}^{\beta} = \{x \in \mathbb{R}^n_+, \beta \in \mathbb{R}^m \mid Ax - I\beta = 0\}$$

- Then project out the original variables to obtain $\mathcal{C}$.

$$\mathcal{C} = \mathrm{Proj}_{\beta}(\mathcal{P}^{\beta})$$

- In other words, $\mathcal{C}$ is just the set of values of β for which the linear system $Ax = \beta$ has a solution.
- Alternatively, $\mathcal{C}$ consists of the feasible members of a parametric family of linear optimization problems (LPs).
- Therefore, if we can separate b from $\mathcal{C}$, we prove that $\mathcal{P} = \emptyset$ (corresponding instance is infeasible).

Example 2

$$\mathcal{P}^\beta = \left\{ \begin{aligned} 2y_1 - 7y_2 + y_3 &= \beta_1 \\ 6y_1 + 7y_2 + 5y_3 &= \beta_2 \\ y_1, y_2, y_3 &\in \mathbb{R}_+ \end{aligned} \right\} \qquad C = \left\{ \begin{aligned} \beta_1 + \beta_2 &\geq 0 \\ -3\beta_1 + \beta_2 &\geq 0 \\ \beta &\in \mathbb{R}^2 \end{aligned} \right\}$$

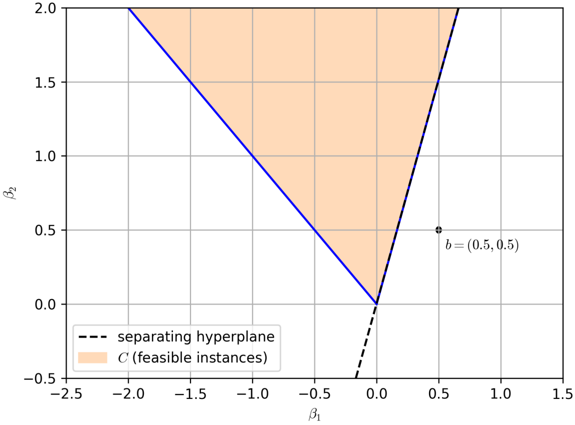

Connection to Complexity Theory

- On one hand, this is a “trick” for recasting a question about an empty set as one about a non-empty convex set (universal $\rightarrow$ existential), but there’s a bigger picture.
- We are embedding a single theorem into a *parametric class* containing both TRUE and FALSE theorems.
- The questions we are asking is being re-cast as a question of where this theorem lies relative to the set of all TRUE theorems (in the class).
- To prove the theorem is FALSE, we separate it from the set of theorems that are TRUE—this is a “dual” proof based on a separation argument.
- In the terminology of complexity theory, the set of true theorems is called a *language*.

Proofs of Optimality

- The problem (OPT) is *not* a decision problem as stated.
- We can nevertheless build a proof that the optimal solution value is K using proofs for two related theorems.

1. $\exists x \in \mathcal{S} : f(x) = K$
2. $\nexists x \in \mathcal{S} : f(x) < K \Leftrightarrow \forall x \in \mathcal{S} : f(x) \geq K$

- The fact that one of these statements is universally quantified is the reason why short proofs of optimality cannot be expected in general.

Farkas Proof of Optimality

- We now consider the case of an LP, constructed as follows.
  - Convert a^1 (the first row of A) from a constraint to the objective function.
  - Let $M = \{1, \ldots, m\}$ and $b_{M\setminus\{1\}} \in \mathbb{R}^{m-1}$ be all but the first element of b.
  - The resulting LP is $\min_{x \in \mathbb{R}^n_+} \{a^1 x \mid A_{M\setminus\{1\}} x = b_{M\setminus\{1\}}\}$.
- The problem of finding the optimal value can then be recast as $z^* = \min\{z \mid (z, b_{M\setminus\{1\}}) \in \mathcal{C}\}$.
- To prove optimality, we need to show that $(z^*, b_{M\setminus\{1\}})$ is not only a member of $\mathcal{C}$, but on its *boundary*.
- The LP optimality conditions imply $\exists u_{M\setminus\{1\}} \in \mathbb{R}^{m-1}$ s.t. ${u_{M\setminus\{1\}}}^\top A_{M\setminus\{1\}} \leq a^1$, ${u_{M\setminus\{1\}}}^\top b_{M\setminus\{1\}} = z^*$.
- This is equivalent to $\exists u \in \mathbb{R}^m$ s.t. $u^\top A \leq 0$, $u^\top (z^*, b_{M\setminus\{1\}}) = 0$, $u_1 = -1$, implying
  - $(z^*, b_{M\setminus\{1\}})$ is on the boundary of $\mathcal{C}$ and
  - the boundary is one that is in the “right direction” ($u_1 < 0$).
- The vector u is a solution to the usual LP dual problem.

Example 3

$$z^* = \min 6y_1 + 7y_2 + 5y_3$$
$$\text{s.t. } 2y_1 - 7y_2 + y_3 = 1/2$$
$$y_1, y_2, y_3 \in \mathbb{R}_+$$

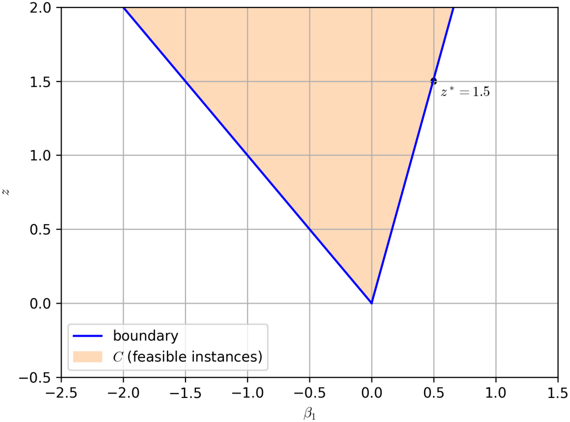

Example 4

- Note that our choice of objective was arbitrary and the same set $\mathcal{C}$ can yield proofs for other objectives.
- The figure shows that $\min_{y \in \mathbb{R}^3_+} \{2y_1 - 7y_2 + y_3 \mid 6y_1 + 7y_2 + 5y_3 = 1\} = -1$

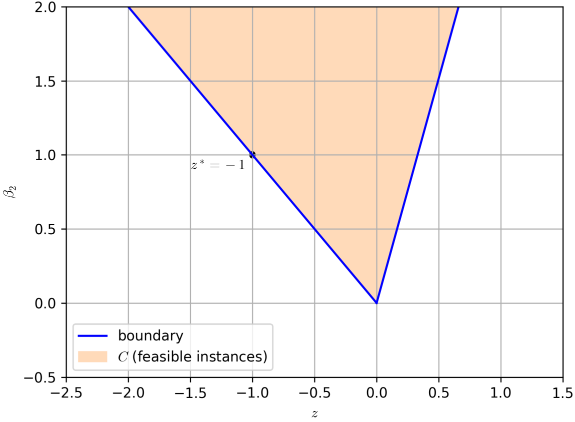

The Boundary of $\mathcal{C}$ and Optimality Conditions

- Under mild conditions, a point $\beta \in \mathbb{R}^m$ is on the boundary of $\mathcal{C}$ if and only if

$$\begin{aligned}\beta_i &= \min / \max \ \{z \in \mathbb{R} \mid (z, \beta_{M\setminus\{i\}}) \in \mathcal{C}\} \\ &= \min_{x \in \mathbb{R}^n_+} / \max_{x \in \mathbb{R}^n_+} \{a^i x \mid A_{M\setminus\{i\}} x = \beta_{M\setminus\{i\}}\}\end{aligned}$$

  for some $i \in \{1, \ldots, m\}$.

- The boundary is comprised of the graphs of the following *value functions* of LPs associated with the rows of A over their finite domain.

$$\begin{aligned}\phi_i^+(\tau) &= \max_{x \in \mathcal{P}_{\{i\}}(\tau)} a^i x \quad \forall \tau \in \mathbb{R}^{m-1} \\ \phi_i^-(\tau) &= \min_{x \in \mathcal{P}_{\{i\}}(\tau)} a^i x \quad \forall \tau \in \mathbb{R}^{m-1},\end{aligned}$$

  where

$$\mathcal{P}_{\{i\}}(\tau) = \{x \in \mathbb{R}^n_+ \mid A_{M\setminus\{i\}} x = \tau\}.$$

- The boundary is essentially a *parametric collection* of proofs of optimality.

Example 5

$$\min 6y_1 + 7y_2 + 5y_3$$
$$\text{s.t. } 2y_1 - 7y_2 + y_3 = 1/2$$
$$y_1, y_2, y_3 \in \mathbb{R}_+$$

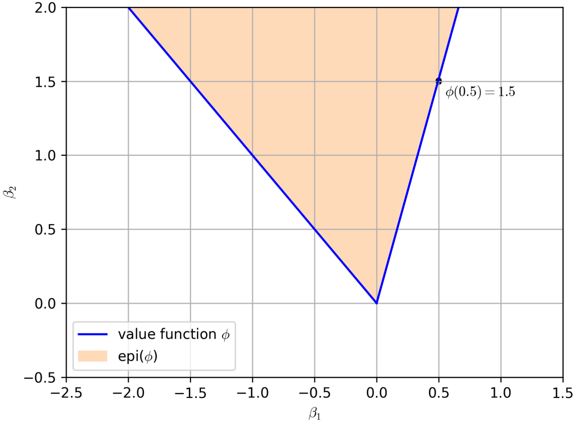

The Boundary of $\mathcal{C}$ and the Dual Problem

- The connection to the value function means that inequalities valid for $\mathcal{C}$ correspond to solutions to the LP dual.
- Hyperplanes that support $\mathcal{C}$ at a given point on the boundary are optimal to the dual.
- Facets of $\mathcal{C}$ are the basic optimal solutions (which are precisely the extreme rays of the dual cone $\mathcal{C}^*$).
- This is equivalent to complementary slackness conditions.

The Boundary of $\mathcal{C}$ and Pareto Optimality

- Conditions for Pareto optimality of solutions to a closely related multiobjective LP can also be derived by considering the boundary of a related set.
- For this, we interpret some of the rows of A as *multiple* objectives and define

$$\mathcal{P}_K^\beta = \{x \in \mathbb{R}_+^n, \beta_K \in \mathbb{R}^K \mid A_K x - I\beta_K = 0, A_{\bar{K}} x = b_{\bar{K}}\}$$

where $K \subseteq \{1, \ldots, m\}$ and $\bar{K} = \{1, \ldots, m\} \setminus K$.

- Project out the original variables to obtain $\mathcal{C}_K$.

$$\mathcal{C}_K = \operatorname{Proj}_{\beta_K}(\mathcal{P}_K^\beta)$$

- $\mathcal{C}_K$ is a slice of higher-dimensional cone from the pure feasibility case.
- It is shown in Fallah et al. [2023] that the efficient frontier for the multiobjective LP with objectives being the rows of A_K is (contained in) the boundary of $\mathcal{C}_K$.

Example 6

$$
\begin{aligned}
\text{vmin} \quad & 7x_2 + 10x_3 + 2x_4 + 10x_5 \\
& 10x_1 - 8x_2 + x_3 - 7x_4 + 6x_5 \\
\text{s.t.} \quad & 9x_1 + 3x_2 + 2x_3 + 6x_4 - 10x_5 = 4 \\
& x_2 + x_6 \leq 5 \\
& x_5 + x_7 \leq 5 \\
& x_j \in \mathbb{R}_+ \quad \forall j \in \{1, \dots, 7\},
\end{aligned}
$$

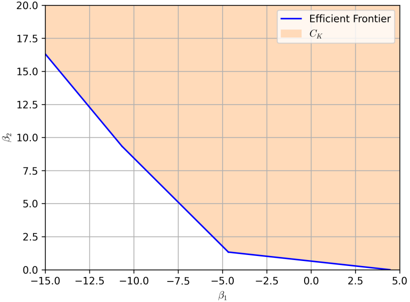

Generalizing to the MILP Case

- The very same logic extends easily to the MILP case.

$$
\begin{aligned}
\mathcal{S} &= \{x \in \mathbb{Z}^r_+ \times \mathbb{R}^{n-r}_+ \mid Ax = b\} \\
\mathcal{S}^\beta &= \{x \in \mathbb{Z}^r_+ \times \mathbb{R}^{n-r}_+, \beta \in \mathbb{R}^m \mid Ax - I\beta = 0\} \\
\mathcal{C} &= \mathrm{Proj}_\beta(\mathcal{S}^\beta) \\
\mathcal{S}^\beta_K &= \{x \in \mathbb{Z}^r_+ \times \mathbb{R}^{n-r}_+, \beta_K \in \mathbb{R}^K \mid A_K x - I\beta_K = 0, A_{\bar{K}} x = b_{\bar{K}}\} \\
\mathcal{C}_K &= \mathrm{Proj}_{\beta_K}(\mathcal{S}^\beta_K)
\end{aligned}
$$

A Generalized Farkas Lemma

$$\mathcal{S} = \emptyset \Leftrightarrow b \notin \mathcal{C}$$

- This is similar to other discrete Farkas lemmas [Bachem and Schrader, 1980, Blair and Jeroslow, 1982].
- Naturally, for this to be useful, we must replace the condition $b \notin \mathcal{C}$ with something that can be verified in practice.

Example 7

$$\mathcal{S}^{\beta} = \left\{ \begin{array}{r} 6x_1 + 5x_2 - 4x_3 + 2x_4 - 7x_5 + x_6 = \beta_1 \\ 3x_1 + \dfrac{7}{2}x_2 + 3x_3 + 6x_4 + 7x_5 + 5x_6 = \beta_2 \\ x_1, x_2, x_3 \in \mathbb{Z}_+,\ x_4, x_5, x_6 \in \mathbb{R}_+ \end{array} \right\}$$

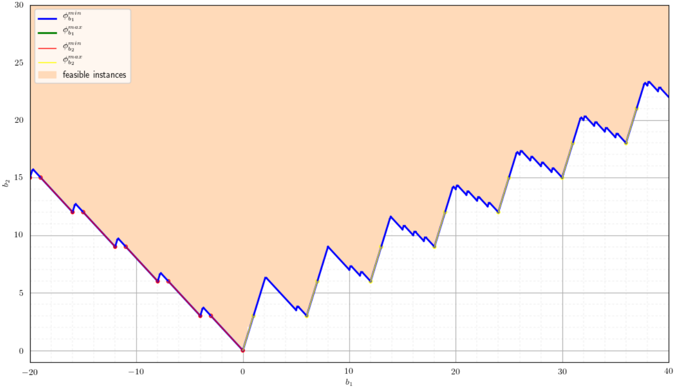

The Boundary of $\mathcal{C}$ in the MILP Case

- The set $\mathcal{C}$ and its boundary have the same interpretation and properties in the MILP case as in the LP case.
- However, since $\mathcal{C}$ may be a discrete set, the definition of "boundary" is not the usual set-theoretic one.
- As in the LP case, we say that β is on the boundary of $\mathcal{C}$ if either

$$\beta_i = \min_{x \in \mathcal{S}_{\{i\}}(\beta_{M \setminus \{i\}})} a^i x$$

or

$$\beta_i = \max_{x \in \mathcal{S}_{\{i\}}(\beta_{M \setminus \{i\}})} a^i x,$$

where

$$\mathcal{S}_{\{i\}}(\tau) = \{x \in \mathbb{Z}^r_+ \times \mathbb{R}^{n-r}_+ \mid A_{M \setminus \{i\}} x = \tau\}.$$

The Boundary of $\mathcal{C}$ and Optimality Conditions in MILP

The set $\mathcal{C}$ and its boundary have the same interpretation and properties in the MILP case as in the LP case.

- Proving $\mathcal{S} = \emptyset$ is equivalent to separating b from $\mathcal{C}$ (but not with a hyperplane!).
- For $i \in \{1, \ldots, m\}$ and $\tau \in \mathbb{R}^{m-1}$, we have that if

$$z^* = \min_{x \in \mathcal{S}_{\{i\}}(\tau)} / \max_{x \in \mathcal{S}_{\{i\}}(\tau)} a^i x. \qquad \text{(MILP)}$$

  then (z^*, τ) is on the boundary of $\mathcal{C}$ (converse holds under mild conditions).
- Thus, the boundary of $\mathcal{C}$ can be described by functions

$$\phi_i^-(\tau) = \min\{z \mid (z, \tau) \in \mathcal{C}\},$$
$$\phi_i^+(\tau) = \max\{z \mid (z, \tau) \in \mathcal{C}\},$$

  for $i \in \{1, \ldots, n\}$, which are the *value functions* of problems (MILP).
- Once again, for $K \subseteq \{1, \ldots, m\}$, the efficient frontier of the MILP with objectives being the rows of A_K is contained in the boundary of $\mathcal{C}_K$.

Example 8

$$\mathcal{S}^\beta = \left\{ \begin{array}{r} 6x_1 + 5x_2 - 4x_3 + 2x_4 - 7x_5 + x_6 = \beta_1, \\ 3x_1 + \frac{7}{2}x_2 + 3x_3 + 6x_4 + 7x_5 + 5x_6 = \beta_2, \\ x_1, x_2, x_3 \in \{0, 1\}, \\ 0 \leq x_4, x_5, x_6 \leq 3, \end{array} \right\}$$

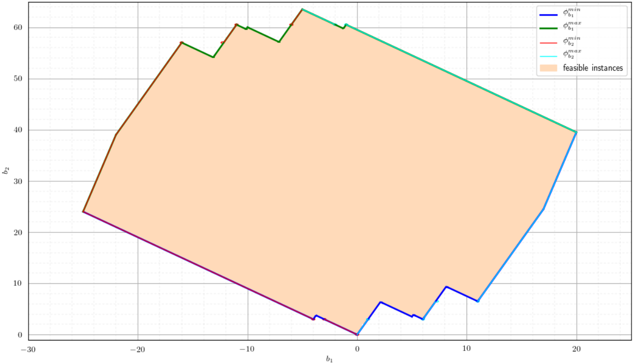

Separation for $\mathcal{C}$ and the Dual Problem

- Methods of constructing both the classical value function and the efficient frontier of a multiobjective MILP involve describing the boundary of $\mathcal{C}$.
- Algorithmically, this can be done by iteratively generating "separating functions," as in a cutting plane method.

Separating Functions

A *separating function* $F : \mathbb{R}^{m-1} \to \mathbb{R}$ for $\mathcal{C}$ is one that satisfies either

$$\begin{aligned} F(\tau) &\leq \phi_i^-(\tau) \quad \forall \tau \in \mathbb{R}^{m-1} \text{ or} \\ F(\tau) &\geq \phi_i^+(\tau) \quad \forall \tau \in \mathbb{R}^{m-1}. \end{aligned}$$

- Just as in the LP case, these separating functions are solutions to a dual problem and are called *dual functions* in that context.
- Finding a separating function for which $F(\tau) \approx \phi_i(\tau)$ for $\tau \in \mathbb{Q}^{m-1}$ is the *general dual problem* associated with (MILP) [Tind and Wolsey, 1981].

$$\max \{F(b) \mid F(\tau) \leq \phi_i^-(\tau),\ \tau \in \mathbb{R}^{m-1}, F \in \Upsilon^{m-1}\},$$

where $\Upsilon^m \subseteq \{f \mid f : \mathbb{R}^{m-1} \to \mathbb{R}\}$.

Discrete Farkas Lemma [Blair and Jeroslow, 1982]

Assuming $r = n$ (pure integer case), exactly one of the following holds:

1. $\mathcal{S} \neq \emptyset$
2. $\exists\, F : \mathbb{R}^m \rightarrow \mathbb{R}$ subadditive such that $F(A^j) \leq 0, j = 1, \ldots, n$ and $F(b) > 0$.

- Primal-dual pairs of MILPs have the same relationship as in the LP case.
- Since the value function of an MILP is subadditive, so there always exists a dual/separating function that is subadditive.
- When F is subadditive, the conditions for F to be a separating function reduce to the above.
- The result then says that $\mathcal{S}$ is empty if and only if we can separate b from $\mathcal{C}$ with a separating function.
- Alternatively, this is equivalent to F certifying that the dual problem is unbounded (with primal objective 0).

Outer Approximating $\mathcal{C}$

Using the machinery described so far, we can outer approximate $\mathcal{C}$ with separating functions.

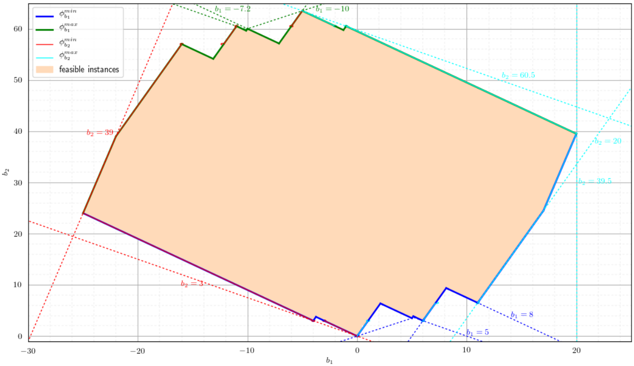

When $\mathcal{C}$ is bounded, we can describe it with a finite number of piecewise affine functions.

Multiobjective MILP in a Single Branch-and-Bound Tree

- For the remainder of the talk, we focus on an algorithm for generating the efficient frontier for a general multiobjective MILP.
- Surprisingly, this can be done within a single branch-and-bound by exploiting the ideas discussed so far.
- As earlier, let K be the index set of the rows of A that we interpret as multiple objectives.
- We arbitrarily choose one of these objectives as primary and treat the others as constraints.

Disjunctive Approximation of the Efficient Frontier

Let T be the set of terminating nodes of a branch-and-bound tree. The LP relaxation at node $t \in T$ is:

$$\begin{aligned} \phi^t(\tau) = \min\ & a^1 x \\ \text{s.t. } & A_{K\setminus\{1\}} x \leq \tau, \\ & A_{\bar{K}} x = b_{\bar{K}}, \\ & l^t \leq x \leq u^t, x \geq 0 \end{aligned} \tag{BB.VF}$$

By LP duality, we then have that:

$$\begin{aligned} \phi^t(\tau) = \max\ & v\tau + w b_{\bar{K}} + \underline{\pi} l^t + \bar{\pi} u^t \\ \text{s.t. } & v A_{K\setminus\{1\}} + w A_{\bar{K}} + \underline{\pi} + \bar{\pi} \leq a^1 \\ & \underline{\pi} \geq 0, v, \bar{\pi} \leq 0 \end{aligned} \tag{BB.LP.D}$$

Given a collection D of solutions feasible to (BB.LP.D), we obtain the following dual function, which approximates the value function and the efficient frontier from below.

$$F(\tau) = \min_{t \in T} \max_{(v,w,\underline{\pi},\overline{\pi}) \in D} u\tau + vb + \underline{\pi} l^t + \overline{\pi} u^t, \qquad \forall \zeta \in \mathcal{C}, \tag{1}$$

Example: Constructing the Separating/Dual Function

$$\phi(\beta) = \min 6x_1 + 4x_2 + 3x_3 + 4x_4 + 5x_5 + 7x_6$$
$$\text{s.t. } 2x_1 + 5x_2 - 2x_3 - 2x_4 + 5x_5 + 5x_6 = \beta$$
$$x_1, x_2, x_3 \in \mathbb{Z}_+, x_4, x_5, x_6 \in \mathbb{R}_+.$$

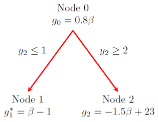

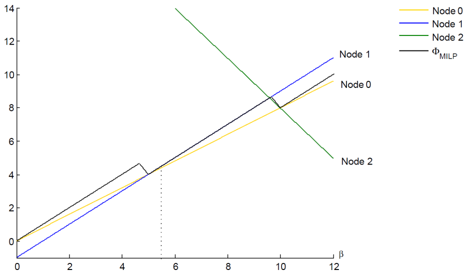

Example: Continuing with a Different Right-hand Side

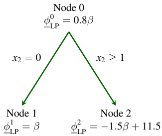

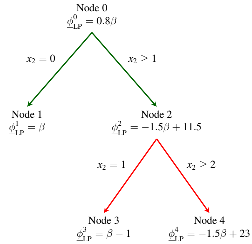

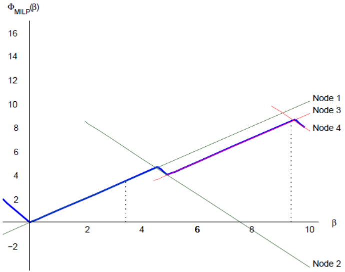

Convergence of the Algorithm

- We execute the branch-and-bound for a sequence of right-hand sides.
- Instead of re-starting each time, we continue in the same tree.
- We collect the dual solutions generated by solving the LP relaxations.
- There is a sequence of right-hand sides for which the algorithm converges finitely to the exact frontier.
- The key is finding the right set of right-hand sides.

Tree Representation of the Value Function

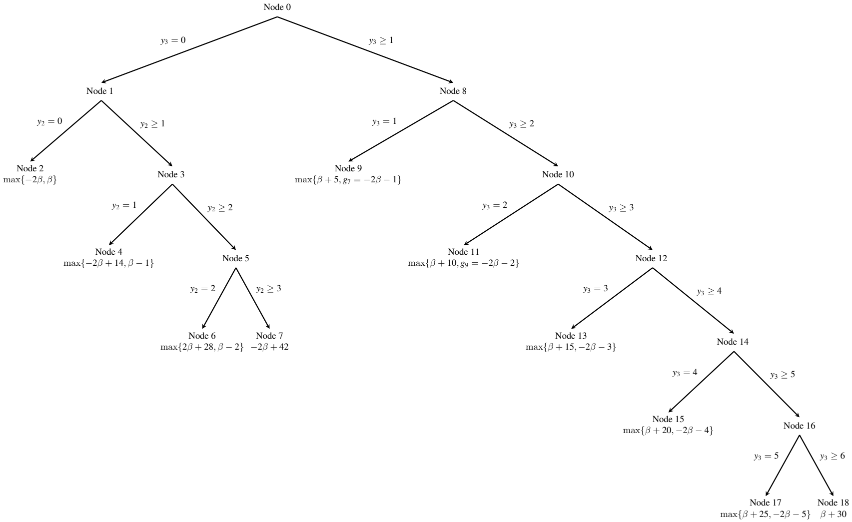

Correspondence of Nodes and Local Stability Regions

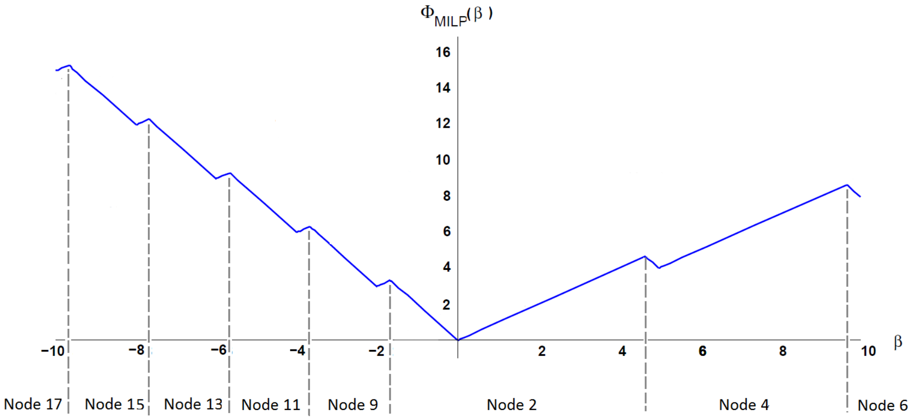

Another Example

$$
\begin{aligned}
\text{vmin} \quad & 2x_1 + 5x_2 + 7x_4 + 10x_5 + 2x_6 + 10x_7 \\
& -x_1 - 10x_2 + 10x_3 - 8x_4 + x_5 - 7x_6 + 6x_7 \\
\text{s.t.} \quad & -x_1 + 4x_2 + 9x_3 + 3x_4 + 2x_5 + 6x_6 - 10x_7 = 4 \\
& x_4 + 5x_2 \leq 5 \\
& x_7 + 5x_2 \leq 5 \\
& x_j \in \{0, 1\} \quad \forall j \in \{1, 2\}, \\
& x_j \in \mathbb{R}_+ \quad \forall j \in \{3, \ldots, 7\},
\end{aligned}
$$

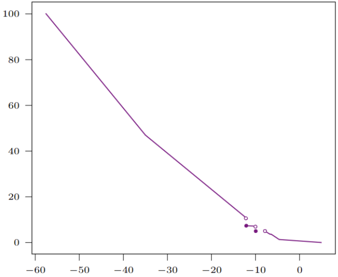

Evolution of Approximation

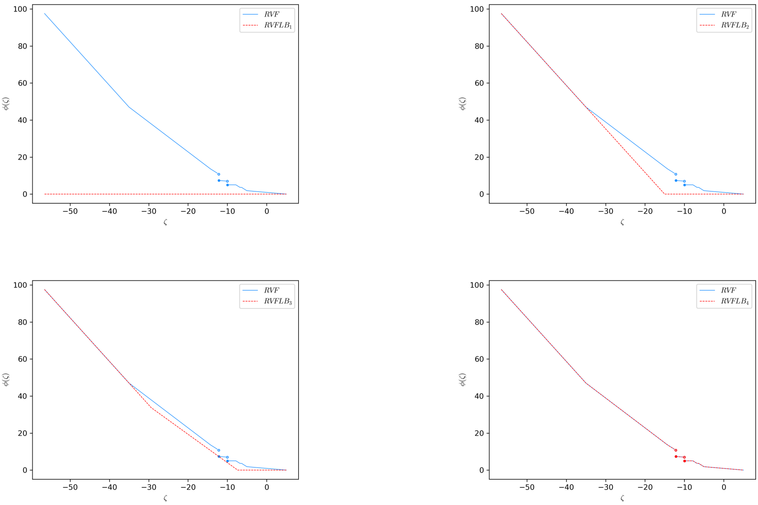

Evolution of Approximation

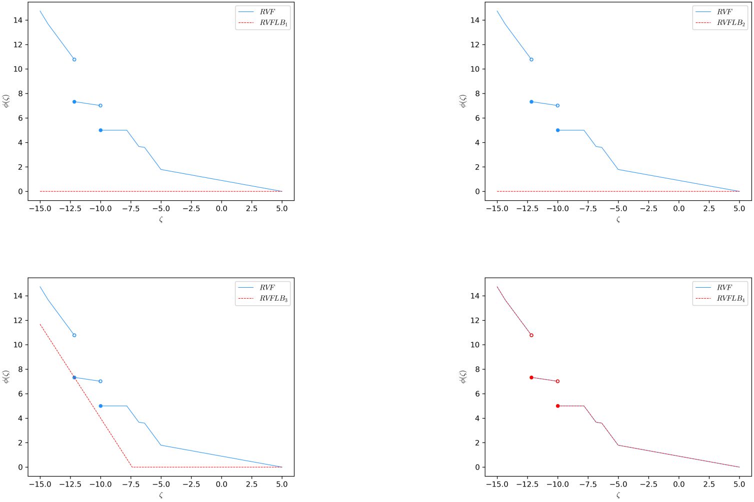

SYMPHONY

- SYMPHONY is an open source MILP solver framework with unique capabilities.
  - Can output formal proofs of optimality in the form of dual functions.
  - Can warm-start solution of a modified instance in the same tree.
  - Can be used to construct the value function or efficient frontier.
- The algorithm for constructing the efficient frontier was implemented in only a few dozen lines of code.
- SYMPHONY is also the subsolver for the bilevel solver MibS an can be used to warm-start the feasibility check, among other things.
- A generalized Benders algorithm for two-stage stochastic mixed integer linear optimization with recourse is also being revived.

Implications and Applications

Optimality conditions and value functions [Bolusani et al., 2020]

Yields optimality conditions for the follower's problem in bilevel optimization, which can be exploited to generate valid inequalities.

Construction of the efficient frontier [Fallah et al., 2023]

We derive a class of algorithms that generates the efficient frontier of a multiobjective mixed integer optimization problem in a single branch-and-bound tree.

Generalized Benders [Hassanzadeh and Ralphs, 2014]

Benders for two-stage stochastic optimization and bilevel optimization.

Lagrangian relaxation and Dantzig-Wolfe decomposition [Bodur et al., 2016]

Alternative methods for computing bounds in decomposition methods.

Warm-starting solution of MILPs [Ralphs and Güzelsoy, 2005]

Improved efficiency when solving sequences of related MILPs.

References I

A. Bachem and R. Schrader. Minimal equalities and subadditive duality. *Siam J. on Control and Optimization*, 18(4):437–443, 1980.

C.E. Blair and R.G. Jeroslow. The value function of an integer program. *Mathematical Programming*, 23:237–273, 1982.

Merve Bodur, Shabbir Ahmed, Natashia Boland, and George L Nemhauser. Decomposition of loosely coupled integer programs: A multiobjective perspective. *Optimization Online*, 2016.

S. Bolusani, S. Coniglio, T.K. Ralphs, and S. Tahernejad. A Unified Framework for Multistage Mixed Integer Linear Optimization. In S. Dempe and A. Zemkoho, editors, *Bilevel Optimization: Advances and Next Challenges*, chapter 18, pages 513–560. Springer, 2020. ISBN 978-3-030-52118-9. doi: 10.1007/978-3-030-52119-6. URL `https://arxiv.org/abs/2104.09003`.

References II

Samira Fallah, Ted K. Ralphs, and Natashia L Boland. On the Relationship Between the Value Function and the Efficient Frontier of a Mixed Integer Linear Optimization Problem. Technical report, COR@L Laboratory, Lehigh University, 2023. URL `https://arxiv.org/abs/2303.00785`.

A. Hassanzadeh and T.K. Ralphs. A Generalization of Benders' Algorithm for Two-Stage Stochastic Optimization Problems with Mixed Integer Recourse. Technical Report 14T-005, COR@L Laboratory, Lehigh University, 2014. URL `http://coral.ie.lehigh.edu/~ted/files/papers/SMILPGenBenders14.pdf`.

T.K. Ralphs and M. Güzelsoy. The SYMPHONY Callable Library for Mixed Integer Programming. In *Proceedings of the Ninth INFORMS Computing Society Conference*, pages 61–76, 2005. doi: 10.1007/0-387-23529-9_5. URL `http://coral.ie.lehigh.edu/~ted/files/papers/SYMPHONY04.pdf`.

J. Tind and L. Wolsey. An elementary survey of general duality theory in mathematical programming. *Mathematical Programming*, 21:241–261, 1981.