

## Complexity of Pure Nash Equilibria in Congestion Games

ETH Zürich

Paul Dütting

We have seen that congestion games always have a pure Nash equilibrium, and that in singleton congestion games pure Nash equilibria can be found in polynomially many steps using (best response) improvement steps. What about more general congestion games? Do best response sequences always converge in polynomial time? Or can we at least compute a pure Nash equilibrium using a different algorithm in polynomial time? If we cannot, how would we give evidence of computational intractability?

Why do we care? Positive results, such as quick convergence of best response dynamics, clearly speak in favor of an equilibrium concept. Negative results, in turn, cast a shadow on the predictive power of an equilibrium concept. A famous quote in this context is Kamal Jain's "If your laptop cannot find it, then neither can the market."

## 1 Polynomial-time Algorithm for Symmetric Network Games

An interesting mid-ground are symmetric network congestion games.

**Definition 2.1.** *In a network congestion game, the set of resources is the set of edges  $E$  of a directed graph  $G = (V, E)$ . Player  $i$ 's strategies correspond to the set of paths between a fixed  $s_i \in V$  and  $t_i \in V$ .*


**Definition 2.2.** *A network congestion game is symmetric if all players share the same source  $s \in V$  and have a common target  $t \in V$ .*

It is known that there are instances of symmetric congestion games in which there are states such that every improvement sequence from this state to a pure Nash equilibrium has exponential length. Hence, applying improvement steps is not an *efficient* (i.e., polynomial time) algorithm for computing Nash equilibria in these games. However, there is another algorithm which finds pure Nash equilibria in polynomial time.

**Theorem 2.3** (Fabrikant, Papadimitriou, Talwar 2004). *In symmetric network congestion games with non-decreasing delays on each edge, a pure Nash equilibrium can be computed in polynomial time.*

*Proof.* Find a min-cost flow in the following graph:

- Each edge is replaced by  $n$  parallel edges of capacity 1 each.
- The  $i$ -th copy of edge  $e$  has cost  $d_e(i)$ ,  $1 \leq i \leq n$ .



The optimal solution minimizes Rosenthal's potential function and, hence, is a pure Nash equilibrium.  $\square$

## 2 The Complexity Class PLS

It turns out that for more general congestion games, computing a pure Nash equilibrium is a computationally hard problem. For this we will interpret the problem of finding a pure Nash equilibrium as the problem of finding a local optimum, and we will show that it is as hard as any other local search problem.

**Definition 2.4.** A local search problem  $\Pi$  is given by its set of instances  $\mathcal{I}_\Pi$ . For every instance  $I \in \mathcal{I}_\Pi$ , we are given a finite set of feasible solutions  $\mathcal{F}(I)$ , an objective function  $c : \mathcal{F}(I) \rightarrow \mathbb{Z}$ , and for every feasible solution  $s \in \mathcal{F}(I)$ , a neighborhood  $N(s, I) \subseteq \mathcal{F}(I)$ . A feasible solution  $s \in \mathcal{F}(I)$  is a local optimum if the objective value  $c(s)$  is at least as good as the objective value  $c(s')$  of every other feasible solution  $s' \in N(s, I)$ .

How hard is it to compute a local optimum? To make this a well defined question we need the following definition.

**Definition 2.5** (Johnson, Papadimitriou, Yannakakis 1988). A local search problem  $\Pi$  belongs to the class PLS, for Polynomial Local Search, if the following polynomial-time algorithms exist:

*Algorithm A:* Given an instance  $I \in \mathcal{I}_\Pi$ , return a feasible solution  $s \in \mathcal{F}(I)$ .

*Algorithm B:* Given an instance  $I \in \mathcal{I}_\Pi$  and a feasible solution  $s \in \mathcal{F}(I)$ , return the objective value  $c(s)$ .

*Algorithm C:* Given an instance  $I \in \mathcal{I}_\Pi$  and a feasible solution  $s \in \mathcal{F}(I)$ , either certify that  $s$  is a local optimum or return a solution  $s' \in N(s, I)$  with better objective value.

For every problem in PLS one can apply the following natural heuristic:

### Local Search Algorithm

1. Use Algorithm A to find a feasible solution  $s \in \mathcal{F}(I)$ .
2. Iteratively, use Algorithm C to find a better feasible solution  $s' \in N(s, I)$  until a locally optimal solution is found.

Note that this local search procedure is guaranteed to terminate because there are only finitely many candidate solutions, and in each iteration the objective function strictly improves. However, because there can be exponentially many feasible solutions, the local search algorithm need not run in polynomial time.

**Question:** For a given local search problem  $\Pi$ , is there a polynomial-time algorithm (not necessarily local search) for finding a local optimum?


## 3 Max-Cut and PLS-Completeness

**Definition 2.6** (Max-Cut). The search problem Max-Cut is defined as follows.

*Instances:* Graph  $G = (V, E)$  with edge weights  $w : E \rightarrow \mathbb{N}$ .  
*Feasible solutions:* A cut, which partitions  $V$  into two sets Left and Right.  
*Objective function:* The value of a cut is the weighted number of edges with one endpoint in Left and one endpoint in Right.  
*Neighborhood:* Two cuts are neighboring if one can obtain one from the other by moving only one node from Left to Right or vice versa.

**Observation 2.7.** (Membership in PLS) *Max-Cut is a PLS problem.*

**Example 2.8.** *Consider the following instance of Max-Cut:*



Consider weights  $w_{e_i} = i$  for  $1 \leq i \leq 6$ . Then the cut that separates  $Left = \{v_1, v_4\}$  from  $Right = \{v_2, v_3\}$  has a weight of 15. The four neighboring cuts each separate one vertex from the other three vertices. These cuts have weight 11, 8, 11, and 12. So the indicated cut is a local optimum. This cut, however, is not a global optimum: the cut that separates  $v_1, v_2$  from  $v_3, v_4$  has weight 17.

Max-Cut unlike Min-Cut is an NP-hard problem, but we are not interested in a globally optimal solution. We only want a local optimum. Intuitively, computing a locally optimal solution should be easier.

A concrete example is Max-Cut in graphs  $G = (V, E)$  with weights  $w_e = 1$  for all  $e \in E$ . Here, we can find a local optimum with the local search algorithm in at most  $|E|$  steps, while computing a maximum cut remains NP-hard.

Quite surprisingly, for Max-Cut in graphs with general weights no polynomial-time algorithm for computing a local optimum is known; and we can show that this problem is “as hard as” any other local search problem.

**Definition 2.9** (PLS-reduction). *Given two PLS problems  $\Pi_1$  and  $\Pi_2$ , there is a PLS-reduction (written  $\Pi_1 \leq_{PLS} \Pi_2$ ) if there are two polynomial-algorithms  $f$  and  $g$ :*

- Algorithm  $f$  maps every instance  $I \in \mathcal{I}_{\Pi_1}$  to an instance  $f(I) \in \mathcal{I}_{\Pi_2}$ .
- Algorithm  $g$  maps every local optimum  $s$  of  $f(I) \in \mathcal{I}_{\Pi_2}$  to a local optimum  $g(s)$  of  $I \in \mathcal{I}_{\Pi_1}$ .

As usual, one can read the symbol  $\leq_{PLS}$  as “is not harder than”. The reduction  $\Pi_1 \leq_{PLS} \Pi_2$  gives us a way to derive an algorithm for  $\Pi_1$  from an algorithm for  $\Pi_2$ .

**Definition 2.10** (PLS-completeness). *A problem  $\Pi^*$  in PLS is called PLS-complete if, for every problem  $\Pi$  in PLS, it holds  $\Pi \leq_{PLS} \Pi^*$ .*

It is generally assumed that there are problems in PLS that cannot be solved in polynomial time. For this reason, showing PLS-completeness effectively shows that presumably there is no polynomial-time algorithm.

**Theorem 2.11** (Schäffer and Yannakakis 1991). *Max-Cut is PLS-complete.*

As in the theory of NP-completeness, showing PLS-completeness requires an “initial” PLS-complete problem. Such a problem was given by Johnson et al.; PLS-completeness of other problems such as Max-Cut can then be established through reduction. It is worth pointing out that in the original problem, local search takes (worst-case) exponential time and that all known reductions preserve these bad instances.

## 4 PLS-Completeness of Pure Nash in General Congestion Games

Let us now interpret the problem of finding a pure Nash equilibrium in a congestion game as a local search problem. In this case, the feasible solutions are strategy profiles. They are neighboring if they differ in the choice of only a single player. We have seen before that in a unilateral deviation step, the Rosenthal potential changes by the same amount as the respective player’s cost. For this reason, pure Nash equilibria correspond to local minima of Rosenthal’s potential function.

**Observation 2.12** (Membership in PLS). *It is a PLS problem to find a pure Nash equilibrium in a congestion game.*


**Theorem 2.13.** *Max-Cut  $\leq_{PLS}$  Pure Nash Equilibrium in Congestion Games*

*Proof.* For this reduction, we have to map instances of Max-Cut to congestion games. Given a graph  $G = (V, E)$  with edge weights  $w: E \rightarrow \mathbb{N}$ , this game is defined as follows. Players correspond to the vertices  $V$ . For each edge  $e \in E$ , we add two resources  $r_e^{\text{left}}$  and  $r_e^{\text{right}}$ . The delays are defined by

$$d_{r_e^{\text{left}}}(k) = d_{r_e^{\text{right}}}(k) = \begin{cases} 0 & \text{for } k = 1 \\ w_e & \text{for } k \geq 2 \end{cases} .$$

Each player  $v \in V$  has two strategies, namely either to choose all the “left” resources for its incident edges  $\{r_e^{\text{left}} \mid v \in e\}$  or all the “right” resources for its incident edges  $\{r_e^{\text{right}} \mid v \in e\}$ .

This way, cuts in the graphs are in one-to-one correspondence to strategy profiles of the game. A cut of weight  $W$  is mapped to a strategy profile of Rosenthal potential  $\sum_{e \in E} w_e - W$  and vice versa. To see this, consider an edge  $e \in E$ . If its endpoints are in different sets of the cut, its resources contribute nothing to the potential; if its endpoints are in the same set, then the contribution is  $0 + w_e = w_e$ .



**Figure 1:** Potential game instance derived from the Max-Cut instance in Example 2.8. Players are color-coded. A player can either choose all solid or all dashed boxes in his color. The solid boxes correspond to the locally optimal cut that separates  $v_1, v_3$  from  $v_2, v_4$ . The only non-zero contributions to the potential function come from the edges that are contained in either Left or Right, namely  $e_2$  and  $e_4$  with weight 2 and 4. Their sum, 6, is equal to the total edge weight minus the weight of the cut,  $21 - 15$ .

Consequently, local maxima of Max-Cut correspond to local minima of the Rosenthal potential, which are exactly the pure Nash equilibria. Therefore, the second part of the reduction is again trivial.  $\square$

Note that the above reduction was to a congestion game in which the players’ strategy sets are not identical (so it is asymmetric) and required strategies to be arbitrary subsets of the resources (as opposed to, e.g., paths in a network). It turns out that either restriction can be dropped and the problem remains PLS-complete; but, as we have seen, dropping both turns the problem into one that we can solve in polynomial time.

## Recommended Literature

- D. S. Johnson, C. H. Papadimitriou, M. Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79-100, 1988. (Class PLS, Cook-Like Theorem for CircuitFlip)
- A. Fabrikant, C. H. Papadimitriou, K. Talwar. The complexity of pure Nash equilibria. *STOC 2004*. (First Proof of PLS-Completeness of Pure Nash in Congestion Games)
- A. A. Schäffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM Journal on Computing*, 20(1):56-87, 1991. (PLS-Completeness in Congestion Games via Max-Cut)
- H. Ackermann, H. Röglin, B. Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM*, 55(6), 2008. (Further Results on PLS-Completeness)