Introduction to Mathematical Programming
IE406

Lecture 8

Dr. Ted Ralphs

Reading for This Lecture

- Bertsimas 3.3, 3.5-3.7

The Tableau Method

- This is the standard method for solving LPs "by hand."
- We update the matrix $B^{-1}[b|A]$ instead of just B^{-1}.
- In addition, we keep track of the reduced costs in row "zero."
- This method is more expensive then revised simplex and is really just used for illustration.
- The method for updating the matrix is the same as in revised simplex.

What the Tableau Looks Like

- The tableau looks like this

$-c_B^\top B^{-1}b$	$c^\top - c_B^\top B^{-1}A$
$B^{-1}b$	$B^{-1}A$

- In more detail, this is

$-c_B^\top x_B$	$\bar{c}_1$	$\cdots$	$\bar{c}_n$
$x_{B(1)}$ $\vdots$ $x_{B(m)}$	$B^{-1}A_1$	$\cdots$	$B^{-1}A_n$

Parts of the Tableau

- Row zero contains the reduced costs.
- Column zero contains the values of the current basic variables.
- The upper left-hand corner entry is the opposite of the current objective function value.
- Each nonbasic column contains the feasible direction corresponding to increasing the given nonbasic variable.
- The basic columns are the columns of $B^{-1}B = I$, i.e., they are the unit vectors.
- All the information needed to perform an iteration of the simplex method is readily available.
- If variable j is to enter the basis, perform elementary row operations to turn column j of the tableau into the i^{th} unit vector, where i is the variable leaving the basis.

Implementing the Tableau Method

1. Start with the tableau associated with a specified BFS and associated basis B.
2. Examine the reduced costs in row zero and select a *pivot column* with $\bar{c}_j < 0$ if there is one. Otherwise, the current BFS is optimal.
3. Consider $u = B^{-1}A_j$, the j^{th} column of the tableau. If no component of u is positive, then the LP is unbounded.
4. Otherwise, compute the step size using the minimum ratio rule and determine the *pivot row*.
5. Scale the pivot row so that the *pivot element* becomes one.
6. Add a constant multiple of the pivot row to each other row of the tableau so that all other elements of the pivot column become zero.
7. Go to Step 2.

Finding an Initial Basic Feasible Solution

- Recall that we need an *initial BFS* to start the simplex methods.
- Ideally, the initial basis matrix would also be the identity matrix (easy to invert).
- In some cases, this is easy to achieve.
  - If the formulation contains only inequalities and the origin is feasible, then there is no problem.
  - The initial BFS is simply $x = 0, s = b$ and the initial basis matrix is the identity matrix (easy to invert).
- For problems with equalities or where the origin is not feasible, things are more difficult.
- To deal with this, we need the concept of an *artificial variable*.
- Artificial variables are used in the equality rows like slack variables, but then forced to zero.

Obtaining a Feasible Solution

- Suppose we are given an LP $min\{c^\top x | Ax = b, x \geq 0\}$ already in standard form.
- Assume without loss of generality that $b \geq 0$.
- Then we can obtain an initial BFS by solving the following auxiliary LP.

$$
\begin{aligned}
min \quad & \sum_{i=1}^{m} y_i \\
s.t. \quad & Ax + y = b \\
& x \geq 0 \\
& y \geq 0
\end{aligned}
$$

- If the optimal value for this problem is zero, then we obtain a feasible solution.
- Otherwise, the original problem was infeasible.
- This is usually called the *Phase I LP*.

Obtaining a Basis

- The only difficulty occurs when the solution to Phase I is degenerate and there are still artificial variables in the basis.
- In this case, we need to perform some *degenerate pivots* to rotate the artificial variables out of the basis.
- This is done in the normal way except that the pivot element doesn't have to be positive.
- If the pivot element is zero, then that row is linearly dependent on the others and we can eliminate it.
- After we rotate all the artificial variables out of the basis, we are left with a BFS and the corresponding basis.

Two-phase Simplex Method

1. By multiplying some of the constraints by -1, change the problem so that $b \geq 0$.
2. Introduce auxiliary variables $y_1, \ldots, y_m$, if necessary, and solve the auxiliary problem (Phase I).
3. If the optimal cost for the auxiliary problem is zero, then a feasible solution has been found. Otherwise, the problem is infeasible.
4. Eliminate any zero rows.
5. Rotate any remaining basic artificial variables out of the basis.
6. Drop the (nonbasic) artificial variables from the problem and compute the reduced costs for the nonbasic variables with the original objective function.
7. Continue the simplex algorithm as usual (Phase II).

The Big M method

- In the above method, we first solve the Phase I LP to obtain a BFS.
- Using that BFS, we start Phase II, which is solving the original problem.
- Another approach is to combine Phase I and Phase II.
- In this approach, we add the artificial variables and then change the objective function to

$$\sum_{j=1}^{n} c_j x_j + M \sum_{i=1}^{m} y_i$$

- M has to be large enough to force the artificial variables to zero.
- If the optimal solution has any artificial variables at nonzero level, then the original problem was infeasible.
- Otherwise, we obtain an optimal solution.
- In practice, two-phase simplex is usually used.

Computational Efficiency of the Simplex Method

- The efficiency of the method depends on the number of iterations.
- The number of iterations depends on how many extreme points are visited.
- It is easy to construct an example where there are 2^n extreme points and all of them are visited.
- This means in the worst case, the simplex method requires an exponential number of iterations.
- Of course, this depends on the pivoting rule, etc.
- What is the best case?

The Diameter of a Polyhedron

- The *distance* between two vertices of a polyhedron is the minimum number of pivots to get from one to the other.
- The *diameter* of a polyhedron is the maximum of the distances between all pairs of vertices.
- The diameter tells us something about the number of pivots we might be forced to make in the best case.
- Unfortunately, we don't know much about the diameter of general polyhedra.
- On average, the simplex algorithm performs extremely well if implemented correctly.