AGTA Tutorial 5 Solutions

Exercise 1. Consider the *atomic network congestion game*, with three players, described by the following directed graph.

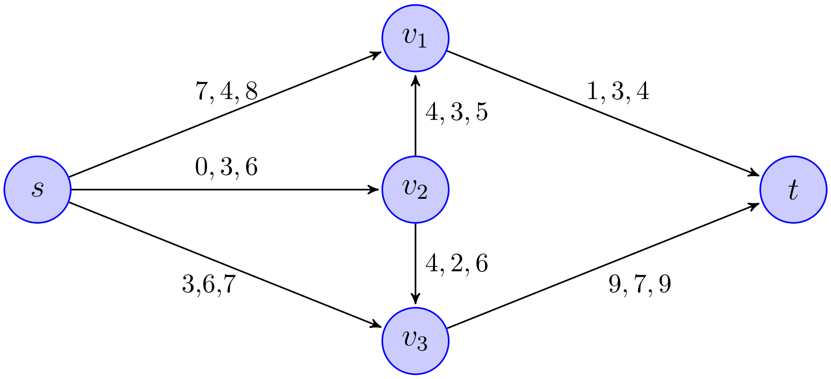

In this game, every player i (for $i = 1, 2, 3$) needs to choose a directed path from the source s to the target t. Thus, every player i's set of possible actions (i.e., its set of pure strategies) is the set of all possible directed paths from s to t.

Each edge e is labeled with a sequence of three numbers (c_1, c_2, c_3). Given a profile $\pi = (\pi_1, \pi_2, \pi_3)$ of pure strategies (i.e., s-t-paths) for all three players, the *cost* to player i of each directed edge, e, that is contained in player i's path π_i, is c_k, where k is the total number of players that have chosen edge e in their path. The total cost to player i, in the given profile π, is the sum of the costs of *all* the edges in its path π_i from s to t. Each player of course wants to minimize its own total cost.

Compute a pure strategy Nash Equilibrium in this atomic network congestion game.

Solution 1. As we are working with a congestion game, we can find a pure Nash Equilibrium by starting at any pure strategy profile, and iteratively improving it until we cannot improve any further. To get a concrete starting point, let's say all players take the route $s \to v_3 \to t$. Then we can do iterative improvements for example[1] as follows:

A. Player 1 switches to $s \to v_2 \to v_1 \to t$

B. Player 2 switches to $s \to v_2 \to v_1 \to t$

C. Player 3 switches to $s \to v_1 \to t$.

D. Player 2 switches to $s \to v_1 \to t$

At **D.** no further improvements can be made, so we reached the following PNE:

Player 1: $s \to v_2 \to v_1 \to t$
Player 2: $s \to v_1 \to t$
Player 3: $s \to v_1 \to t$

[1] At many stages in the process there is more than one option on who improves and how.

Note that in the above sequence we were not done at stage **C**, even though every player had switched once. Other starting points will take through other sequences of steps, and they might end up in a different PNE, although it turns out that in this game all pure Nash equilibria send two players via the route $s \to v_1 \to t$ and one via $s \to v_2 \to v_1 \to t$, differing only in which player chooses the path $s \to v_2 \to v_1 \to t$.

Note that there are "different" pure PNEs in this game for a trivial reason: we can arbitrarily rename the players in any permutation. However, this is only a different pure PNE in a trivial sense, by renaming which player plays which strategy. It turns out that in this game there is no other "fundamentally different" PNE, meaning a PNE consisting of an entirely different set of 3 paths from s to t.

In general however, such an atomic network congestion came can have multiple genuinely different PNEs, where the players have completely different costs in the different PNEs.

Exercise 2. Consider the class of linear network congestion games, i.e., network congestion games with linear cost functions $c_e(x) = a_e \cdot x + b_e$. For these games, the cost functions can be represented by providing a_e and b_e for each edge $e \in E$ by its binary representation. The strategy sets can be represented *implicitly*, by using the graph G as an input (e.g., represented as an adjacency list) and the origin-destination (source-target) pair (o_i, d_i) for each player i.

- Explain why the best response dynamics algorithm for finding a pure Nash equilibrium of such a game is still a pseudopolynomial time, under this representation of the cost functions.
- Now that the strategies are implicitly represented rather than explicitly, can the best response of an agent be computed in polynomial time? Justify your answer.

Solution 2. Before we discuss linear network congestion games, let us first discuss the representation of the inputs in the general class of congestion games. The input consists of the number of players n and the number of resources m. Since we need to specify strategies for each agent and costs for each resource, we need to be able to iterate over n and m, so these will not be represented in binary. So an algorithm that runs in time $\text{poly}(n, m)$ is indeed polynomial. The strategy set of each agent is given explicitly in the input as a set $S_i \subseteq 2^m$, i.e., it could be exponential in the number of resources m. Still, an algorithm that runs in time $\text{poly}(m, n, \max_i |S_i|)$ is polynomial.

Furthermore, the cost functions are part of the input, and these need to be represented as well. How do we represent a function? If the function is *general*, i.e., it does not have any specific structure that allows us to represent it more succinctly, we will have to represent it as a map (e.g., a two-dimensional array) by explicitly stating for each possible input (i.e., each number $n_i \in 1, 2, \dots, n$) the corresponding output, i.e., a number c_r, such that $c_r = c_r(n_i)$, where c_r is represented in binary. If we look at one entry of the map, the input is bounded by n, and the output is bounded by the binary representation of c_r, i.e., the number of bits needed to represent c_r. In terms of space, this is $\Theta(\log c_r)$. So, an algorithm that runs in time $\text{poly}(m, n, \max_i |S_i|, \log \max_r c_r)$ is polynomial in the size of the input, but an algorithm that runs in time $\text{poly}(m, n, \max_i |S_i|, \max_r c_r)$ is exponential (and in fact, pseudopolynomial, since c_r are the numerical parameters of the input).

- Now let's look at linear congestion games. Now the cost functions have a more explicit representation: for each resource (edge) e, we instead give a_e and b_e in the input; crucially, these are again represented in binary. Therefore, a polynomial time algorithm would need to run in time $\text{poly}(m, n, \log \max_e(a_e \cdot n + b_e))$. The best-response dynamics algorithm however runs in time $\text{poly}(m, n, \max_e(a_e \cdot n + b_e))$ and thus it is only pseudopolynomial. This fact is easy to verify, by looking at the maximum value of the potential function in Rosenthal's potential.
- The best response of an agent can still be computed in polynomial time. The potential issue here is that we no longer have an input representation of the strategies that is exponential in the number of resources m. In other words, we no longer provide all possible paths in the input, but only the origin and the destination for each player. Note that this (origin, destination) pair in itself does not

specify a strategy, but the strategy can be recovered by this pair as well as a path from the origin to the destination. To compute the best response, a player would need to be able to find a (o_i, d_i) pair that minimises her cost, fixing the strategies (paths) of the other players. This boils down to a minimum-cost path computation on the network, which can be done e.g., using the Bellman-Ford algorithm.

Exercise 3.

A. Consider any linear congestion game (i.e., a congestion game with linear cost functions), and any strategy profile $\mathbf{s} \in S_1 \times \ldots \times S_n$ in this game. Let $\Phi(\mathbf{s})$ be the value of Rosenthal's potential function on input $\mathbf{s}$, and let $\mathrm{SC}(\mathbf{s})$ be the social cost of $\mathbf{s}$. Show that

$$\frac{1}{2}\mathrm{SC}(\mathbf{s}) \leq \Phi(\mathbf{s}) \leq \mathrm{SC}(\mathbf{s})$$

Recall that Rosenthal's potential function is defined as:

$$\Phi(\mathbf{s}) = \sum_{r\in R} \sum_{j=1}^{n_r(\mathbf{s})} c_r(j),$$

where $n_r(\mathbf{s}) = \#(r, \mathbf{s})$ is the number of players that use resource r under strategy profile $\mathbf{s}$.

B. Consider a congestion game for which we have the following guarantee: The cost functions c_r are such that no resource is every used by more than λ players. Use the Potential Method (and Rosenthal's potential function) to show that the Price of Stability of any such game is at most λ.

Solution 3.

A. Recall that the social cost and Rosenthal's potential are defined, respectively, by

$$C(\mathbf{s}) = \sum_{r\in R} n_r c_r(n_r) \qquad \varPhi(\mathbf{s}) = \sum_{r\in R}\sum_{j=1}^{n_r} c_r(j),$$

where for ease of notation we denote $n_r = n_r(\mathbf{s})$ to be the set of players using resource r according to the strategy profile $\mathbf{s}$.

For the right-hand side inequality first, since the cost functions of the resources are nondecreasing, we have

$$\varPhi(\mathbf{s}) = \sum_{r\in R}\sum_{j=1}^{n_r} c_r(j) \leq \sum_{r\in R}\sum_{j=1}^{n_r} c_r(n_r) = \sum_{r\in R} n_r c_r(n_r) = C(\mathbf{s}).$$

For the left-hand side inequality, observe that for any natural number n and for $k = 0, 1, \ldots$, it holds that

$$\frac{1}{k+1} n^{k+1} \leq 1^k + 2^k + \cdots + n^k = \sum_{j=1}^{n} j^k. \tag{1}$$

Since each cost function can be expressed as a linear function $c_r(x) = a_r \cdot x + b_r$ we can upper bound the social cost by:

$$
\begin{aligned}
C(\mathbf{s}) &= \sum_{r\in R} n_r c_r(n_r) \\
&= \sum_{r\in R} n_r(a_r n_r + b_r) \\
&= \sum_{r\in R} (a_r \cdot n_r^2 + b_r \cdot n_r) \\
&\le \sum_{r\in R} \left(2a_r \sum_{j=1}^{n_r} j + b_r \sum_{j=1}^{n_r} 1 \right), && \text{due to (1),} \\
&\le 2 \sum_{r\in R} \sum_{k=0}^{1} a_{r,k} \sum_{j=1}^{n_r} j^k && \text{letting } a_{r,1} = a_r \text{ and } a_{r,0} = b_r \\
&= 2 \sum_{r\in R} \sum_{j=1}^{n_r} \sum_{k=0}^{1} a_{r,k} \cdot j^k \\
&= 2 \sum_{r\in R} \sum_{j=1}^{n_r} c_r(j) \\
&= 2\mathit{\Phi}(\mathbf{s}).
\end{aligned}
$$

B. Let $\mathbf{s}$ be a strategy profile and let $SC(\mathbf{s})$ be the social cost of that strategy profile. It holds that

$$
\begin{aligned}
SC(\mathbf{s}) \quad &= \quad \sum_{i=1}^{n} \mathrm{cost}_i(\mathbf{s}) = \sum_{i=1}^{n} \sum_{r\in s_i} c_r(n_r) = \sum_{r\in R} \sum_{i:r\in s_i} c_r(n_r) = \sum_{r\in R} n_r \cdot c_r(n_r) \\
&\le \quad \lambda \sum_{r\in R} c_r(n_r) \le \lambda \sum_{r\in R} \sum_{i=1}^{n_r} c_r(i) = \lambda\Phi(\mathbf{s}).
\end{aligned}
$$

where the first inequality follows from the fact that at most λ players are using resource r in any profile. As shown in Part **A.** above, we also have that

$$
\Phi(\mathbf{s}) = \sum_{r\in R} \sum_{i=1}^{n_r} c_r(i) \le \sum_{r\in R} \sum_{i=1}^{n_r} c_r(n_r) = \sum_{r\in R} n_r \cdot c_r(n_r) = SC(\mathbf{s})
$$

Now we start from a profile s^* with minimum social cost. If s^* is a pure Nash equilibrium, then the Price of Stability is 1 and we are done. Otherwise, we let the players deviate and best-respond until the reach a pure Nash equilibrium s. Since in each step some player decreases her cost and Φ is a potential function, we know that $\Phi(s) \le \Phi(s^*)$ and we have

$$
SC(s) \le \lambda\Phi(s) \le \lambda\Phi(s^*) \le \lambda SC(s^*),
$$

and the Price of Stability bound follows.

Exercise 4. Consider the class of *singleton congestion games*, i.e., congestion games in which the strategies of the players consist of single resources. Show that in singleton congestion games, a pure Nash equilibrium can be computed in polynomial time.

Hint: Starting from a singleton congestion game G, construct an "equivalent" game $\tilde{G}$ with $\tilde{c}_{\max} = poly(n, m)$, and argue that the best response dynamics algorithm converges in $\tilde{G}$ in polynomial time.

Solution 4. Fix a singleton congestion game $\mathcal{G} = (N, R, \{S_i\}, \{c_e\})$. Our goal is to construct an "equivalent" game $\tilde{\mathcal{G}} = (N, R, \{S_i\}, \{\tilde{c}_e\})$ with $\tilde{c}_{\max} = \mathsf{poly}(n, m)$. Then, we can deploy the best-response dynamics convergence theorem to get convergence in polynomially many steps, concluding our proof.

Let $K = \{c_r(j) \mid e \in R, j \in [n]\}$ be the set of all possible resource-cost values of $\mathcal{G}$. Notice that $|K| \leq nm$. We define modified cost functions $\tilde{c}_r$, for all resources $r \in R$, via

$$\tilde{c}_r(j) = k \quad \iff \quad |\{\kappa \in K \mid \kappa < c_r(j)\}| = k - 1. \tag{2}$$

In other words, $\tilde{c}_r(j)$ gives the position of the value $c_r(j)$ within K (when taken ordered from smallest to largest). By this, we can immediately deduce that indeed $\tilde{c}_r(j) \leq nm = \mathsf{poly}(n, m)$ for all resources r.

The equivalence of $\mathcal{G}$ and $\tilde{\mathcal{G}}$ can now be taken to mean that the two games preserve the better-response incentives of the players. Formally, we need to show that for all pure strategy profiles $\mathbf{s}$, every player i and any deviation s'_i, we have that:

$$C_i(s'_i, \mathbf{s}_{-i}) < C_i(\mathbf{s}) \quad \iff \quad \tilde{C}_i(s'_i, \mathbf{s}_{-i}) < \tilde{C}_i(\mathbf{s}).$$

Making use of the fact that our games are singleton, and denoting with r, r' the (single) resource in strategies s_i and s'_i, respectively, this condition is equivalently written as

$$c_{e'}(n_{r'} + 1) < c_r(n_r) \quad \iff \quad \tilde{c}_{r'}(n_{r'} + 1) < \tilde{c}_r(n_r), \tag{3}$$

where for simplicity we are using the shortcut notation $n_r = n_r(\mathbf{s})$, $n_{r'} = n_{r'}(\mathbf{s})$. Now it is not hard to establish the validity of (3), directly by observing that the definition of the modified costs $\tilde{c}_r$ in (2) preserves the relative order of the cost values.