21 Solving Linear Programs

In this chapter we will discuss the simplex method, the simplest and most commonly used algorithm for solving arbitrary linear programs.

21.1 Equational Form

In Section 20.1 we defined standard form for LPs where the objective is maximization, all variables are constrained to be non-negative, and all other constraints are $\leq$-constraints. We also saw how to convert an arbitrary LP to standard form. The simplex method, however, operates on LPs in a slightly different form known as **equational form**. We say that an LP is in equational form if:

- It has n constraints of the form $x_i \geq 0$
- The remaining m constraints are each of the form $a_{j1}x_1 + a_{j2}x_2 + \cdots + a_{jn}x_n = b_j$, where $a_{j1}, \ldots, a_{jn}$, and b_j are constants (i.e. the remaining constraints are all equations with a constant on the right hand side).
- The objective is maximization (rather than minimization).

Note that the only difference between equational form and standard form is that our constraints (other than the non-negativity constraints) are equations rather than $\leq$-inequalities.

We have previously seen how to convert equations into inequalities to get an LP into standard form. We can also convert inequalities into equations to convert a standard form LP into equational form. The key idea is to introduce additional variables which we call *slack variables* which will show up in only one constraint each, and store the difference between the left and right hand sides of the original constraint. We introduce a new variable s_j for each of the $m \leq$ constraints in the standard form LP. We replace each constraint $f_j(x) \leq b_j$ with the new constraints $f_j(x) + s_j = b_j$ and $s_j \geq 0$.

Since we've converted all the inequalities into equations and added a non-negativity constraint for each new variable, our new LP is in equational form. Every feasible solution x in our original standard form LP corresponds to a solution to the equational form LP where we set each x_i to the same value as in x and set $s_j = b_j - \sum_{i=1}^{n} a_{ji}x_i$ for each $1 \leq j \leq m$. Since x satisfied the $\leq$ constraints, each s_j will be non-negative. On the other hand, if we have a solution x, s to the equational form LP, then we can ignore s and x is a solution to the original standard form LP. Since each slack variable is non-negative, x must satisfy the $\leq$ constraints.

For example, if we take the standard form LP from the end of Section 20.1, reproduced below,

$$\begin{array}{lrcl}
\text{maximize} & -y_1 + z_1 - y_2 + z_2 & & \\
\text{subject to} & y_1 - z_1 & \leq & 5 \\
& -2y_1 + 2z_1 - y_2 + z_2 & \leq & -10 \\
& -y_1 + z_1 - y_2 + z_2 & \leq & 3 \\
& y_1 - z_1 + y_2 - z_2 & \leq & -3 \\
& y_1 \ , \ z_1 \ , \ y_2 \ , \ z_2 & \geq & 0
\end{array}$$

we can convert it to the following equivalent LP in equational form:

$$\begin{array}{lrcl}
\text{maximize} & -y_1 + z_1 - y_2 + z_2 & & \\
\text{subject to} & y_1 - z_1 + s_1 & = & 5 \\
& -2y_1 + 2z_1 - y_2 + z_2 + s_2 & = & -10 \\
& -y_1 + z_1 - y_2 + z_2 + s_3 & = & 3 \\
& y_1 - z_1 + y_2 - z_2 + s_4 & = & -3 \\
& y_1 \ , \ z_1 \ , \ y_2 \ , \ z_2 \ , \ s_1 \ , \ s_2 \ , \ s_3 \ , \ s_4 & \geq & 0
\end{array}$$

A **linear combination** of a set of equations $L_1 = R_1, L_2 = R_2, \ldots, L_k = R_k$ is an equation in the form $\sum_{i=1}^{k} c_i L_i = \sum_{i=1}^{k} c_i R_i$ where $c_1, \ldots, c_k$ are constants. That is, a linear combination of equations is formed by multiplying each equation by a constant and then adding them together.

For example, the equation $8y = 16/3$ is a linear combination of the equations $x + 2y = 3$ and $6x - 12y = 2$, because we can obtain it by adding 2 times $x + 2y = 3$ to $-1/3$ of $6x - 12y = 2$ ($2(x + 2y) + (-1/3)(6x - 12y) = 8y$ and $2(3) + (-1/3)(2) = 16/3$).

We say that a constraint is **redundant** if it is a linear combination of other constraints in the LP. We will also assume for the rest of this chapter that our LPs in equational form do not contain any redundant constraints. This is a natural restriction, because if we have redundant constraints, we can delete one of them without changing the set of feasible solutions for the LP (any assignment to the variables which satisfies the other constraints will also satisfy a linear combination of them).

21.2 The Geometry of Linear Programs

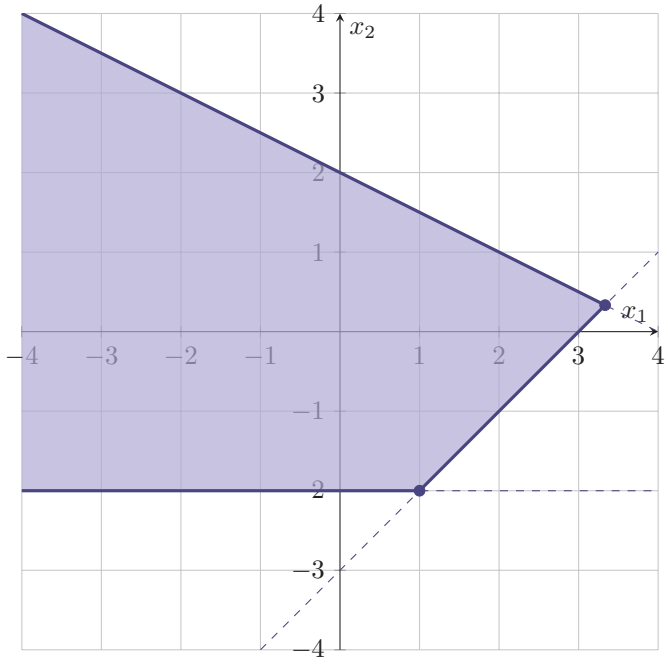
In this section we will discuss a geometric perspective on linear programming in order to develop an intuition for what the simplex is doing and why it works. If P is a linear program with variables $x_1, \ldots, x_n$, we can view a solution $x_1 = y_1, x_2 = y_2, \ldots, x_n = y_n$ to P as representing a point $y = (y_1, y_2, \ldots, y_n)$ in n-dimensional space. We define the **feasible region** of a linear program P as the set of points in n-dimensional space which are feasible solutions to P (i.e. the set of points whose coordinates satisfy all of the inequalities).

For example, Figure 23 depicts the feasible region of the LP below.

Figure 22: A linear program

$$\begin{array}{lrcl} \text{maximize} & x_1 - 2x_2 & & \\ \text{subject to} & x_1 + 2x_2 & \leq & 4 \\ & x_1 - x_2 & \leq & 3 \\ & x_2 & \geq & -2 \end{array}$$

Figure 23: The feasible region for the LP in Figure 22



For an LP with two variables, each constraint divides the Cartesian plane in half along a line. The shaded region is the part of the plane which is on the correct side of all three lines. Note that the

feasible region extends infinitely off to the left of the picture. The feasible region of an LP depends only on the constraints, not the objective function. We can view the objective function as giving us a direction that we want to travel, and an optimal solution is a point that is as far as possible in the objective direction while remaining in the feasible region.

The feasible region of a linear program has some nice properties. In particular, it is *convex*, which means that for any two points in the feasible region, any point on the line between them is also in the feasible region. We state this formally in the following theorem (the proof is not complicated, but we omit it to avoid getting bogged down in linear algebra).

Theorem 93. *If $y = (y_1, \ldots, y_n)$ and $z = (z_1, \ldots, z_n)$ are feasible solutions to a linear program P, then for every $0 \leq \alpha \leq 1$, $\alpha y + (1 - \alpha)z$ (i.e. the solution $x_i = \alpha y_i + (1 - \alpha)z_i$) is also a feasible solution to P.*

We say that a solution x is a **vertex** of P (not to be confused with a vertex of a graph) if x lies on a corner of the feasible region, i.e. if x does not lie on a line between any two distinct points in the feasible region for P. For example, the vertices for the LP in Figure 22 are $(10/3, 1/3)$ and $(1, -2)$. Vertices of the LP in Figure 22 correspond to points where two of the lines in Figure 23 intersect, i.e. points where the left hand side is exactly equal to the right hand side for two of the three constraints in the LP.

The unique optimal solution to the LP in Figure 22 is the vertex $(1, -2)$. This is not a coincidence, it turns out that when an LP has an optimal solution, it *always* has an optimal solution at a vertex.

Theorem 94. *If P is an LP with an optimal solution (i.e. if P is feasible and bounded) and P has a vertex, then P has an optimal solution which is a vertex.*

We will not prove this theorem either, but an intuition is that if we start anywhere in the feasible region and move in the direction of the objective function until we get stuck, then we will either be at a corner of the feasible region or we will hit a boundary of the feasible region which is directly perpendicular to our objective direction, in which case any point on that boundary is an optimal solution, including one of its corners (e.g. if we change the objective of the LP in Figure 22 to "maximize $x_1 + 2x_2$", then there are infinitely many optimal solutions along the line $x_1 + 2x_2 = 4$, but one of them is the vertex $(10/3, 1/3)$).

Theorem 94 is the inspiration for the simplex method. We give high level pseudocode for the method below (this version assumes that the input LP is feasible and bounded, although as we'll see later the simplex method can be used to verify both of those conditions).

```
SIMPLEX(P)
  Find a vertex y of the feasible region for P
  While there exists a vertex of the feasible region adjacent to y with a larger objective value:
      y ← an adjacent vertex with a higher objective value
  Return y
```

That is, the idea of the simplex method is to simply walk around the border of the feasible region, moving from corner to corner in a way that always makes progress in the direction of the objective function. This is essentially a greedy algorithm, but Theorem 93 guarantees that it can never get stuck in a local optimum. The line between our current point y and an optimal solution for the LP always lies within the feasible region, so we can always make progress in the direction of the objective function. Since every iteration of the while loop improves the objective value, we will never repeat vertices, so the simplex method is guaranteed to terminate (although it may take exponential time), because an LP can have only finitely many vertices (each one corresponds to a subset of the constraints).

In Section 21.3 we will explain how we can find an adjacent vertex with higher objective value (we call the process of moving from one vertex to an adjacent vertex "pivoting"), and in Section 21.4 we will explain how to implement the first line and find a vertex in the first place (or determine that the LP is infeasible).

21.3 Pivoting

Vertices of LPs in equational form have a particularly nice structure. Let P be an LP in equational form with n variables and m equations. Recall that vertices occur where boundaries of the feasible region intersect. Since P is in equational form, it is bounded by the non-negativity constraints, and each of its vertices will be a point where at most m of the variables are non-zero.

To find a vertex for P, we can choose a set of m variables which we will not force to be 0 (we call these variables a **basis**), and solve our system of equations to obtain values for these variables assuming that the other variables are 0. If the values that we obtain are all non-negative, then the solution that we've found is feasible, and we call it a **basic feasible solution**. By restricting our input to LPs in equational form, the simplex method consists of pivoting between basic feasible solutions. Pivoting to an adjacent vertex involves bringing a new variable into the basis and kicking an old variable out (so that our basis always has size m, where m is the number of equations in the LP).

So how do we find the solution associated with a given basis? The idea is to take linear combinations of our equations in order to obtain equations which involve only 1 variable from our basis (you may have learned this as the "elimination method" in middle or high school algebra). For example, suppose that we have the following LP and want to find a solution with basis x_1, x_2, x_3.

$$\begin{array}{lrcl}
\text{maximize} & x_1 + 2x_2 + 3x_3 + 4x_4 + 5x_5 & & \\
\text{subject to} & x_1 + 2x_3 + 3x_4 + x_5 & = & 7 \\
& 2x_1 + x_2 + x_4 + x_5 & = & 4 \\
& x_1 + 4x_2 + 3x_3 - 3x_5 & = & 18 \\
& x_1, x_2, x_3, x_4, x_5 & \geq & 0
\end{array}$$

We start by adding or subtracting multiples of the equations from each other to eliminate x_1 from all but 1 of the equations. If we subtract the first equation from the third, we eliminate x_1 from the third equation, obtaining the equivalent LP:

$$\begin{array}{lrcl}
\text{maximize} & x_1 + 2x_2 + 3x_3 + 4x_4 + 5x_5 & & \\
\text{subject to} & x_1 + 2x_3 + 3x_4 + x_5 & = & 7 \\
& 2x_1 + x_2 + x_4 + x_5 & = & 4 \\
& 4x_2 + x_3 - 3x_4 - 4x_5 & = & 11 \\
& x_1, x_2, x_3, x_4, x_5 & \geq & 0
\end{array}$$

Then if we subtract twice the first equation from the second equation, we get an equivalent LP where only the first equation involves x_1:

$$\begin{array}{lrcl}
\text{maximize} & x_1 + 2x_2 + 3x_3 + 4x_4 + 5x_5 & & \\
\text{subject to} & x_1 + 2x_3 + 3x_4 + x_5 & = & 7 \\
& x_2 - 4x_3 - 5x_4 - x_5 & = & -10 \\
& 4x_2 + x_3 - 3x_4 - 4x_5 & = & 11 \\
& x_1, x_2, x_3, x_4, x_5 & \geq & 0
\end{array}$$

Next we eliminate x_2 from all equations but the second by subtracting four times the second equation from the third equation to obtain:

$$\begin{array}{lrcl}
\text{maximize} & x_1 + 2x_2 + 3x_3 + 4x_4 + 5x_5 & & \\
\text{subject to} & x_1 + 2x_3 + 3x_4 + x_5 & = & 7 \\
& x_2 - 4x_3 - 5x_4 - x_5 & = & -10 \\
& 17x_3 + 17x_4 & = & 51 \\
& x_1, x_2, x_3, x_4, x_5 & \geq & 0
\end{array}$$

We divide the third equation by 17 so that it has only a single copy of x_3:

$$\begin{array}{lrcrcrcrcl}
\text{maximize} & x_1 & + & 2x_2 & + & 3x_3 & + & 4x_4 & + & 5x_5 \\
\text{subject to} & x_1 & & & + & 2x_3 & + & 3x_4 & + & x_5 = 7 \\
& & & x_2 & - & 4x_3 & - & 5x_4 & - & x_5 = -10 \\
& & & & & x_3 & + & x_4 & & = 3 \\
& x_1 & , & x_2 & , & x_3 & , & x_4 & , & x_5 \geq 0
\end{array}$$

Finally, we eliminate x_3 from the first two equations, by subtracting twice the third equation from the first equation and adding four times the third equation to the second equation. This at last gives us an LP equivalent to our original LP where each variable of our basis appears in an equation without the other two:

$$\begin{array}{lrcrcrcrcl}
\text{maximize} & x_1 & + & 2x_2 & + & 3x_3 & + & 4x_4 & + & 5x_5 \\
\text{subject to} & x_1 & & & & & + & x_4 & + & x_5 = 1 \\
& & & x_2 & & & - & x_4 & - & x_5 = 2 \\
& & & & & x_3 & + & x_4 & & = 3 \\
& x_1 & , & x_2 & , & x_3 & , & x_4 & , & x_5 \geq 0
\end{array}$$

Setting x_4 and x_5 to 0, we can read values for x_1, x_2, x_3 off the right hand side and we find that $x = (1, 2, 3, 0, 0)$ is a basic feasible solution for the original linear program, with an objective value of $1(1) + 2(2) + 3(3) = 14$.

To move to an adjacent vertex, we need to replace one of x_1, x_2, x_3 in our basis with either x_4 or x_5. We could try repeating the above process from scratch for all 3×2 combinations of a variable leaving the basis and variable entering the basis, but it's possible to do things a bit faster.

Let's suppose that we decide that we want to bring x_4 into the basis. Then, we are going to increase x_4 (which is 0 in our current solution) while keeping $x_5 = 0$. In order to figure out which variable leaves, we can express x_1, x_2, x_3 in terms of x_4 using the last version of the LP that we wrote above. We get that

$$\begin{aligned}
x_1 &= 1 - x_4 \\
x_2 &= 2 + x_4 \\
x_3 &= 3 - x_4
\end{aligned}$$

Thus, if we increase x_4 while keeping the equations true, x_1 and x_3 will decrease, while x_2 will increase. We reach a new basic feasible solution when one of our current basis variables hits 0 and drops out of the basis. Which one hits 0 first? Solving the equations above, we find that x_1 would hit 0 at $x_4 = 1$, while x_3 would hit 0 at $x_4 = 3$ (x_2 will never hit 0 as we increase x_4). Thus, x_1 hits 0 first, and so in our new basic feasible solution, x_4 replaces x_1. We have $x_4 = 1$, $x_2 = 2 + 1 = 3$, and $x_3 = 3 - 1 = 2$. Thus, if we bring x_4 into the basis, our new basic feasible solution would be $(0, 3, 2, 1, 0)$ with an objective value of $2(3) + 3(2) + 1(4) = 16$. This is larger than the objective value of our previous basic feasible solution $(1, 2, 3, 0, 0)$, so we pivot to $(0, 3, 2, 1, 0)$ (if bringing in x_4 had decreased the objective value we would have to try bringing in x_5 instead).

From our previous version of the LP where x_1, x_2, x_3 are by themselves, we can easily obtain a new version where x_2, x_3, and x_4 are by themselves instead. We add the first equation to the second and subtract the first equation from the third, so that our first equation now isolates x_4 instead of x_1:

$$\begin{array}{lrcrcrcrcl}
\text{maximize} & x_1 & + & 2x_2 & + & 3x_3 & + & 4x_4 & + & 5x_5 \\
\text{subject to} & x_1 & & & & & + & x_4 & + & x_5 = 1 \\
& x_1 & + & x_2 & & & & & & = 3 \\
& -x_1 & & & + & x_3 & & & - & x_5 = 2 \\
& x_1 & , & x_2 & , & x_3 & , & x_4 & , & x_5 \geq 0
\end{array}$$

To continue the simplex method we need to find another adjacent vertex (basic feasible solution to pivot to). Since we just removed x_1 from our basis, adding it back would return us to the same

vertex that we already tried with the basis x_1, x_2, x_3. Therefore, the only neighboring vertex we need to consider is the one we would reach by introducing x_5 to the basis.

Like we did before, we represent our current variables in terms of x_5 to see how they change when x_5 enters the basis:

$$\begin{aligned} x_2 &= 3 \\ x_3 &= 2 + x_5 \\ x_4 &= 1 - x_5 \end{aligned}$$

This time, x_4 is the only variable that decreases as x_5 increases, so x_4 will be the variable to leave the basis, hitting 0 at $x_5 = 1$. Our new basic feasible solution is $(0, 3, 3, 0, 1)$ with an objective value of $2(3) + 3(3) + 5(1) = 20$. Once again, our objective value has increased, so we pivot to this new vertex, and rewrite our LP to isolate x_5 instead of x_4 (by adding the first equation to the third):

$$\begin{array}{lrcrcrcrcl} \text{maximize} & x_1 & + & 2x_2 & + & 3x_3 & + & 4x_4 & + & 5x_5 \\ \text{subject to} & x_1 & & & & & + & x_4 & + & x_5 = 1 \\ & x_1 & + & x_2 & & & & & & = 3 \\ & & & & & x_3 & + & x_4 & & = 3 \\ & x_1 & , & x_2 & , & x_3 & , & x_4 & , & x_5 \geq 0 \end{array}$$

Again adding x_4 to the basis would only send us back to the vertex we were just at, so we continue by looking at the vertex we get by adding x_1 into the basis. Expressing the basis variables in terms of x_1 we have:

$$\begin{aligned} x_2 &= 3 - x_1 \\ x_3 &= 3 \\ x_5 &= 1 - x_1 \end{aligned}$$

So as we increase x_1, x_5 hits 0 first at $x_1 = 1$. The new basic feasible solution would be $(1, 2, 3, 0, 0)$, with an objective value of 14 (this is in fact the same vertex that we started at). The objective value of this vertex is smaller than our current objective value of 20. Since introducing x_1 or x_4 into the basis sends us to a solution with a smaller objective value, we conclude that we have reached an optimal solution for the LP, $(0, 3, 3, 0, 1)$ with objective value 20.

Note that our initial choice to bring x_4 into the basis instead of x_5 was arbitrary and ended up making things harder for ourselves. If we had started by bringing in x_5, we would have immediately reached the optimal solution, using only 1 pivot step instead of 2. There is no known method to choose pivots that get us to the optimal solution as fast as possible (or even in a guaranteed polynomial number of pivots).

Nevertheless, there are a variety of "pivot rules" that we can use to try to make smarter pivoting decisions with the simplex method. One simple pivot rule is to look at all adjacent vertices and choose the one which increases the objective value the most: e.g. instead of just bringing x_4 into the basis, we look at the new basic feasible solution we get by introducing x_4 and the one we get by introducing x_5 instead and pivot to whichever one has the larger objective value. This works pretty well in practice, although it is still possible to construct linear programs for which this rule takes exponentially many pivots to reach the optimal solution.

If the LP was unbounded then at some point while running the simplex method we would have reached a point where there is a variable that we could bring into the basis and increase as much as we want without any other variables going to 0.

21.4 Finding a Basic Feasible Solution

One step that we omitted in the previous section was how we found our starting basis in the first place. In the previous section we guessed that there would be a basic feasible solution with basis x_1, x_2, x_3 and we turned out to be right. This is not guaranteed to work, however.

For example, if we had started by trying to obtain a basic feasible solution for the LP in the previous section with basis x_1, x_2, x_4, then using the elimination method to isolate each variable would have given us the LP:

$$\begin{array}{lrcrcrcrcl} \text{maximize} & x_1 & + & 2x_2 & + & 3x_3 & + & 4x_4 & + & 5x_5 & \\ \text{subject to} & x_1 & & & - & x_3 & & & + & x_5 & = -2 \\ & & & x_2 & + & x_3 & & & - & x_5 & = 5 \\ & & & & & x_3 & + & x_4 & & & = 3 \\ & x_1 & , & x_2 & , & x_3 & , & x_4 & , & x_5 & \geq 0 \end{array}$$

Reading out the values for x_1, x_2, x_4 here we get the vector $x = (-2, 5, 0, 3, 0)$. But while these value satisfy the equations, this is *not* a feasible solution, because it does not satisfy the non-negativity constraints.

In order to find a basic feasible solution without having to guess the correct basis, we can use "two-phase" simplex method, where we create a new LP which has an obvious basic feasible solution, and whose solution gives us a basic feasible solution for the original LP. In particular, we use a similar technique to the one that we used for converting a standard form LP to equational form.

First, if there are any equations where the right hand side is negative, we multiply both sides of those equations by -1 (this will yield an equivalent equation). Then, we convert each equation $a_{j1}x_1 + a_{j2}x_2 + \cdots + a_{jn}x_n = b_j$ in the LP to $a_{j1}x_1 + a_{j2}x_2 + \cdots + a_{jn}x_n + z_j = b_j$ where z_j is a new variable which is only used in that equation. Since we have already guaranteed that the constants b_j on the right side of each equation are non-negative, this guarantees that our new LP has a basic feasible solution with basis $z_1, \ldots, z_m$ where we set all of the variables of the original LP to 0 and set each $z_j = b_j$. In order to try to find a basis for the original LP we make the objective of our new LP "minimize $z_1 + z_2 + \cdots + z_m$". For example, if our original LP is

$$\begin{array}{lrcrcrcrcl} \text{maximize} & x_1 & + & 2x_2 & + & 3x_3 & + & 4x_4 & + & 5x_5 & \\ \text{subject to} & x_1 & & & - & x_3 & & & + & x_5 & = -2 \\ & & & x_2 & + & x_3 & & & - & x_5 & = 5 \\ & & & & & x_3 & + & x_4 & & & = 3 \\ & x_1 & , & x_2 & , & x_3 & , & x_4 & , & x_5 & \geq 0 \end{array}$$

then our new LP would be

$$\begin{array}{lrcrcrcrcrcrcrcl} \text{minimize} & & & & & & & & & z_1 & + & z_2 & + & z_3 & \\ \text{subject to} & -x_1 & & & + & x_3 & & & - & x_5 & + & z_1 & & & = 2 \\ & & & x_2 & + & x_3 & & & - & x_5 & & & + & z_2 & = 5 \\ & & & & & x_3 & + & x_4 & & & & & & + & z_3 = 3 \\ & x_1 & , & x_2 & , & x_3 & , & x_4 & , & x_5 & , & z_1 & , & z_2 & , & z_3 \geq 0 \end{array}$$

If the new LP has an optimal solution with objective value 0 (i.e. a solution where all of our new variables are 0) we can use it as a starting basic feasible solution to run the simplex method on the original LP with the original objective function. On the other hand, if our new LP has an optimal solution with objective value greater than 0, then the original LP must be infeasible.

21.5 LP Duality

We will conclude our discussion of linear programming by talking about an important property of LPs known as *duality*, which can allow us to relate LPs for seemingly unrelated problems.

Suppose that we want to upper bound the optimum value for an LP in standard form. To take a concrete example, let's say that we want to give an upper bound for the best possible objective value in the LP below:

$$\begin{array}{lrcrcrcl} \text{maximize} & x_1 & + & x_2 & & & & \\ \text{subject to} & 2x_1 & + & x_2 & + & 4x_3 & \leq & 2 \\ & x_1 & + & 2x_2 & + & 3x_3 & \leq & 3 \\ & x_1 & , & x_2 & , & x_3 & \geq & 0 \end{array}$$

One observation that we can immediately make is that the optimal objective value is at most 2. This is because our non-negativity constraint guarantees that our objective function is less than the left hand side of our first constraint. That is,

$$x_1 + x_2 \leq 2x_1 + x_2 + 4x_3 \leq 2.$$

We can do better than this, though, by taking a linear combination of the constraints. If we take 1/3 of the first constraint plus 1/3 of the second constraint we get

$$x_1 + x_2 + (7/3)x_3 \leq 5/3.$$

Since all of the coefficients on the left hand side of this inequality are greater than or equal to the corresponding coefficients in our objective function (in this case the coefficients in the objective function are $1, 1$, and 0), the left hand side of this inequality gives us another upper bound on the objective function, and we can say that the optimal value is at most $5/3$.

What is the best upper bound we could obtain by taking a linear combination of the constraints like this? To determine this we can construct a new LP, where the variables represent the coefficients in a linear combination of the constraints of the original LP. We call the original LP the **primal** and this new LP the **dual**. The dual has a variable y_j for each constraint $\sum_{i=1}^{n} a_{ji}x_i \leq b_j$ in the primal. For example, the dual to our LP above will have two variables y_1 and y_2

In order to guarantee that our linear combination of the constraints actually forms an upper bound on the objective function, we need to guarantee that the coefficient of each variable in the left hand side of our linear combination of constraints is at least the coefficient of that variable in the objective function, i.e. $\sum_{j=1}^{m} a_{ji}y_j \geq c_i$. For our LP above this means that we need:

- $2y_1 + y_2 \geq 1$ (so that y_1 times the first constraint plus y_2 times the second will have a coefficient of at least 1 for x_1),
- $y_1 + 2y_2 \geq 1$ (so that y_1 times the first constraint plus y_2 times the second will have a coefficient of at least 1 for x_2),
- $4y_1 + 3y_2 \geq 0$ (so that y_1 times the first constraint plus y_2 times the second will have a coefficient of at least 0 for x_3),

We also need all of our new variables to be non-negative (since multiplying a constraint by a negative value flips the direction of the inequality).

Finally, in order for our linear combination to form the best possible upper bound on the objective value, we want the right hand side to be as small as possible. i.e. we want to minimize $\sum_{j=1}^{m} b_j y_j$. Putting it all together, we find that the dual of the LP from the start of this section is the following LP:

$$\begin{array}{ll} \text{minimize} & 2y_1 + 3y_2 \\ \text{subject to} & 2y_1 + \ y_2 \geq 1 \\ & y_1 + 2y_2 \geq 1 \\ & 4y_1 + 3y_2 \geq 0 \\ & y_1 \ , \ \ y_2 \geq 0 \end{array}$$

Recall that an LP in standard form can be expressed as:

$$\begin{array}{ll} \text{maximize} & c^T x \\ \text{s.t.} & Ax \leq b \\ & x \geq 0. \end{array}$$

The corresponding dual LP is:

$$\begin{array}{ll} \text{minimize} & b^T y \\ \text{s.t.} & A^T y \geq c \\ & y \geq 0. \end{array}$$

The reason we care about duality is that there is a close relationship between the optimal solutions to the primal and the dual. In particular, they always have equal value, as the following theorem states (this theorem is often known as "strong LP duality").

Theorem 95. *If an LP P in standard form is feasible and bounded then so is its dual D (and vice-versa). Moreover, their optimal values are equal.*

Since we want to avoid getting too into the algebraic weeds of linear programming, we will not prove this theorem. It has a variety of applications. In particular, we can view the max-flow min-cut theorem (Theorem 85) as a special case of LP duality.

21.6 Comprehension Questions

1. Convert the following LP to equational form:

$$\begin{array}{ll} \text{maximize} & 2x_1 + 3x_2 \\ \text{subject to} & x_1 - 2x_2 \leq 4 \\ & x_1 \ , \quad x_2 \geq 0 \end{array}$$

2. Suppose that $x = (1, 2, 3)$ and $x = (1, 2, 5)$ are both feasible solutions to an LP P with three variables. Give another solution that must be feasible for this LP.

3. True or false? If P is an LP with n variables and m constraints then the number of vertices in P's feasible region is $O(mn)$.

4. Suppose that P is a feasible and bounded LP in standard form with a feasible solution that achieves an objective value of 14. What can we say about the optimal objective value for the dual of P?

   (a) The optimal objective value for the dual is at most 14.

   (b) The optimal objective value for the dual is exactly 14.

   (c) The optimal objective value for the dual is at least 14.

   (d) There is not enough information provided to say anything about the optimal objective value for the dual.

5. How many variables does the dual of the LP below have?

$$\begin{array}{lrrrrrr} \text{maximize} & x_1 & - 2x_2 & + & x_3 & & \\ \text{subject to} & 2x_1 & & + 7x_3 & + 12x_4 & \leq 17 \\ & 3x_1 & + 8x_2 & & + 19x_4 & \leq 28 \\ & x_1 \ , & x_2 \ , & x_3 \ , & x_4 \geq & 0 \end{array}$$

21.7 Exercises

1. Convert the LP below to equational form

$$\begin{array}{lrrrrrr} \text{minimize} & x_1 & - 2x_2 & + & x_3 & & \\ \text{subject to} & 2x_1 & & + 7x_3 & + 12x_4 & \leq 50 \\ & 3x_1 & + 8x_2 & & + 19x_4 & \geq 79 \end{array}$$

2. Suppose we are running the simplex method on the LP in equational form below, using the pivot rule "pivot to whichever adjacent vertex increases the objective value the most". If we are currently at the vertex $x = (5, 8, 0, 0)$, which vertex will we pivot to next? What is the optimal solution to the LP?

$$\begin{array}{lrrrrrr} \text{maximize} & x_1 & - 2x_2 & + & x_3 & & \\ \text{subject to} & 2x_1 & + 5x_2 & + & 7x_3 & + 12x_4 = 50 \\ & 3x_1 & + 8x_2 & + & 11x_3 & + 19x_4 = 79 \\ & x_1 \ , & x_2 \ , & x_3 \ , & x_4 \geq & 0 \end{array}$$

3. Find the dual of each of the following LPs (you do not need to solve the LP).

   (a)

$$\begin{array}{ll} \text{maximize} & x_1 - 2x_2 + x_3 \\ \text{subject to} & 2x_1 + 5x_2 + 7x_3 + 12x_4 \leq 50 \\ & 3x_1 + 8x_2 + 11x_3 + 19x_4 \leq 79 \\ & x_1 \ , \ x_2 \ , \ x_3 \ , \ x_4 \geq 0 \end{array}$$

   (b)

$$\begin{array}{ll} \text{maximize} & x_1 + 2x_2 + 3x_4 - 2x_5 \\ \text{subject to} & 3x_2 - x_3 + 7x_5 \leq 14 \\ & -x_1 + 4x_3 + 3x_5 \leq 8 \\ & x_1 + x_2 + x_3 + x_4 + x_5 \leq 12 \\ & x_1 \ , \ x_2 \ , \ x_3 \ , \ x_4 \ , \ x_5 \geq 0 \end{array}$$

4. Suppose we are running the simplex method on the LP in equational form below, using the pivot rule "pivot to whichever adjacent vertex increases the objective value the most". If we are currently at the vertex $x = (3, 1, 0, 0)$, which vertex will we pivot to next? (You do not need to solve the LP)

$$\begin{array}{ll} \text{maximize} & x_3 + x_4 \\ \text{subject to} & x_1 + 3x_2 + x_3 = 6 \\ & 2x_1 + 4x_2 + x_4 = 10 \\ & x_1 \ , \ x_2 \ , \ x_3 \ , \ x_4 \geq 0 \end{array}$$

5. Find an optimal solution to the LP below. Show your work.

$$\begin{array}{ll} \text{maximize} & 2x_1 + 3x_2 \\ \text{subject to} & x_1 + 2x_2 \leq 4 \\ & 3x_1 - x_2 \leq 3 \\ & x_1 \ , \ x_2 \geq 0 \end{array}$$

6. An integer program is a linear program with the additional constraint that all variables must take integer values. This additional constraint allows integer programs to model some problems that linear programs cannot. Unfortunately solving integer programs turns out to be much harder than solving linear programs (at least with the best known algorithms). Thus, when the natural way to model a problem is with an integer program, computer scientists often consider the *LP relaxation* of that IP, which is the LP obtained by removing the constraint that the variables must have integer values. Sometimes it is possible to argue that this LP relaxation will have the same optimum value as the integer program, meaning that the LP relaxation still has an integral optimum solution.

   (a) Give an integer program for the bipartite matching problem (where we are given a bipartite graph G and need to find a matching in G with as many edges as possible). [Hint: create one variable for each edge of G]

   (b) Prove that the LP relaxation of your integer program from part (a) has an integral optimum solution [Hint: It might be helpful to apply what you know about flows]

   (c) Compute the dual of your LP relaxation from part (b)

   (d) Can you figure out what problem the dual you found in part (c) represents?