LP RELAXATIONS AND INTEGRALITY GAP

CSCI 4113/6101
INSTRUCTOR: NORBERT ZEH
SEPTEMBER 9, 2025

Given that we have polynomial-time algorithms for solving LPs but solving ILPs is NP-hard, an important question is whether we can use the machinery for solving LPs to obtain suboptimal but good solutions for ILPs. We will explore this question in the third part of this book. Here, we introduce an important tool for translating ILPs into related LPs: LP relaxation.

For every ILP, one can define a corresponding LP by simply dropping the requirement that a feasible solution must assign integer values to all variables—in the LP solution, variables are allowed to take on any real values as long as they satisfy the constraints in the LP. This LP is called the **LP relaxation** of the ILP. As an example, here is the LP relaxation of the first ILP formulation of the MST problem we developed earlier. It is identical to the ILP, except that the constraint $x_e \in \{0,1\}$ has been replaced with the constraint $0 \le x_e \le 1$—x_e is no longer required to be integral:

$$\begin{aligned}
&\text{Minimize } \sum_{e \in E} w_e x_e\\
\text{s.t. } &\sum_{e \in E} x_e = n - 1\\
&\sum_{e \in C} x_e \le |C| - 1 \quad \forall \text{cycle } C \text{ in } G\\
&0 \le x_e \le 1 \quad \forall e \in E.
\end{aligned}$$

In general, the LP relaxation of an ILP may have a much better optimal solution than the ILP itself, but it can never have a worse optimal solution:

LEMMA 1. *If $\hat{x}$ is an optimal solution to a minimization ILP P and $\tilde{x}$ is an optimal solution to its LP relaxation P', then $f(\tilde{x}) \le f(\hat{x})$, where f is the objective function of P and P'.*

Proof. Since $\hat{x}$ is a feasible solution of P, it is also a feasible solution of P'. Since $\tilde{x}$ is an optimal solution of P', this implies that $f(\tilde{x}) \le f(\hat{x})$. □

The ratio $\frac{f(\hat{x})}{f(\tilde{x})}$ between the optimal objective function values of the ILP and its relaxation is called the **integrality gap** of the ILP. In general, this gap can be arbitrarily large. However, for problems where the integrality gap is bounded, the LP relaxation can often be used to obtain an approximation of the optimal solution of the ILP efficiently, a fact that is at the heart of LP rounding, technique for developing polynomial-time approximation algorithms discussed later in this course. Another useful property exploited by some approximation algorithms is **half-integrality** of the optimal solution of some LPs: in the optimal solution, every variable takes on a value that is a multiple of $\frac{1}{2}$. We will discuss this later, as port of the discussion of LP rounding.

1 Integrality Gaps of Our ILP Formulations of MST

In some instances, the LP relaxation of an ILP has an integral optimal solution, that is, every variable in this optimal solution has an integer value even though this is not explicitly enforced by the LP. By Lem. 1, this implies that the optimal solution of the LP relaxation is also an optimal solution of the ILP. Is the MST problem such a problem? The answer depends on the chosen ILP formulation. For each of the ILP formulations of the MST problem that we have discussed so far, there exists an input that has a fractional solution with a lower objective function value than the best possible integral solution, as shown in the following two examples.

Example 2. The first ILP formulation of the MST problem we developed was based on the characterization that a tree on n vertices has $n-1$ edges and is acyclic:

$$\begin{aligned} \text{Minimize } & \sum_{e \in E} w_e x_e \\ \text{s.t. } & \sum_{e \in E} x_e = n - 1 \\ & \sum_{e \in C} x_e \le |C| - 1 \quad \forall \text{cycle } C \text{ in } G \\ & x_e \in \{0,1\} \quad \forall e \in E. \end{aligned} \tag{1}$$

As observed above, its LP relaxation is

$$\begin{aligned} \text{Minimize } & \sum_{e \in E} w_e x_e \\ \text{s.t. } & \sum_{e \in E} x_e = n - 1 \\ & \sum_{e \in C} x_e \le |C| - 1 \quad \forall \text{cycle } C \text{ in } G \\ & 0 \le x_e \le 1 \quad \forall e \in E. \end{aligned} \tag{2}$$

Any MST of the graph in Fig. 1 has weight 10. One such MST is shown in red. On the other hand, the

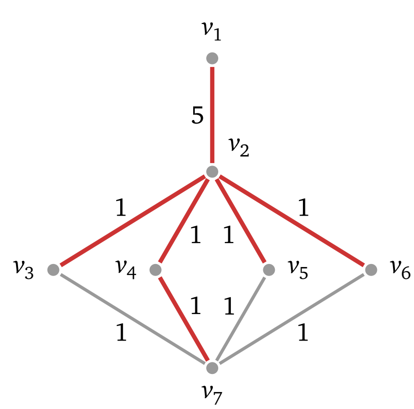

Figure 1: The graph used in Ex. 2. The edge labels are edge weights.

following fractional solution with objective function value 6 satisfies all constraints of (2): [1]

$$\begin{aligned} x_{12} &= 0 \\ x_{23} = x_{24} = x_{25} = x_{26} &= \frac{1}{2} \\ x_{37} = x_{47} = x_{57} = x_{67} &= 1 \end{aligned}$$

Thus, the integrality gap of (1) is at least 5/3.

Example 3. The second ILP formulation of the MST problem we developed was based on the characterization that a tree on n vertices has $n-1$ edges and must be connected, which we expressed by requiring that every cut in the tree must be crossed by at least one edge:

$$\begin{aligned} &\text{Minimize } \sum_{e \in E} w_e x_e \\ \text{s.t. } &\sum_{e \in E} x_e = n-1 \\ &\sum_{e \text{ crosses } S} x_e \geq 1 \qquad \forall \text{cut } S \text{ in } G \\ &x_e \in \{0,1\} \quad \forall e \in E. \end{aligned} \tag{3}$$

Its LP relaxation is

$$\begin{aligned} &\text{Minimize } \sum_{e \in E} w_e x_e \\ \text{s.t. } &\sum_{e \in E} x_e = n-1 \\ &\sum_{e \text{ crosses } S} x_e \geq 1 \qquad \forall \text{cut } S \text{ in } G \\ &0 \leq x_e \leq 1 \quad \forall e \in E. \end{aligned} \tag{4}$$

Any MST of the graph in Fig. 2 has weight 2. One such MST is shown in red. On the other hand, the following fractional solution with objective function value $\frac{3}{2}$ satisfies all constraints of (4):

$$\begin{aligned} x_{12} = x_{13} = x_{24} = x_{34} &= \frac{1}{2} \\ x_{25} = x_{45} &= 1 \end{aligned}$$

Thus, the integrality gap of (3) is at least 4/3.

2 An ILP Formulation of MST with Integrality Gap 1

As the two examples in the previous section show, integrality matters for all the ILP formulations of the MST problem discussed so far. The following extension of (1) finally is an ILP whose LP relaxation has

[1] For the sake of simplicity, we write x_{ij} instead of x_{v_i,v_j} in this example and the next two.

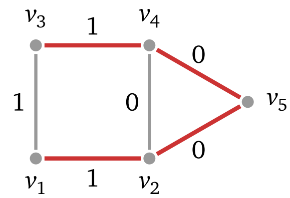

Figure 2: The graph used in Ex. 3. The edge labels are edge weights.

an optimal integral solution, that is, this ILP formulation has integrality gap 1:

$$
\begin{aligned}
&\text{Minimize } \sum_{e \in E} w_e x_e \\
&\text{s.t.} \quad \sum_{e \in E} x_e \geq n - 1 \\
&\qquad \sum_{\substack{(u,v) \in E \\ u,v \in S}} x_{u,v} \leq |S| - 1 \quad \forall \emptyset \subset S \subseteq V \\
&\qquad x_e \in \{0,1\} \qquad \forall e \in E.
\end{aligned}
\tag{5}
$$

This is an extension of (1) because it does not only require that at most $k-1$ edges of every k-vertex cycle in G are chosen as part of the spanning tree but that at most $k-1$ edges between *any* subset of k vertices are chosen. Given that not including any cycle in T is what makes T a tree, it is surprising that this seemingly unimportant strengthening of the LP ensures that the LP relaxation has an integral optimal solution. Let's first establish the correctness of the *ILP* (5).

Lemma 4. *For any undirected edge-weighted graph (G, w), $\hat{x}$ is an optimal solution of the ILP* (5) *if and only if the subgraph $T = (V, E')$ of G with $E' = \{e \in E \mid \hat{x}_e = 1\}$ is an MST of (G, w).*

Proof. By Lem. 2 in our discussion of the MST problem, we only need to prove that $w(E') = \sum_{e \in E} w_e \hat{x}_e$ and that $\hat{x}$ is a feasible solution of (5) if and only if $T = (V, E')$ is a spanning tree of G. The former is obvious. So let us prove the latter.

First note that, if $\hat{x}$ is a feasible solution of (5), then E' contains at least $n-1$ edges, and it contains at most $|C|-1$ edges from every cycle C in G because it contains at most $|C|-1$ of the edges between the vertices in C whether these edges are part of the cycle or not. For the set $S = V$, (5) ensures that E' contains at most $|V|-1 = n-1$ of the edges connecting vertices in V. But that's all edges of G, that is, E' contains at most $n-1$ edges. Since E' contains at least $n-1$ edges, it therefore contains exactly $n-1$ edges. Thus, $T = (V, E')$ is an acyclic graph with $n-1$ edges, that is, it is a spanning tree of G.

Conversely, if $T = (V, E')$ is a spanning tree of T', then it contains exactly $n-1$ edges, that is, $\hat{x}$ satisfies the first constraint of (5). To prove that $\hat{x}$ also satisfies the other constraints of (5), consider any non-empty subset $S \subseteq V$. Since T is acyclic, so it $T[S] = (S, \{(u,v) \in E' \mid u, v \in S\})$ (we call $T[S]$ the **subgraph of T induced by S**). Therefore, $T[S]$ is a forest. Each tree in this forest has one less edge than vertices. Thus, $T[S]$ has at least one less edge than vertices. But the vertex set of $T[S]$ is S, and the edges of $T[S]$ are exactly the edges in T connecting vertices in S. Thus, $\hat{x}$ satisfies the constraint of (5) corresponding to S. Since this is true for every non-empty subset $S \subseteq V$, $\hat{x}$ is a feasible solution of (5).

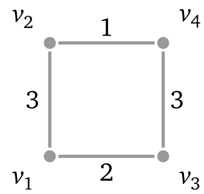

Figure 3: A graph whose relaxation of (5) has a fractional optimal solution.

By Lem. 2 in the discussion of the MST problem, this implies that the set E' corresponding to an optimal solution $\hat{x}$ of (5) defines an MST $T = (V, E')$ of G. □

As the following example shows, if there are edges of equal weight in (G, w), then the LP relaxation of (5) may once again have an optimal fractional solution.

EXAMPLE 5. Consider the graph in Fig. 3. The corresponding LP is

$$\begin{aligned} \text{Minimize } & 3x_{12} + 2x_{13} + x_{24} + 3x_{34} \\ \text{s.t. } & x_{12} + x_{13} + x_{24} + x_{34} \geq 3 \\ & x_{12} + x_{13} + x_{24} + x_{34} \leq 3 \\ & x_{12} + x_{13} \leq 2 \\ & x_{12} + x_{24} \leq 2 \\ & x_{13} + x_{34} \leq 2 \\ & x_{24} + x_{34} \leq 2 \\ & 0 \leq x_e \leq 1 \quad \forall e \in E \end{aligned}$$

which has the optimal fractional solution

$$x_{12} = \alpha \qquad x_{34} = 1 - \alpha \qquad x_{13} = x_{24} = 1,$$

for any $0 \leq \alpha \leq 1$.

However, it is easily verified that, at least in this example, any optimal fractional solution has the same objective function value as an optimal integral solution, that is, the optimal solution of the ILP is also an optimal solution of its relaxation. To prove that this is always the case, we will need the concept of complementary slackness, which we will discuss in an upcoming lecture. We will revisit this ILP formulation then and prove that its integrality gap is indeed 1.

This ILP formulation has an even stronger property, which we will *not* prove: If we use the Simplex Algorithm to find a solution of the LP relaxation of (5), the computed solution is always integral. The proof uses two facts that we do not prove here: (1) Recall that the feasible region of an LP is a convex polytope in $\mathbb{R}^n$. A geometric view of the Simplex Algorithm is that it moves between vertices of this polytope until it reaches a vertex that represents an optimal solution. (2) Every vertex of the polytope defined by the LP relaxation of (5) represents an integral solution.