15-451/651 Algorithm Design & Analysis
Spring 2023, Recitation #8

Objectives

- To review zero-sum games and how to determine their values and optimal strategies.
- To understand dominant strategies
- To practice constructing linear programs
- To review NP-hardness reductions

Recitation Problems

1. **(Zero Sum Games & Dominance)** Consider the following zero-sum game:

		column player	
		A	**B**
row player	**1**	$(0, 0)$	$(1, -1)$
	2	$(3, -3)$	$(2, -2)$

(a) Let $\mathbf{p} = (p, 1 - p)$ be the row player's mixed strategy. As a function of p, what is the expected value to the row player $V_R(\mathbf{p}, \mathbf{q})$ when the column player plays column A or column B? Write both expressions, and then plot them against p.

(b) Based on your answer to part (a), what is the minimax optimal strategy for the row player, and what value is the value of the game?

It turns out that for this particular game, there's an even faster way to find this optimal strategy by observing that in every column, row 2 gives a better payoff for the row player than row 1. When this occurs, we say that row 2 dominates row 1.

In general we can say that pure strategy α dominates pure strategy β when

	row strategies	**column strategies**
strictly dominates	$R_{\alpha j} > R_{\beta j} \quad \forall j$	$C_{i\alpha} > C_{i\beta} \quad \forall i$
weakly dominates	$R_{\alpha j} \geq R_{\beta j} \quad \forall j$	$C_{i\alpha} \geq C_{i\beta} \quad \forall i$

and if a strategy is dominated by another strategy, we can say there is no reason to ever play it instead of that strategy. This can enable us to solve games faster, or even solve larger games than we would otherwise be able to. Consider for example, the following game:

		column player		
		A	**B**	**C**
row player	**1**	$(1, -1)$	$(2, -2)$	$(4, -4)$
	2	$(2, -2)$	$(-1, 1)$	$(0, 0)$
	3	$(2, -2)$	$(3, -3)$	$(1, -1)$

(c) Find and eliminate any dominated strategies. What game matrix does this leave you with?

(d) Now solve the resulting matrix to find the minimax optimal strategy for the row and column players and the value of the game.

2. **(Shortest Paths As An LP)**

Let's code up the s-t shortest-path problem as an LP.

The input is a directed graph G with edge weights $w(e) \geq 0$, start node s, and a target t. We want to find a path from s to t of least weight.

Suppose we have a variable f_e for every edge e, where $0 \leq f_e \leq 1$. We want these variables to somehow represent which edges are on the shortest path and which are not. Another way to say this is that we can think of f_e *like a flow* on the edges e, and our goal is to send a unit of flow from s to t along the shortest path.

(a) Keeping in mind that the f_e variables should behave like a flow, write down some suitable constraints and an objective function.

(b) We can now solve this LP using a polynomial-time LP solver like ellipsoid or Karmarkar's, but might get some weird solutions. How does this LP fail to perfectly capture the problem statement?

(c) In order to solve this issue, we decide to add functionality to our LP where we can constrain numbers to be integers, i.e. We can now add to our LP the constraints

$$f_e \in \mathbb{Z} \qquad \text{for every edge } e$$

Great! We now seem to have solved the problem. If we modify our LP solver to solve LP's with this additional constraint we can use it to find s-t shortest-paths. However, there is reason to believe that our modified LP solver is too good to be true. Give a good reason to believe that we might not be able to construct a poly-time LP solver using our modified LP.

Further Review

1. **(More Zero-Sum Games)** Consider the following zero-sum game:

		column player			
		A	**B**	**C**	**D**
row player	**1**	$(3, -3)$	$(2, -2)$	$(1, -1)$	$(3, -3)$
	2	$(4, -4)$	$(1, -1)$	$(3, -3)$	$(0, 0)$

   (a) What is the optimal strategy for the row player?

   (b) What is the optimal strategy for the column player?

   (c) What is the value of the game?

2. **(Non-Zero-Sum Games & Dominance)** In a non-zero-sum game, we remove the zero-sum requirement that $R_{ij} + C_{ij} = 0$, and so may have some outcomes that are good for both players or bad for both players. This makes them harder to solve, since neither player can assume their opponent is trying to minimize their score. We are no longer solving a minimax.

   However, the logic of dominance that we defined in problem 1 is still valid. There's still no reason why a player would pick a strategy that is dominated rather than the strategy that dominates it. Therefore, we can use dominance to solve for the outcome and payoffs of such games.

   (a) What are the dominant strategies and resulting payoffs of the following game, commonly known as the "prisoner's dilemma"?

		column player	
		A	**B**
row player	**1**	$(10, 10)$	$(11, -1)$
	2	$(-1, 11)$	$(0, 0)$

   (b) What are the dominant strategies and resulting payoffs of the following game?

		column player		
		A	**B**	**C**
row player	**1**	$(1, 2)$	$(3, 1)$	$(-1, 3)$
	2	$(2, 2)$	$(0, -1)$	$(1, 1)$
	3	$(1, 1)$	$(1, 0)$	$(2, 1)$

   Hints: A strategy that does not initially appear to be dominated may become dominated once you eliminate other strategies as possibilities.

3. **(Nash Equilibria & Quadratic Programming)** A *Nash equilibrium* is a pair of strategies $(\mathbf{p}, \mathbf{q})$ for the row and column players where $\mathbf{p}$ is an optimal response to $\mathbf{q}$, and $\mathbf{q}$ is an optimal response to $\mathbf{p}$, therefore if both players played these strategies, neither one would have an incentive to deviate and pick a different strategy.

   In zero-sum games, all Nash equilibria give the same payoffs, the value of the game, and both players playing optimally as we have traditionally solved for, is a Nash equilibrium. However, in non-zero-sum games, the situation is more complex. For example, consider the following game, commonly known as "the battle of the sexes"

		column player	
		A	**B**
row player	**1**	$(2, 1)$	$(-1, -1)$
	2	$(-1, -1)$	$(1, 2)$

   Note that $((1, 0), (1, 0))$ and $((0, 1), (0, 1))$ are both Nash equilibria for this game, as neither player has incentive to unilaterally deviate and risk a -1 payoff, but the two lead to different expected row and column player payoffs of $(2, 1)$ and $(1, 2)$ respectively, so the row player would prefer to end up in the first Nash equilibrium, while the column player would prefer to end up in the second.

   (a) It turns out there's one more Nash equilibrium for the game above. Find the strategies and expected payoffs corresponding to that equilibrium and justify that it truly is an equilibrium.

   In lecture, we discussed a way to compute optimal strategies, and thus Nash equilibria, for zero-sum games by representing it as a linear programming problem and then solving that linear program. Unfortunately, it turns out that computing Nash equilibria in general for non-zero-sum games is much harder. To do so, we'll need more powerful tools.

   (b) Quadratically constrained quadratic programs (QCQPs) are like linear programs, except that their objective functions and constraints may be quadratic rather than just linear, so they can include terms of up to degree 2. If you want to consider this formally, although you by no means need to, this means they have the form

   $$\begin{aligned} \text{maximize/minimize} \quad & \frac{1}{2}\mathbf{x}^\intercal C\mathbf{x} + \mathbf{c}^\intercal \mathbf{x} \\ \text{subject to} \quad & \frac{1}{2}\mathbf{x}^\intercal A_i \mathbf{x} + \mathbf{a}_i^\intercal \mathbf{x} \le b_i \quad \forall i \in [m] \end{aligned}$$

   where $C, A_i \in \mathbb{R}^{n\times n}$ represent the coefficients on the quadratic terms, $\mathbf{c}, \mathbf{a}_i \in \mathbb{R}^n$ represent the coefficients on the linear terms, $b_i \in \mathbb{R}$ represent the bounds, and $\mathbf{x} \in \mathbb{R}^n$ represents the variables to be optimized over.

   Write a QCQP that solves for the best Nash equilibrium for the row player for a game with arbitrary row and column payoff matrices R and C. Test it on the game above.

Hints: $\mathbf{p}$ and $\mathbf{q}$ must be optimal responses to each other, and may need to be mixed strategies. It might be helpful to know, that among the optimal responses to some strategy, there is always a pure strategy.

(c) Quadratic programs (QPs) are like linear programs, except that their objective functions may be quadratic, while their constraints must still be linear. If you want to consider this formally, this means they have the form

$$\begin{aligned} \text{maximize/minimize} \quad & \frac{1}{2}\mathbf{x}^\intercal C\mathbf{x} + \mathbf{c}^\intercal\mathbf{x} \\ \text{subject to} \quad & A\mathbf{x} \leq \mathbf{b} \end{aligned}$$

Write a QP that solves for the best Nash equilibrium for the row player for a game with arbitrary row and column payoff matrices R and C. Test it on the game above.

Hints: How did you find maxima and minima in calculus?

(d) Suppose you wanted to solve for the Nash equilibrium of a 3-player game. What sort of programming would you need in order to encode that problem?

4. **(Shortest Paths As An LP, Again)** It turns out that every problem you can represent as an LP, can actually be represented as two different, intuitively opposite LPs. We will explore this notion of "duality" further and generalize it in the next lecture, but here's a little previous.

In problem 2, we created an LP to find length of the shortest path from s to t with one variable for each edge, intuitively representing whether that edge was in the shortest path. Now try to create an LP for the same problem with one variable for each vertex, intuitively representing the length of the shortest path from s to that vertex, and justify its correctness.

Hints: The "dual" of a minimization problem will always be a maximization problem