# Welfare Maximization in Congestion Games

Liad Blumrosen Shahar Dobzinski*

**Abstract**

Congestion games are non-cooperative games where the utility of a player from using a certain resource depends on the total number of players that are using the same resource. While most work so far took a distributed game-theoretic approach to this problem, this paper studies centralized solutions for congestion games. The first part of the paper analyzes the problem from a computational perspective. We analyze the computational complexity of the welfare-maximization problem, for which we provide both approximation algorithms and lower bounds. We study this optimization problem under different kinds of congestion effects (externalities) among the players: positive, negative, and unrestricted. Our main algorithmic result is a constant approximation algorithm for congestion games with unrestricted externalities. In the second part of the paper, we also take the strategic behavior of the players into account, and present centralized truthful mechanisms for congestion-game environments. Our main result in this part is an incentive-compatible mechanism for $m$-resource $n$-player congestion games that achieves an $O(\sqrt{m}\log n)$ approximation to the optimal welfare.

We also describe an important and useful connection between congestion games and combinatorial auctions. This connection allows us to use insights and methods from the combinatorial-auction literature for solving congestion-game problems.

*The School of Computer Science and Engineering, The Hebrew University of Jerusalem, {liad, shahard}@cs.huji.ac.il. The authors thank Noam Nisan, Moshe Babaioff, Michal Feldman, Ahuva Mu'alem and Michael Schapira for helpful discussions. The research was supported by grants from the Israel Science Foundation and the USA-Israel Bi-national Science Foundation.

# 1 Introduction

In congestion games, the players are presented with a set of resources (alternatives) to choose from, and the payoff associated with each resource may depend on the number of other players that use this resource. Such externalities effects may appear in different directions. For example, as the number of other drivers that choose to drive in the same road I took increases, my utility clearly decreases. In such cases, we say that the congestion game admits *negative externalities*. On the other hand, my utility increases as more users choose to use my favorite file-sharing system. Such congestion games admit *positive externalities*.

Congestion games were introduced in a seminal work by Rosenthal [30], and several extensions have been suggested ever-since. This paper adopts a model first presented by Milctaich [25] for *player-specific* congestion games, i.e., games where the preferences of the players may be non-identical. More formally, given a set $N$ of players, and a set $M$ of resources ($|N| = n$ and $|M| = m$), denote by $v_i(j,k)$ the non-negative value that player $i$ receives from using resource $j$ when a total number of $k$ players, including $i$, are using resource $j$. We refer to this model as the *anonymous* model (players do not care about the identity of the other players that share the resource with them). In environments with positive externalities (henceforth the POS model), for every player $i$ and resource $j$, and for every $1 \leq k \leq n-1$ we have $v_i(j,k) \leq v_i(j,k+1)$. If the valuations are monotone in the opposite direction (i.e., $v_i(j,k) \geq v_i(j,k+1)$), we have negative externalities (henceforth the NEG model).

In recent years, congestion games have been extensively studied by computer scientists. Since congestion games are known to always admit pure Nash equilibrium, most of this work studied the worst-case performance of these pure Nash equilibria ("the price of anarchy", see, e.g., the survey [31]), or the computational hardness of computing such equilibria (e.g., [14, 17]). It turns out that in many reasonable environments, a pure Nash equilibrium may achieve arbitrarily bad results or it may be hard to find. Therefore, other methods for coordination among players may be required. Chakrabarty, Mehta, Nagarajan and Vazirani [6] recently studied centralized solutions for congestion games. They characterized environments where cost minimization in congestion games can be obtained in polynomial time, and presented hardness of approximation results. In this paper, we study a similar problem of welfare *maximization* in congestion games, i.e., algorithms that maximize the sum of values of the players, $\Sigma_{j=1}^{m}\Sigma_{i\in S_j}v_i(j,|S_j|)$, where $S_1,...,S_n$ is an assignment of players to resources (i.e., $S_j \subseteq N$ denotes the set of players assigned to resource $j$). Our main challenge is to construct computationally-feasible centralized mechanisms for congestion games. Applications include, e.g., routing in networks, E-commerce and P2P systems.

## 1.1 Our Contribution

Designing such centralized mechanisms involve several difficulties. First, we should design algorithms for finding the optimal allocation, or when this is computationally hard, design approximation algorithms. In addition, we have to handle *incentive compatibility*, that is, ensure that our mechanisms achieve the desired results even if each player acts selfishly. We first study the computational questions, and then we initiate the study of incentive-compatible centralized mechanisms for congestion games, and present such mechanisms that approximate the optimal welfare.

Our first contribution is to point out a close connection between congestion games and combinatorial auctions. This connection may look surprising at first glance, but is actually simple: for every congestion game we build a "dual" combinatorial auction, where any assignment of players in the congestion game defines an equally-valued allocation in the dual combinatorial auction. This enables us to use intuitions and algorithms from the literature of combinatorial auctions for designing welfare-maximizing algorithms in congestion games. We also show that *negative* externalities in congestion games imply *substitutabilites* among the items in the dual combinatorial auction. Moreover, such combinatorial auctions belong to the special class of combinatorial auctions with "XOS" valuations, defined in [24]. We also observe that *positive* externalities in congestion games imply that the preferences in their dual combinatorial auctions will have *complementarities*, and in particular, they include all the supermodular valuations. Despite the similarities, there are significant differences between congestion games and combinatorial auctions (see Subsection 2.1).

We now mention our results regarding the computational complexity of congestion games:

- Our main algorithmic result:

**Theorem:** There exists an 18-approximation algorithm for approximating the welfare in congestion games with unrestricted externalities.

The algorithm constructs an initial infeasible assignment, using randomized rounding on the solution of the linear program. Given the congestions on the resources after the randomized-rounding procedure, we induce a feasible allocation with similar congestions. We show that this transition maintains a good percentage of the welfare using the following process: We pick a random set of players that hold, on expectation, a constant fraction of the infeasible assignment. Then, we make the assignment feasible by replacing the unchosen players with "dummy players" (these are players that were not assigned in the first assignment.) The replacement procedure is possible due to the special structure of the preferences in congestion games.

- Assuming negative externalities, we can improve this bound and achieve an $\frac{e}{e-1}$ approximation ratio based on the algorithm defined in [12] for $XOS$ valuations in combinatorial auctions. For using the combinatorial-auction algorithm we have to adjust the algorithm to our case. For example, we must make sure that the queries that the algorithm uses can be simulated in polynomial time in our model.
- Hardness results: we provide constant lower bounds for these approximation problems both for preferences with negative externalities and with positive externalities.
- Cost minimization: Although this paper centers on welfare maximization, we also study cost minimization in games with positive externalities. We provide an $O(\log n)$-approximation algorithm for cost minimization in the POS model, under a "cost-monotonicity" assumption. This algorithm may be viewed as a generalization of algorithms for non-metric facility-location problems, and indeed the analysis uses similar methods. As for a lower bound, we show that it is impossible to achieve better than an $O(\log n)$-approximation, unless $P = NP$.

Incentives issues are considered in the second part of the paper. Our main contributions are:

- We provide an incentive-compatible mechanism that guarantees a $O(\sqrt{m}\log n)$ approximation to the optimal welfare. The approximation can be achieved for players that have monotone preferences (either with positive or with negative externalities). The algorithm restricts the set of the possible assignments, and finds in polynomial time the best assignment over this restricted set. Therefore, VCG-based payments can be used to guarantee incentive compatibility.

- Inspired by a mechanism designed in [12] for combinatorial auctions, we construct an $O(\sqrt{n})$-approximation truthful mechanism for environments with negative externalities.
- We identify several special cases for which the optimal welfare can be determined in polynomial time. For these families, the VCG mechanism can be directly applied.

We also consider other congestion-game models, and show that they are either easy to solve or hard to approximate. This supports the claim that the agent-specific model is computationally the most interesting. For example, with *non-anonymous* valuations (when the players' preferences depend on the identity of the players that share the resources with them, not only on their number), only a trivial approximation ratio can be guaranteed.

## 1.2 Related Work

The closest work in the literature to ours is the recent work by Chakrabarty et al. [6] who were the first to study centralized solutions for congestion games. The questions we study in this paper are different than those studied in there in several substantial ways: first, we mainly aim to *maximize* the social welfare where [6] studied cost minimization. Optimizing the welfare may be considered economically more reasonable in many environments, and it also turns out to be more interesting from a computational perspective: While proving NP-hardness in both models is equivalent, approximating the optimal solution is a different task; For example, [6] proves that the optimal cost in congestion games with negative externalities cannot be approximated by any factor, where we show that the maximal welfare can be approximated within a factor of $\frac{e}{e-1} \approx 1.58$. Another difference is that [6] assumes negative externalities, where we also study environments with positive or unrestricted externalities. Our paper is the first to study algorithmic mechanism-design issues in congestion games, and we provide non-trivial solutions for these questions. Moreover, our incentive-compatible mechanisms adds to the set of rare examples of approximately-optimal truthful mechanisms for multi-parameter domains, i.e., domains where the secret information of each participant is composed of several values.[1]

Various kinds of congestion games and related settings were recently studied by computer scientists. For example, in selfish routing problems (see survey in [31]) and scheduling problems (e.g., [3]) negative externalities clearly exist. Other papers studied environments with positive externalities in the background, such as facility location and other cost-sharing environments (e.g., [16, 2]) and negotiations [7, 8]. However, as far as we know, this is the first paper to computationally analyze positive externalities in a general congestion-game model.

Congestion games were introduced by Rosenthal [30]. Monderer and Shapley [27] showed that they are isomorphic to the class of "potential games", and Milchtaich [25] extended the model to player-specific payoffs, and proved that with negative externalities a pure Nash equilibrium always exists also in this more general model. With positive externalities, pure Nash equilibria may not exist in this model [23]. The models we study also relate to the vast economic literature on network formation and group formation (see, e.g., [22, 13]).

The organization of the paper is as follows: in Section 2 we characterize the relation between congestion games and combinatorial auctions. In Section 3, we present our main approximation and hardness results. Finally, Section 4 describes mechanism design for congestion games. All the missing proofs appear in the appendix.

# 2 Congestion Games and Combinatorial Auctions

In this section we connect between welfare optimization in congestion games and in combinatorial auctions. In a *combinatorial auction*, a set of items are for sale to a set of bidders, and every player $j$ has a value $v_j(S) \in \Re^+$ for every bundle $S$ of items. The function $v_j$ depends only on the bundle $S$, and is called the valuation of player $j$. We assume that the valuation is *normalized* ($v_j(\emptyset) = 0$) and *monotone* (if $A \subseteq B$ then $v_j(A) \leq v(B)$). The goal is to find a partition of the items $S_1, .., S_n$, such that the welfare, $\Sigma_i v_i(S_i)$, is maximized. For a survey on combinatorial auctions see [9].

This section considers a general model of congestion games that allows *non-anonymous*

[1] The only examples for such deterministic mechanisms that we know of are [4, 11].

valuations: the value $v_i(j,S)$ of player $i$ for using the resource $j$ is a function of the identity of the players that use this resource, and not only of their quantity. Observe that with non-anonymous valuations, the private data of each player may have an exponential size.

Given a congestion game, we define a "dual" combinatorial auction, where the roles of the players and the resources are reversed: for each *resource* in the congestion game we define a bidder, and the players of the congestion game will be allocated to the resources.[2] We now have to define the valuations of the bidders in the combinatorial auction. As a first try, we define the valuation for every new bidder $j$ for a bundle $S$ as the sum of the values of the "items" in $S$ ($\Sigma_{i\in S} v_i(j,S)$). However, such valuations may not be monotone, because with negative externalities an additional player may decrease the total value. Therefore, since our model does not require assigning all bidders to resources, we define the value of a bundle $S$ by the subset $T \subseteq S$ with the maximal sum of values.[3]

**Definition 2.1** *Given a congestion game with a set $N$ of players, and a set $M$ of resources, define its dual combinatorial auction to be the following: the set of items is $N$, the set of bidders is $M$, and for each bidder $j$ we have the following valuation $\tilde{v}_j$*

$$\tilde{v}_j(S) = \max_{T\subseteq S} \Sigma_{i\in T} v_i(j,T)$$

Figure 2 illustrates a simple example for a congestion game and its dual combinatorial auction. A dual congestion game to a combinatorial auction can be defined in a similar way (see the proof for Theorem 2.4). An important property of the dual combinatorial auction, easily derived from its definition, is that an allocation in the combinatorial auction defines an assignment to the congestion game with the same value; this helps us using methods designed for approximating combinatorial auctions for solving congestion games.

[2]Recently, Monderer [26] also observed that congestion games are similar to combinatorial auctions with "strategic goods".

[3]This is another difference from the cost minimization problem; in cost minimization we must assign all players, otherwise the solution is trivial. Yet, allocating all players in our variant is not too restrictive: by assigning all unassigned players to the least valuable resource, we lose at most $\frac{1}{m}$ of the welfare.

Valuation of Player a

| | | Congestion | |
|---|---|---|---|
| | | 1 player | 2 players |
| Resource: | x | 3 | 1 |
| | y | 0 | 0 |

Valuation of Player b

| | | Congestion | |
|---|---|---|---|
| | | 1 player | 2 players |
| Resource: | x | 2 | 1 |
| | y | 2 | 1 |

Figure 1: The figure describes a congestion game with two players and two resources. Its dual combinatorial auction has two items, $a$ and $b$, two bidders, $x$ and $y$ with the following valuations: $v_x(\{a\}) = 3$, $v_x(\{b\}) = 2$, $v_x(\{ab\}) = 3$, and $v_y(\{a\}) = 0$, $v_y(\{b\}) = 2$, $v_y(\{ab\}) = 2$.

## 2.1 Combinatorial Auctions and Congestion Games: Differences

Despite the similarities, congestion games and combinatorial auctions are different both algorithmically and game theoretically.

First, the natural classes of players' valuations are different. For example, the class of combinatorial auctions that are dual to congestion games with positive externalities does not have an interesting semantic meaning. Another important difference is in the set of queries that are reasonable in each model; since bidders in combinatorial auctions may hold private data with an exponential size, many algorithms for combinatorial auctions are described in the "oracle" model – that is, the auctioneer queries the bidders with specific types of queries, and the outcome is determined accordingly. However, in congestion games, a reasonable model will query the players, where combinatorial-auction based algorithms for congestion games will actually query the resources. This difference has several implications. For example, while combinatorial auctions can be approximated by a factor of a square root of the number of items [5], we show that non-anonymous congestion games cannot be approximated by a better than a linear factor (in the number of players – the congestion-game equivalent to the combinatorial-auction items), even if the players are computationally unbounded.

**Proposition 2.2** *In non-anonymous congestion games, approximating the welfare to within a better factor than* $\frac{n}{2}$ *requires an exponential amount of communication. This result holds even if the players are computationally unlimited.*

The difference stems from the inability to simulate efficiently standard combinatorial-auction queries in congestion games. One widely-used query is the "demand query"[4]. Fortunately, we will be able to show that for *anonymous* congestion games, the simulation of demand queries can be done in polynomial time. One useful implication is that the linear-programming relaxation of the assignment problem can be solved in polynomial time [29, 5].

**Proposition 2.3** *Let $j$ be a resource in an anonymous congestion game, and let $\tilde{v}_j$ the corresponding valuation in the dual combinatorial auction problem. Then, given a vector of prices $(p_1, p_2, ..., p_n)$ a demand query for $\tilde{v}_j$ can be simulated in polynomial time.*

The idea behind the simulation is as follows: obviously, the size of the most profitable bundle is between 1 and $m$. We will check each possible bundle size, and find the most profitable bundle of that size. Finding the most profitable bundle of size $k$ can be done by calculating the value of $v_i(j, k) - p_i$ for each player $i$, sorting the values, and taking the first $k$ players. It is clear that the algorithm runs in polynomial time, and returns the correct answer.

The last crucial difference between combinatorial auctions and congestion games concerns dealing with the players' incentives. In congestion games, we think of the players as strategic. Under our reduction, this means that we think of the combinatorial-auction items as strategic. An immediate conclusion is that incentive-compatible mechanisms for combinatorial auctions do not necessarily imply the existence of incentive-compatible mechanisms for congestion games. Moreover, as opposed to standard combinatorial-auction models where bidders only care about their own allocation, the preferences in congestion games admit externalities of various sorts; this makes the construction of incentive-compatible mechanisms even harder.

## 2.2 Classes of Valuations

A question that arises is what types of combinatorial auctions are generated as duals of congestion games. For that, we define the following classes of combinatorial auctions. The first two classes are defined as duals of classes of congestion games, and the rest are known from the combinatorial-auction literature.

[4]In a demand query, the auctioneer presents a price $p_i$ per each item $i$, and the bidder responds with his favorite bundle under these prices, i.e., the bundle $S$ that maximizes her utility $v_j(S) - \sum_{i \in S} p_i$.

- Let **CG-POS** denote all combinatorial auctions that are dual to non-anonymous congestion game *with positive externalities*. Let **CG-NEG** denote all the combinatorial auctions that are dual to non-anonymous congestion game *with negative externalities*.
- Let **XOS** denote the family of combinatorial auctions with *XOS valuations* (defined in [24], see appendix A.1).
- Let **SF** denote the family of combinatorial auctions with *substitute-free* valuations, that is, for every two disjoint bundles $S, T$ we have $v(S)+v(T) \leq v(S \cup T)$. Similarly, let **CF** denote *complement-free* valuations where for all $S, T$ we have $v(S) + v(T) \geq v(S \cup T)$.
- Let **SUPM** denote the family of combinatorial auctions with *supermodular* valuations, that is, for every two bundles $S, T$ we have $v(S) + v(T) \leq v(S \cup T) + v(S \cap T)$. Similarly, let **SUBM** be *sub-modular* valuations, where for all bundles $S, T$ we have $v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$.

An easy conclusion from the reductions that we defined between congestion games and combinatorial auctions is that non-anonymous congestion games can express any combinatorial auction, and actually, the classes of congestion games and combinatorial auctions are isomorphic. Next, we ask what classes of combinatorial auctions are created from congestion games when these are restricted to have negative (or positive) externalities. We show that *negative externalities* imply that the dual combinatorial auctions have "substitutabilities". Specifically, we show that this class of auctions is contained[5] in the class of XOS valuations, known to lie between the "complement-free" class and the "sub-modular" class [24] (that is, $SUBM \subsetneq$CG-NEG$\subseteq XOS \subsetneq CF$). Luckily, approximation algorithms are known for XOS valuations in combinatorial auctions [12], and will be helpful in our context (see Section 3.2). As for *positive externalities*, we show that the class *CG-POS* lies in a similar hierarchy for valuations with complementarities, that is, it lies strictly between "substitute-free" valuations and "super-modular" valuations.[6]

**Theorem 2.4** *(1)* $SUBM \subsetneq CG\text{-}NEG \subseteq XOS$ *(2)* $SUPM \subsetneq CG\text{-}POS \subsetneq SF$

[5]We do not know if this inclusion is strict.

[6]Although super-modular valuations are hard to approximate in combinatorial auctions, we show in Section 3 a constant approximation for anonymous congestion games with positive externalities.

# 3 Approximating the Welfare

This section explores the approximability of the welfare in congestion games. We first present our main algorithmic result: an $O(1)$-approximation algorithm for anonymous congestion games without any restrictions on the externalities. We first outline the algorithm, and provide intuitions about the analysis.

A quick overview of the algorithm: we first solve the linear relaxation of the problem, and then we use randomized rounding to find the congestion (i.e., the number of players) on each resource. If the result of the randomized rounding procedure is "good enough", although possibly not feasible, our algorithm will find the best feasible assignment of players to resources over all the assignments with exactly the same congestions on each resource as in the result of the randomized rounding. This can be done using classic matching results, as shown in [6]. As a last step, we check that this result cannot be outperformed by a trivial solution (where the players are allowed to be assigned only to a single resource).

The analysis goes as follows: the randomized rounding implicitly defines a "pre-assignment" that has a value close to the optimum. The caveat is that this assignment may not be feasible, as a single player may be assigned to multiple resources, instead of just one.

The main challenge in the analysis of the algorithm is in analyzing the way the implicit pre-assignment can be converted into a feasible assignment. In the proof, we will show the *existence* of a feasible solution with a value close to the optimal with the same congestions on the resources as in the pre-assignment. Given that such a solution exists, we find an explicit (and possibly different) assignment with at least that value.

The existence of such a feasible solution is shown by presenting a randomized "process" that converts the pre-assignment into a feasible solution, taking advantage of the special structure of the players' preferences in congestion games. The key idea in proving the existence is to keep aside during the analysis a sufficient number of *dummy players*. That is, we make sure that "a large" number of players are not assigned to any resource at all in the pre-assignment. Thus, these dummy players they contribute nothing to the welfare in the pre-assignment, and we are free to assign them to resources as we want. For each player that is assigned to multiple resources in the pre-assignment, we re-assign all his resources but one to dummy players. Not only that the pre-assignment is now a feasible solution, but also the congestion on each resource hasn't been changed. Since in anonymous congestion games the value of each player only depends on the number of the other players that are assigned to a certain resource and not on their identity, a player that is assigned to a resource both in the pre-assignment and in the current assignment will face the same congestion in both. Therefore, his value for being assigned to that resource will not change.

We are still left with one important issue of selecting the single resource a player will be assigned to. The selection process used in the analysis is inspired by a technique from a recent paper of Feige [15]. This technique is based on a lottery process that selects a resource uniformly at random, from the set of resources a player was assigned to in the pre-assignment. We prove that the *expected* value achieved after the second lottery is close to that of the optimal solution.

## 3.1 An Algorithm for Unrestricted Valuations

**Input:** A congestion game with $n$ unrestricted valuations $v_i$ for $m$ resources.

**Output:** An assignment which is an $O(1)$-approximation to the optimal assignment.

1. Solve the linear-programming relaxation of the problem.

   *Maximize:* $\sum_{j,\, S,\, i \in S} x_{j,S} v_i(j, |S|)$

   | | | |
   |---|---|---|
   | *Subject To:* | For each resource $j$: | $\Sigma_S \, x_{j,S} \leq 1$ |
   | | For each player $i$: | $\Sigma_{j,\, S \mid i \in S} \, x_{j,S} \leq 1$ |
   | | For each $j$, $S$: | $x_{j,S} \geq 0$ |

   Denote the optimal value of the linear programming by $OPT^*$, and the solution by $x_{j,S}$ (for every resource $j$ and set of players $S$).

2. For each resource $j$, select at most one set to assign, such that the set $S$ of players is chosen with probability $\frac{x_{j,S}}{2.2}$, and the empty set is chosen with probability $1 - \Sigma_{j,S} \frac{x_{j,S}}{2.2}$. Denote this pre-assignment by $S_1, ..., S_m$.

3. Proceed only if the following two constraints hold:

   (a) The total number of assignments $\Sigma_{j=1}^{m} |S_j|$ is at most $n$.

(b) The total value of the pre-assignment is large. I.e., $\Sigma_{j,S,i\in S} v_i(j,|S_j|) \geq \frac{OPT^*}{6}$

4. Return the assignment with the maximal welfare over the next two:

   (a) The best assignment in which the congestion on item $j$ is exactly $|S_j|$ (see Claim B.1).

   (b) The best assignment in which all the assigned players are assigned to the same resource, where the other players are not assigned at all.

**Theorem 3.1** *The algorithm above is a polynomial-time* 18*-approximation for the optimal welfare in congestion games with unrestricted preferences with high probability.*

**Proof:** We prove that the approximation guarantee is correct, and later we show that the algorithm runs in polynomial time. We first claim that, with good probability, the conditions of Stage 3 hold (the claim is proved in Appendix B).

**Claim 3.2** *The two conditions in Stage 3 hold with probability of at least* 0.24*.*

We define the *pre-assignment* to be the result of the (not necessarily feasible) assignment achieved after Stage 2 of the algorithm. We will show that if the conditions of Stage 3 hold, then the following process provides an assignment with an expected value of at least $\frac{1}{18}$ of the optimal one, with the same congestion on each resource as in the pre-assignment. Since in Stage 4a we are maximizing over all such assignments that have exactly the same number of players in each resource as in the result of the randomized rounding, the theorem follows. We prove this assuming that no player can capture more than a $\frac{1}{18}$ fraction of the optimal welfare (otherwise, the approximation result is guaranteed by Stage 4).

**The Process:**

1. Let $N_W$ be the set of players assigned to resources in the pre-assignment. I.e., $N_W = \cup_{j=1}^{m} S_j$. For each player $i \in N_W$, choose uniformly at random a resource $r(i)$ from the resources to which he was assigned $\{S_j | i \in S_j\}$.

2. For each player $i$ in $N_W$, and for each resource $j$ assigned to player $i$ in the pre-assignment, i.e., $i \in S_j$ (initialize $FILL = \emptyset$):

- If $r(i) \neq j$, assign an arbitrary player not in $N_W$ instead of player $i$. More formally:
  (a) Arbitrarily choose a player $p$ from $N \setminus (N_W \cup FILL)$.
  (b) Assign player $p$ to resource $j$ instead of player $i$, that is, $r(p) = j$.
  (c) Add $p$ to $FILL$, i.e., $FILL = FILL \cup \{p\}$.

Notice that the process is well defined (namely, we always have a player which is not assigned to any resource yet during Step 2 of the process) as the conditions of Stage 3 imply that the total number of assignments is at most $n$.

Let PRE be the welfare achieved by the pre-assignment (after the randomized-rounding stage and given that Condition 3 holds). Denote by PROC the value achieved by the process. We will prove that $E[PROC] \geq \frac{E[PRE]}{2.89}$. Due to the restriction in Stage 3 of the algorithm, PRE holds at least 1/6 of OPT*. Thus, there must be an assignment which holds at least a fraction of $\frac{1}{18}$ of the optimal fractional welfare, and the lemma will follow.

Let C3 be the event: "the conditions in Stage 3 of the algorithm hold". Let $S^{C3}$ be the set of events that correspond to all possible pre-assignments by the randomized-rounding procedure for which C3 holds. For every $\sigma \in S^{C3}$, we use the notation $\sigma_j$ to denote the set of players assigned to resource $j$ in the pre-assignment corresponding to $\sigma$, and $\Pr(\sigma)$ denote the probability for choosing the pre-assignment $\sigma$. Using these notations we have that (enumerating on all players $i \in N$, resources $j \in M$ and sets $S \subseteq N$),

$$E[PRE] = \Sigma_{S,j,i} \Sigma_{\sigma | i \in S, S = \sigma_j} \Pr(\sigma) v_i(j, |S|) \tag{1}$$

For every $\sigma \in S^{C3}$ we define $n(i, \sigma)$ to be the number of resources player $i$ is assigned to in the pre-assignment $\sigma$. For every player $i$, denote by $n_i$ the random variable that indicates the number of times player $i$ is assigned. Note that in PROC the expected value of player $i$ in the pre-assignment $\sigma$ is divided by $n(i, \sigma)$ (comparing to PRE), due to the uniform lottery of the process. Therefore (we define $v_i(j, |S|) = 0$ if $j \notin S$),

$$\begin{aligned} E[PROC] &= \Sigma_{S,j,i} \Sigma_{\sigma | i \in S, S = \sigma_j} \Pr(\sigma) \frac{v_i(j, |S|)}{n(i, \sigma)} \geq \Sigma_{S,j,i} \frac{v_i(j, |S|)}{\Sigma_{\sigma | i \in S, S = \sigma_j} \Pr(\sigma) n(i, \sigma)} && (2) \\ &= \Sigma_{S,j,i} \frac{\left(\Sigma_{\sigma | i \in S, S = \sigma_j} \Pr(\sigma)\right) v_i(j, |S|)}{E[n_i | S \text{ is assigned to } j, \text{ and } C3]} && (3) \end{aligned}$$

where the first inequality is due to the convexity of $1/x$ and the Jensen's inequality[7], and the last equality is owing to the definition of (conditional) expectation (note that the term $\Sigma_{\sigma|i\in S,S=\sigma_j}\Pr(\sigma)$ is needed for the normalization of the probabilities).

Therefore, looking at Equations 1 and 3, it is sufficient to prove the next claim for proving that $E[PROC] \geq \frac{E[PRE]}{2.89}$.

**Claim 3.3** *For every player $i$, resource $j$, and set $S$ such that $i \in S$, we have that:*

$$E[n_i|S \textit{ is assigned to } j, \textit{ and } C3] \leq 2.89$$

**Proof:** Since player $i$ is assigned at most once to resource $j$ we have that:

$$E[n_i|S \text{ is assigned to } j, \text{ and } C3] \leq E[n_i|C3] + 1$$

It is also easy to see that:

$$E[n_i] = \Pr(C3) \cdot E[n_i|C3] + (1 - \Pr(C3)) \cdot E[n_i|\neg C3) \geq \Pr(C3) \cdot E[n_i|C3]$$

So now we have

$$E[n_i|S \text{ is assigned to } j, \text{ and } C3] \leq E[n_i|C3] + 1 \leq \frac{E[n_i]}{\Pr[C3]} + 1 \leq \frac{1}{2.2 \cdot 0.24} + 1 \leq 2.89$$

where the second-rightmost inequality holds because $E[n_i] \leq \frac{1}{2.2}$, due to the constraints of the LP and because the $x_{j,S}$'s are divided by 2.2, and also due to Claim 3.2. □

We now show that the algorithm runs in polynomial time. It is straightforward to see that the algorithm runs in polynomial time for all stages except for Stages 1, and 4. The linear program has an exponential number of variables, but it can still be solved in polynomial time by solving the dual linear program using the ellipsoid method. As observed in [29, 5], the ellipsoid method requires a "separation" oracle, and this oracle can be simulated using a demand-query oracle. Since all the valuations are anonymous, a demand-query oracle can be implemented in polynomial time (Lemma 2.3). In Stage 4 we calculate the best assignment when the congestion on each resource is known, and this was shown possible in [6]. In stage 4b, we also enumerate on all the possible congestions on each resource. □

[7] Jensen's inequality says that if $X$ is a random variable and $f$ is a convex function, then $E[f(X)] \geq f(E[x])$.

### 3.2 An $\frac{e}{e-1}$-Approximation in the $NEG$ Model

In Section 2 we have proved that the dual combinatorial auction of a congestion game with negative externalities belongs to the special class of combinatorial auction with XOS valuations (see Appendix A.1 for a definition). We would now like to take advantage of the $\frac{e}{e-1}$ approximation algorithm for approximating combinatorial auctions with XOS valuations [12]. However, this algorithm uses a special type of queries (called "XOS queries") in addition to standard demand queries. By showing that these types of queries can be efficiently computed in the anonymous model, we adopt this algorithm to congestion-game environments.

**Theorem 3.4** *There exists an $\frac{e}{e-1}$-approximation algorithm for congestion game, when all valuations are in the anonymous $NEG$ model.*

**Proof:** Theorem 2.4 proves that the every congestion game with negative externalities has a dual combinatorial auction with XOS valuations. To use the algorithm of [12], we need to show that the we can simulate the oracles required by the algorithm (XOS oracles and demand oracles) in polynomial time. We note again that querying the bidders in the dual combinatorial auction is equivalent to querying the resources in the congestion game. Proposition 2.3 shows that demand oracles of the dual combinatorial auction can be simulated in polynomial time. XOS oracles (see Appendix A.1 for definition) can also be simulated in polynomial time: since the combinatorial auction instance was generated by a congestion game, observe that the maximizing additive valuation for some bundle $S$ in the XOS valuation dual to resource $j$ is defined by at most $t < |S|$ non-zero values, where the $i$'th non-zero value is the value of assigning the corresponding player to resource $j$ under congestion $t$. Notice that these values are part of the input. Now all we have to do is to find the value of $t$ that maximizes the sum of the non-zero values, which can be done in a way similar to the simulation of a demand oracle.

We can now run the algorithm of [12], and get an allocation of the items in the dual combinatorial auction. We deduce an assignment of the players of the congestion game with the same value. The algorithm of [12] provides an $\frac{e}{e-1}$-approximation to the optimal welfare, and so we get the same approximation ratio for the congestion game. □

Can a similar approximation ratio be achieved also for valuations with positive externalities? In Proposition B.2 in Appendix B we provide a partial negative answer for this question, and show that the integrality gap in the linear-programming formulation of the problem is at least 2, even in the $POS$ model. Thus, any LP-based algorithm cannot achieve a higher approximation ratio.

### 3.3 Hardness Results

The next theorem proves constant lower bounds for approximating congestion games with positive and negative externalities. The lower bound for the $NEG$ model is inspired by a reduction in [6], but requires a more careful analysis. For proving the lower bound in the $POS$ model, we use a reduction from $k$-Dimensional Matching using a lower bound from [20].

**Theorem 3.5** *There is no $(\frac{32}{31}-\epsilon)$-approximation for congestion games for any $\epsilon > 0$, unless $P = NP$. In particular, unless $P = NP$ and for any $\epsilon > 0$:*

1. *There is no $(\frac{32}{31}-\epsilon)$-approximation for the welfare in the $NEG$ model.*
2. *There is no $(\frac{138}{137}-\epsilon)$-approximation for the welfare in the $POS$ model.*

**Proof:** (1) To prove the first claim, we reduce from MAX-3SAT. Hastad [19] proved that it is NP-hard to distinguish between the following two cases: either a fraction of $(1-\epsilon)$ of the clauses can be satisfied, or at most $\frac{7}{8}$ of the clauses can be satisfied. A closer look on the proof of [19] shows that the result holds even if for each variable $x$, the number of appearances of $\neg x$ and $x$ is equal.

The reduction: Given an instance of MAX-3SAT with $n$ variables and $m$ clauses, define a congestion game with $2n$ resources: a resource for each variable, and another resource for its negation. For each variable $x$, define $t_x$ dummy players, where $t_x$ is the number of appearances of $x$ in the 3SAT formula. Each one of these $t_x$ dummy players receives a value of 1 for being assigned to resources $x$ or $\neg x$, regardless of the congestion. His value for being assigned to any other resource is 0. In addition, define $m$ more players, one for each clause, that get the value of 1 if assigned to a resource which "satisfies" the clause, but only if the congestion on that resource is at most $\frac{t_x}{2}$. Otherwise, these players get the value of 0. Observe that all valuations are indeed in the $NEG$ model.

Analysis: We claim that for each solution of value $t$ of the congestion game produced by our reduction, $t-3m$ clauses of the corresponding 3SAT formula can be satisfied: first, if there exists a dummy player who is not assigned to any resource, or assigned to a resource in which he gets the value of 0, re-assign it to one of the resources that he gets a value of 1 for being assigned to them. The value of the congestion game solution may decreased, since the clause players are sensitive to the congestion. In this case, de-assign exactly one clause player from that resource. Now the solution value is at least as in the beginning of the process.

The assignment itself can be determined as follows: for each clause player $i$ that was assigned to resource $x$ and gets the value of 1, set the value of $x$ to true (if player $i$ was assigned to resource $\neg x$, we set the value of $x$ to false). The assignment is indeed valid: assume, on the way of contradiction, that there is a clause player which gets the value of 1 and is assigned to resource $x$, and another clause player that also gets the value of 1 but is assigned to $\neg x$. Because all the $t_x$ dummy players are assigned either ro resource $x$ or to $\neg x$, either resource $x$ or resource $\neg x$ have $\frac{t_x}{2}$ players assigned to it. Thus, no clause player can get the value of 1 when assigned to that resource – a contradiction. To finish, arbitrarily set the value of variables which their value hasn't been determined.

Now observe that $\Sigma_x t_x = 3m$. If the 3SAT formula can be satisfied by $(1-\epsilon)m$ clauses, than there is an assignment in the congestion game in which all the dummy players, and a fraction of $(1-\epsilon)$ of the clause players, contribute a value of 1 each: look at the optimal assignment to the 3SAT formula. If the variable $x$ gets the value of true, assign all $t_x$ dummy players of $x$ to the resource $\neg x$. Otherwise, assign them to the resource $x$. Assign each clause player to a resource that "satisfies" it, according to the optimal assignment of the 3SAT. Since each variable can satisfy at most $\frac{t_x}{2}$ clauses, all clause players get the value of 1.

If at most a fraction of $\frac{7}{8}$ of the clauses can be satisfied, the optimal assignment to the congestion game has a value of at most $\Sigma_x t_x + \frac{7m}{8}$, as explained before. Now we have that $\frac{OPT}{ALG} \geq \frac{(1-\epsilon)m+\Sigma_x t_x}{\frac{7m}{8}+\Sigma_x t_x} = \frac{(1-\epsilon)+3}{\frac{7}{8}+3} = \frac{32}{31} - \frac{8}{31}\epsilon$.

(2) For proving the hardness of the $POS$ model we reduce from the $k$-dimensional matching problem ($k$-DM). In this problem we are given $k$ disjoint sets of objects $V_1, ..., V_k$, each of size $n$, and a set of $M = |m|$ possible matchings, $|M| \subseteq V_1 \times ... \times V_k$. The goal is to maximize the number of disjoint matchings. That is, assign each $j \in \cup_i V_i$ to at most one matching, to

maximize the number of matchings that are assigned all their $k$ objects.

The reduction: Define a resource for each matching $t \in M$. Also define $k \cdot n$ players, one for each object in $\cup_i V_i$. Let the valuation of a player corresponding to item $j$ be 1 if the player is assigned to a resource corresponding to a matching it belongs, and the congestion is at least $k$, and 0 otherwise.

Analysis: Observe that given a K-DM instance in which at most $v$ matchings can be satisfied, the optimal value of the corresponding congestion game is between $v \cdot k$ and $v \cdot k + (m-v)(k-1)$. For the lower bound, assign the players "according" to the $v$ matchings. For the upper bound, assign the players according to the $v$ matchings, and assign for each of the remaining $(m-v)$ matchings at most $(k-1)$ players that get the value of 1. A solution to the congestion game with a better value implies that the K-DM instance has a value strictly larger than $v$.

In [20] it is proved that for any $\epsilon > 0$ it is NP-hard to distinguish between the following cases of 6-DM: either at least $(1-\epsilon)m$ matchings are satisfied, or at most $\frac{22}{23}m$ matchings are satisfied. Therefore, $\frac{OPT}{ALG} \geq \frac{(1-\epsilon)6m}{(\frac{22}{23})6m+\frac{1}{23}5m} = \frac{138}{137} - \frac{138}{137}\epsilon$. □

We also show that generalizations of player-specific congestion games are hard to approximate. In Proposition 2.2 we have proved that congestion games with unrestricted *non-anonymous* valuations cannot be approximated by a factor of $O(n)$. We show that even if the non-anonymous valuations exhibit only positive or negative externalities, the optimal welfare is still hard to approximate (within better factors than $\sqrt{n}$ and $n$, respectively). We also show that enabling the bidders to choose sets of resources ("non-simple" games) also makes the problem hard. All of these results can be found in Appendix B.1.

## 3.4 Minimizing Cost with Positive Externalities

We now explore cost minimization in congestion games. We focus on the anonymous POS model, but now the value (here, "cost") $v_i(j,k)$ of player $i$ from a resource $j$ is *non-increasing* in the number of users. For every set of players $S \subseteq N$ and any resource $j \in M$, we denote $C_j(S) = \sum_{i \in S} v_i(j,|S|)$. Our focus is the natural case were the game is cost-monotonic.[8]

[8]Cost minimization in congestion games was already studied by [6]. [6] proves that it is NP-hard to approximate the cost to within any factor. However, their proof does not hold for the POS model.

**Definition 3.6** *We say that a congestion game is* cost-monotonic*, if the total cost in each resource increases when more players are using the resource. I.e., for every resource $j$ and every subset of players $A \subseteq B$ we have $C_j(A) \leq C_j(B)$.*

Notice that cost monotonicity does not contradict the existence of positive externalities. A good example, which we use below, is in facility-location problems: assigning an additional city to a facility reduces cost from the other cities but increases the overall cost.

The following theorem uses a reduction from FACILITY-LOCATION to prove a lower bound of $O(\log n)$ for approximating the minimal cost. For cost-monotonic valuations, we present an algorithm that achieves this approximation ratio. The algorithm generalizes existing methods developed for facility-location problems [21, 18], with several subtleties that need to be handled. It is easy to show that that our model of congestion games with positive externalities is indeed more general then the existing facility-location models.

**Theorem 3.7** *In congestion games with positive externalities, an $O(\log n)$-approximation can be found in polynomial time for cost-monotonic games. On the other hand, it is NP-hard to approximate the optimal cost in a congestion game within a better factor than $O(\log n)$, even when the game is cost monotonic.*

# 4 Truthful Mechanisms for Congestion Games

While most of the literature on congestion games discussed Nash equilibria and their properties, this paper studies centralized mechanism-design issues for such environments. In such mechanisms the players send their messages to the auctioneer, and the auctioneer determines the assignment and the price that each player should pay for using the resource allocated to him. As mentioned, treating incentive-compatibility in congestion games is substantially different than in combinatorial auctions, for example, since it resembles combinatorial auctions where the items are strategic, instead of the bidders.

Our goal is to design *incentive-compatible* mechanisms, where a dominant strategy for each player is to report true private data. In other words, whatever are the values and strategies of the other players, each player will never benefit by misreporting his preferences.

## 4.1 Incentive-Compatible Approximations

The basic tool for obtaining incentive-compatible mechanisms in multi-parameter environments is the VCG payment scheme. Unfortunately, this scheme requires the exact solution of the problem in hand. In our case this is not computationally feasible, as it is NP-hard to find the optimal solution of congestion games. Recall that the VCG scheme fails to ensure truthfulness for approximately-efficient outcomes [28].

One way of overcoming this problem is to use *maximal-in-range* mechanisms. That is, mechanisms that their outcome is restricted in advance to some limited range, and they completely optimize over this restricted range of outcomes (see, e.g., [28, 10]). Incentive compatibility is immediately obtained this way by using the corresponding VCG prices, and all we have to consider is the quality of the best solution in the limited range.

Maximal-in-range mechanisms are usually obtained by reducing the problem to some special case that can be optimally solved in polynomial time. For congestion games, it is known that given the exact congestion on each resource the best assignment, with respect to these congestions, can be found in polynomial time ([6], see claim B.1). This is achieved by constructing a corresponding network and (optimally) solving the associated b-flow problem. This construction is a basic building block of our algorithm.

Our mechanism is obtained by repeatedly trying several (fixed) possibilities of the congestions on the resources. For each possibility we find the best solution under the specified congestion. The final solution of the mechanism is the best assignment of all those which we have tried. The mechanism guarantees an $O(\sqrt{m}\log n)$ approximation both for valuations as long as all the valuations are monotone (that is, the valuation of a player for every resource either monotonically increases with the number of other users, or monotonically decreases).[9]

**Input:** A congestion game with $n$ players with *monotone* valuations and $m$ resources.

**Output:** An assignment which is an $O(\sqrt{m}\log n)$-approximation to the optimal assignment.

**The Mechanism:**

1. Arbitrarily partition the resources into $\sqrt{m}$ sets, $G_1, ..., G_{\sqrt{m}}$, each consists of exactly

[9]That is, each player have negative externalities for some resources, and positive externalities for all others.

$\sqrt{m}$ resources[10].

2. For each $t = 1, 2, 4, 8..., \frac{n}{\sqrt{m}}$, and $i = 1, 2, 3, ..., \sqrt{m}$: find the maximum-value assignment in which the congestion on each resource $j \in G_i$ is exactly $t$, and the congestion on each resource $j \notin G_i$ is 0 (see Claim B.1).

3. Return the maximum-value assignment among the following set: all assignments considered in Step 2, and all assignments in which some players are assigned to the same resource, and the others, if any, are not assigned at all.

**Theorem 4.1** *The mechanism above is a polynomial-time incentive-compatible $O(\sqrt{m}\log n)$-approximation mechanism for congestion games where all the valuations are monotone.*

**Proof:** The mechanism clearly runs in polynomial time. To see that it is maximal in its range, observe that the set of the assignments that are considered is fixed and independent in the valuations of the players. Also, the chosen output is the maximum-value assignment of all those considered. Thus, the corresponding VCG prices can be used to ensure truthfulness. All that is left is to show that the algorithm indeed achieves the guaranteed approximation.

Let the optimal assignment be $(O_1, ..., O_n)$, with value $OPT$. For every real number $t$, denote by $S_t$ the set of resources with congestion between $t$ and $2 \cdot t - 1$ in the optimal allocation. Let $r = |S_t|$. Denote by $A_t$ the assignment obtained from the optimal assignment by setting the congestion on each resource $j \notin S_t$ to be $\emptyset$. Since $t$ can get $O(\log n)$ different values, there exists some $t'$ such that the value of $A_{t'}$ is at least $\Omega(\frac{OPT}{\log n})$. The value of the assignment obtained by the algorithm will be estimated with respect to $A_{t'}$.

*Case 1:* $r \leq \sqrt{m}$

We first claim that in this case, the algorithm provides an assignment which is an $O(\sqrt{m}\log n)$-approximation, as required. Let $j \in S_{t'}$ be the resource that contributes at least a fraction of $\frac{1}{r} \geq \frac{1}{\sqrt{m}}$ of the welfare in $A_{t'}$. One of the assignments the algorithm considers is the one in which $|O_j|$ players are assigned to $j$ and no player is assigned to any other resource. Clearly this assignment holds a value of at least $O(\frac{OPT}{\log n\sqrt{m}})$.

[10] If there are not enough resources to include in $G_{\sqrt{m}}$, we may add dummy resources to $G_{\sqrt{m}}$ such that the value of each player for being assigned to these items is 0, regardless of the congestion.

*Case 2:* $r > \sqrt{m}$

Notations: Let $NEG_j$ be the set of players that have negative externalities with respect to resource $j$, and $POS_j$ be the set of players with positive externalities with respect to $j$. Observe that at least one group $G_i \in \{G_1, ..., G_n\}$ contributes a $\frac{1}{\sqrt{m}}$-fraction of the welfare of $A_{t'}$ (I.e., $\Sigma_{k \in S_{t'}, j \in G_i, i} v_k(j, |O_i|) \geq \frac{|A_{t'}|}{\sqrt{m}}$, where $|A_{t'}|$ denotes the welfare obtained by $A_{t'}$). First, assume that most of the value in $G_i$ is achieved by players that has positive externalities with respect to the resources they are assigned to, i.e., $\Sigma_{j \in (S_{t'} \cap G_i),\, k \in NEG_j} v_k(j, |O_i|) \geq \Sigma_{j \in (S_{t'} \cap G_i), k \in POS_j} v_k(j, |O_i|)$ – later we will handle the other case. In Step 2 the algorithm also considers the assignment in which for each resource $j \in G_i$ the congestion is $t'$, and for each resource $j \notin G_i$ the congestion is 0. This assignment has a value of $\Omega(\frac{|A_{t'}|}{\sqrt{m}})$. To see that, set the congestion on each resource in $G_i$ with non-zero congestion (under $A_{t'}$) to $t'$ by removing from $OPT_{t'}$ enough players with positive externalities, and if there are not enough such players also remove enough players with negative externalities which their utility under congestion $t'$ is the lowest. Notice that for each resource we have kept at least half of the players with negative externalities with the "highest" values, and the congestion on each resource can only decrease, comparing to $A_{t'}$. To conclude, we lost only a constant fraction of the welfare comparing to the welfare that $G_i$ held under $A_{t'}$.

The case where most of the value is achieved by positive-externalities players, i.e., $\Sigma_{j \in (S_{t'} \cap G_i), k \in NEG} v_k(j, |O_i|) < \Sigma_{j \in (S_{t'} \cap G_i, k \in POS} v_k(j, |O_i|)$ is handled similarly. Consider again the assignment $A_{t'}$. We now describe a new assignment, with value of at least $\frac{1}{4}$ of $A_{t'}$, in which the congestion on each resource either 0 or $2t'$: set all *negative* valuations to 0 (we have lost at most half of the value), and reorder the resources in decreasing order by the contribution of each resource to the welfare. Set the congestion on resources $\frac{r}{2} + 1$ to $r$ to be $\emptyset$, and arbitrarily assign the newly de-assigned players to the first $\frac{r}{2}$ resources, if needed. Observe that this manipulation is possible since the number of players in each one of resources $\frac{r}{2} + 1, ..., r$ is at least $t'$, and the congestion on each resource $1, ..., \frac{r}{2}$ is between $t'$ and $2t'$. Also observe that the value obtained from the first $\frac{r}{2}$ resources, can only increase. If $r > \sqrt{m}$, then one of the groups $G_1, ..., G_{\sqrt{m}}$ holds at least $\frac{1}{\sqrt{m}}$ of the welfare. Using a similar argument to the previous case, we have a $O(\sqrt{m} \log n)$-approximation. □

For valuations with negative externalities, we are able to achieve an incentive-compatible

$\sqrt{n}$-approximation algorithm by adopting an algorithm by [11] for combinatorial auctions. It turns out that a modification of this algorithm remains truthful even when applied to congestion games. If we are given that all valuations exhibit negative externalities, then we can run both algorithms, to obtain a truthful $\min(O(\sqrt{m}\log n), O(\sqrt{n}))$-approximation mechanism (the VCG prices should be calculated with respect to the range of *both* algorithms).

**Proposition 4.2** *There is an incentive-compatible $O(\sqrt{n})$-approximation mechanism for congestion games in the $NEG$ model.*

**Proof:** The algorithm is an adaptation of the $O(\sqrt{m})$ truthful approximation algorithm for combinatorial auctions with sub-additive valuations of [11]. This is also a maximal-in-range algorithm, so the corresponding VCG prices should be calculated to ensure truthfulness.

1. Find the best assignment such that all bidders that are assigned to some resource are assigned to the same resource.
2. Find the best assignment such that at most one player is assigned to each resource.
3. Choose the best assignment of the two, and assign players to resources accordingly.

Again, the algorithm can be implemented by solving the associated b-flow problems (See Claim B.1). The proof of the approximation ratio is similar to the one of [11], and is omitted.

□

## 4.2 Tractable cases

Finally, we present settings that are solvable in polynomial time. Consequently, truthful mechanisms can be obtained for these cases in polynomial time using the VCG scheme. We now provide some definitions; We say that the players are *symmetric*, if they have identical valuation functions. We say that the resources are *symmetric*, if the valuation of each bidder for resource $j$ is identical to his valuation for any other resource $l$ (i.e., for every player $i$ and every $j, l \in M$ we have $v_i(j, k) = v_i(l, k)$ for every number of players $k$).

**Proposition 4.3** *The optimal welfare can be computed in polynomial time in the following settings: (1) Symmetric players (2) Symmetric resources in the POS model (3) Constant number of resources (4) $O(\log m)$ players (5) Linear valuations in the POS model*

## References


[1] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *Proc. 28th ACM Symp. on Theory of Computing*, pages 20–29, May 1996.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS 04*, pages 295–304, 2004.

[3] B. Awerbuch, Y. Azar, E. F. Grove, M. Kao, P. Krishnan, and J. S. Vitter. Load balancing in the l p norm. In *IEEE Symposium on Foundations of Computer Science*, pages 383–391, 1995.

[4] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *TARK 03*, 2003.

[5] L. Blumrosen and N. Nisan. On the computational power of iterative auctions I: demand queries. Discussion paper no. 381, The Center for the Study of Rationality, The Hebrew University., 2005. An extended abstract in EC'05 contained preliminary results.

[6] D. Chakrabarty, A. Mehta, V. Nagarajan, and V. Vazirani. Fairness and optimality in congestion games. In *EC'05*, 2005.

[7] V. Conitzer and T. Sandholm. Expressive negotiation in settings with externalities. In *AAAI 05*, 2005.

[8] V. Conitzer and T. Sandholm. Expressive negotiation over donations to charities. In *EC 04*, 2005.

[9] P. Cramton, Y. Shoham, and R. S. (Editors). *Combinatorial Auctions.* MIT Press., 2006.

[10] S. Dobzinski and N. Nisan. Approximations by computationally-efficient vcg-based mechanisms. Technical report, The Hebrew University, 2006.

[11] S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *The 37th ACM symposium on theory of computing (STOC).*, 2005.

[12] S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In SODA'06.

[13] B. Ellickson, B. Grodal, S. Scotchmer, and W. R. Zame. Clubs and the market. *Econometrica*, 67(5):1185–1219, 1999.

[14] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *FOCS 04*, pages 604–612, 2004.

[15] U. Feige. On maximizing welfare where the utility functions are subadditive. In *STOC*, 2006.

[16] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of muliticast transmissions. In *ACM Symposium on Theory of Computing*, pages 218–227, 2000.

[17] P. W. Goldberg and C. Papadimitriou. Reducability among equilibrium problems. In *STOC'06*, 2006.

[18] M. Hajiaghayi, M. Mahdian, and V. Mirrokni. Facility location problem with general cost functions. *Networks*, 42(1):42–47, 2003.

[19] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.

[20] E. Hazan, S. Safra, and O. Schwartz. On the complexity of approximating k-dimensional matching. In *RANDOM-APPROX*, 2003.

[21] D. S. Hochbaum. Heuristics for the fixed cost median problem. *Mathematical programming*, 22(2):148–162, 1982.

[22] L. Katz and C. Shapiro. Network externalities, competition, and compatibility. *American Economic Review*, 75(3):424–440, 1985.

[23] H. Konishi, M. L. Breton, and S. Weber. Pure strategy nash equilibrium in a group formation game with positive externalities. *Games and Economic Behavior*, 21(1-2):161–182, 1997.

[24] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behaviour. (A preliminary version appeared in EC'01)*, 55(2):270–296, 2006.

[25] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13:111–124, 1996.

[26] D. Monderer. Generalized congestion games. Presentation in seminar on 'Computing and Markets', Dagstuhl, Germany, January 2005.

[27] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[28] N. Nisan and A. Ronen. Computationally feasible vcg-based mechanisms. In *ACM Conference on Electronic Commerce*, 2000.

[29] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices, 2003. Forthcoming in the Journal of Economic Theory.

[30] R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[31] T. Roughgarden. *Selfish Routing and the Price of Anarchy.* MIT Press, 2005.

[32] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *STOC'06*, 2006.

# A Congestion Games and Combinatorial Auctions

## A.1 Definition of the XOS Class

This section defines the $XOS$ class of valuations for combinatorial auctions, introduced in [24]. First, a valuation $v$ is called *additive* if there are values $b_1, ..., b_m$ such that for every $S \subseteq M$ $v(S) = \Sigma_{j \in S} b_j$. We can now define $XOS$ valuations:

**Definition A.1** *A valuation $v$ for combinatorial auctions is said to be $XOS$ if there is a set of additive valuations $\{w_1, ..., w_t\}$, such that $v(S) = \max_k \{w_k(S)\}$ for all $S \subseteq M$.*

Notice that the size of the valuation might be exponential in $n$ and $m$. Therefore, we define an *XOS oracle* for a valuation $v$ to be a procedure which given a set of items $S$ returns the additive valuations that maximizes the value of $S$. That is, the oracle returns $\arg\max_k \{w_k(S)\}$.

## A.2 Missing proofs

### Proof of Proposition 2.2

**Proof:** We will prove our lower bound by reducing from the approximate disjointness problem. In this problem, there are $n$ computationally unlimited players, each player $i$ holds a string $A^i$ which specifies a subset of $\{1, ..., t\}$. The goal is to distinguish between the following two extreme cases: **(1)** $\cap_{i=1}^{n} A^i \neq \emptyset$ **(2)** $A^i \cap A^j = \emptyset$ for every $i \neq j$.

Alon et al. [1] prove that the amount of bits needed for distinguishing between the two cases is $\Omega(\frac{t}{n^4})$, also holds for randomized protocols and for non-deterministic protocols.

We will build a congestion game with $n$ players and only one resource. Each player will get a string of size $t = \binom{n}{\frac{n}{2}}$ which represents each possible set of the bidders of size $\frac{n}{2}$. We define the non anonymous valuation of the $i$'th bidder as follows:

$$v_i(S,k) = \begin{cases} 1, & \text{if } S \in A^i \text{ and } i \in S; \\ 0, & \text{otherwise.} \end{cases}$$

If $\cap_{i=1}^{n} A^i \neq \emptyset$, then it is easy to see that the maximum welfare is $\frac{n}{2}$. However, if for every $i \neq j$, $A^i \cap A^j = \emptyset$, then the maximum welfare is 1 (assigning some group $S$ of bidders such that $S \in A^i$ and $i \in S$). Since we reduced from approximate disjointness with vectors of size $t = \binom{n}{\frac{n}{2}}$, any approximation better than $\frac{n}{2}$ will require $\binom{n}{\frac{n}{2}}/n^4 = \exp(n)$ bits of communication, and the theorem follows. □

**Proof of Theorem 2.4**

1. We now show that a dual valuation $\tilde{v}$ for a resource $j$ can be represented as a maximum of additive valuations. We define an additive valuation (an "$XOS$" clause) for every subset $S \subseteq N$: let $T \subseteq S$ be the set that maximizes $\Sigma_{i \in T} v_i(j, T)$; For each item $i \in T$, define a value $b_i = v_i(j, |S|)$, and for each item $i \in S \setminus T$, define $b_i = 0$. This is an $XOS$ valuation by definition. For proving that this is indeed the dual valuation, we show that a bundle $S$ indeed receives its value from the respective additive valuation. If a higher value is received in another clause, remove the players not in $T$. Due to the negative externalities, the value obtained from players in $T$ can only increase. Thus, the additive valuation would be defined on this subset – a contradiction.

   The following example shows that the class of combinatorial auctions dual to congestion games with negative externalities is not contained in the class of combinatorial auctions with sub-modular bidders: consider a congestion game with 4 players, and one resource $j$. All bidders have the same valuation: $v(j,1) = v(j,2) = 1$, and $v(j,3) = v(j,4) = \frac{1}{2}$. In the dual combinatorial auction the value for one item is 1, the value for 2 or 3 items is 2, and the value for getting 4 items is 2.5. Clearly, this valuation does not admit decreasing marginal utilities[11], and hence it is not sub-modular.

2. We first observe that the dual combinatorial auctions of congestion games with positive externalities have substitute-free valuations: indeed, the contribution of the players in $S, T$ to the bundle $S \cup T$ will not be smaller than their contribution in the separate bundles (which are disjoint by definition).

   Next, we should show that every combinatorial auction with super-modular valuations has a dual congestion game with positive externalities. Here we will use the property of *increasing marginal values* that super-modular valuations are known to have. Fix some arbitrary ordering of the players. For each player $i$ and a set of players $S$, if $T \subseteq S$ is the set of players before player $i$ in the ordering, define $v_i(x, S) = \tilde{v}(T \cup x) - \tilde{v}(T)$. Positive externalities are now easily derived from the property of increasing marginal valuations. Claim A.3 below shows that the inclusion is strict.

   We will now show that not every $SF$ valuation can be generated from a congestion game with positive externalities.

   **Claim A.2** $SF \setminus (CG\text{-}POS) \neq \emptyset$

   **Proof:** Consider a bidder with the following SF valuation in a combinatorial auction: $\tilde{v}(\{a\}) = \tilde{v}(\{b\}) = \tilde{v}(\{c\}) = 2$, $\tilde{v}(\{ac\}) = \tilde{v}(\{ab\}) = \tilde{v}(\{bc\}) = 10$, $\tilde{v}(\{abc\}) = 12$. We claim that any congestion game that is a dual to the combinatorial auction can not admit positive externalities. Suppose such a congestion game exists. The game will have three players $\{a, b, c\}$ and a single resource $x$. Observe that if a player attaches for a two-item set value of at least 5, then, due to the positive externalities, his value for the set $abc$ will also be at least 5. Since $\tilde{v}(\{abc\}) = 12$, there must be a player (w.l.o.g $b$) where the value for any 2-item set is strictly smaller than 5. Since $\tilde{v}(\{ab\}) = 10$, it follows that $v_a(x, \{ab\}) > 5$ and due to the positive externalities we have that $v_a(x, \{abc\}) > 5$. Similarly, $v_c(x, \{abc\}) > 5$. Therefore, $v_b(x, \{abc\}) < 2$ which contradicts the positive externalities (since $v_b(x, \{b\}) = 2$). □

[11] A valuation has decreasing marginal utilities if for every subsets $A \subseteq B$ and every item $x$ we have $v(B \cup x) - v(B) \leq v(A \cup x) - v(A)$. Increasing marginal utilities are defined analogously.

**Claim A.3** $(CG\text{-}POS) \setminus SUPM \neq \emptyset$

**Proof:** Consider the following congestion game for three players $a, b, c$ and a single resource $x$. Player $a$ values the resource with 1, if he assigned to it together with at least one more player, and with 0 otherwise. Players $b$ and $c$ always value the resource with 0.

The dual is a combinatorial auction with three items: $a, b$ and $c$, and one bidder with valuation $\tilde{v}$ with a zero value of all bundles, except $\tilde{v}(\{ab\}) = \tilde{v}(\{ac\}) = \tilde{v}(\{abc\}) = 1$ This valuation does not admit increasing marginal values, since $\tilde{v}(\{ac\}) - \tilde{v}(\{a\}) = 1$, but $\tilde{v}(\{abc\}) - \tilde{v}(\{ac\}) = 0$. Therefore, $\tilde{v}$ is not super-modular. □

# B Approximating the Welfare

**Claim B.1** *[**Essentially due to [6]**] Consider a congestion game with $m$ resources and $n$ players, and let $n_1, ..., n_m$ be a vector of non-negative integers. The optimal assignment, over all the assignments in which every resource $j$ has exactly $n_j$ players assigned to it, can be computed in a polynomial time in $n$ and $m$ for every profile of players' valuations.*

**Proof:** (sketch) Construct the following bipartite graph: Let the set of all players be the vertices of one side, and the set of resources be the vertices of the other side. Add an edge $(i, j)$ for any two vertices $i \in N, j \in S$. Set the weight of the edge $(i, j)$ to be $v_i(j, t)$.

Solve the maximum flow b-matching problem with the following capacities: one to each vertex corresponding to a player, and $n_j$ to each vertex corresponding to a resource. Clearly, an solution to the flow problem implies an optimal solution to the congestion game under the specified congestions. □

**Proof of Claim 3.2**

We will first assume the value of the highest valuation is 1 (otherwise we can normalize all values accordingly). Therefore, we assume in the rest of the analysis that $OPT^* \geq 18$, otherwise allocating all players to one resource (Stage 4) achieves the desired approximation.

The value $PRE$ of the pre-assignment is a random variable with the expectation $\frac{OPT^*}{2.2}$ (since we divided the $x_{j,S}$'s by 2.2 in Stage 2). We use Chebyshev's inequality[12]:

$$\begin{aligned}
\Pr[PRE < \frac{OPT^*}{6}] &= \Pr[\frac{OPT^*}{2.2} - PRE > OPT^* \cdot (\frac{1}{2.2} - \frac{1}{6})] \\
&\leq \Pr[|PRE - \frac{OPT^*}{2.2}| \geq OPT^* \cdot (\frac{1}{2.2} - \frac{1}{6})] \\
&\leq \frac{\frac{OPT^*}{2.2}}{(\frac{1}{2.2} - \frac{1}{6})^2 (OPT^*)^2} \leq \frac{1}{2.2 \cdot (\frac{1}{2.2} - \frac{1}{6})^2 \cdot 18} \quad = \quad 0.3047
\end{aligned}$$

The expected number of total assignments of players to resources by the randomized-rounding procedure is $\frac{n}{2.2}$. Thus, with probability of at least $\frac{1}{2.2}$ this number of total assignments is at most $n$ (due to Markov's inequality). Therefore, using the union bound:

$$\Pr(C3) \quad \geq \quad 1 - \Pr(X < \frac{OPT^*}{6}) - \frac{1}{2.2} \quad \geq \quad 0.24$$

**Proposition B.2** *The integrality gap of the LP formulation is at least* 2*. This result holds even if we are in the* $POS$ *model.*

**Proof:** We show this result in a congestion game with $n$ players and $2n$ resources. Assume the players are partitioned into two disjoint sets $N_1$ and $N_2$, $|N_1| = |N_2| = \frac{n}{2}$ (we assume $n$ is even). We arbitrarily identify each set of players from $G = \{S | S = |\frac{n}{2} + 1|, (N_1 \subseteq S \text{ or } N_2 \subseteq S)\}$ with one of the resources. We define the value of a player $i$ to be 1 if he is assigned, with at least $\frac{n}{2}$ other players, to a resource $j$, such that $i$ is in the corresponding set of $j$ from $G$. Otherwise, his value is 0.

In this scenario, the value of the optimal integral solution is $\frac{n}{2} + 1$: assign exactly one set in $G$ to the corresponding resource (all other players can be assigned arbitrarily – the welfare remains the same). However, the optimal fractional solution is $n$, by setting $x_{j,g} = \frac{1}{n+1}$, if resource $j$ is identified with the set $g \in G$. Observe, that since each player appears in exactly $n + 1$ sets of $G$, this is indeed a valid solution, and the integrality gap is proven. □

[12] We use the following version of Chebyshev's inequality: Let $X$ be a sum of independent random variables, each of which lies in $[0, 1]$, and let $\mu = E[X]$. Then, for any $\alpha > 0$, $\Pr(|X - \mu| \geq \alpha) \leq \frac{mu}{\alpha^2}$. One of the reasons we normalize the values $v_i(j, k)$ to be in $[0, 1]$ is for using this inequality.

## B.1 Other Models of Congestion Games

**Proposition B.3** *In congestion games with non-anonymous valuations, it is impossible to approximate the welfare to within a factor better than an* $O(n^{\frac{1}{2}-\epsilon})$ *in the POS model, and within a factor of* $O(n^{1-\epsilon})$ *in the NEG mode, with a polynomial number of value queries, for any* $\epsilon > 0$*, unless* $P = NP$*.*

**Proof:** We prove these statements by different reductions from INDEPENDENT-SET/CLIQUE. It is known [32] that for an $n^{1-\epsilon}$-approximation for these problems cannot be achieved in polynomial time (for any fixed $\epsilon > 0$) unless $P = NP$.

In the POS model: Given a graph $G = (V, E)$, we construct the following congestion game. Let $M = V$ and $N = E$. For every vertex $i$, let $S_i$ be the set of the edges containing $i$. Each player $e = (i, j)$ will have a nonzero value for two set of players representing the edges touching both of its vertices: for every $T \supseteq S_i$, $v_e(i, T) = \frac{1}{|S_i|}$, and for every $T \supseteq S_j$, $v_e(j, T) = \frac{1}{|S_j|}$. Clearly, these valuations exhibit positive externalities. Now, the players assigned to some resource $i$ will contribute a value of 1 if and only if all the edges that touch this vertex are assigned to this vertex. However, since these edges will not be assigned to any of the neighbors of vertex $i$, then a zero value will be contributed by each one of these neighbors. Hence, it is easy to see that a solution for the congestion game is also a solution with the same value for the CLIQUE problem. If $n$ is the number of edges, there are $O(\sqrt{n})$ vertices, and the first statement follows.

In the NEG model: Given a graph $G$, it suffices to construct a congestion game with non-anonymous valuations for a single resource. Every vertex in $G$ will be a player who has a zero value for every set $S$ that contains at least one of her neighbors, and a value of 1 otherwise. Clearly, the valuations are in the NEG model. It is easy to see that a solution for the congestion game is exactly the size of the independent set formed by the relevant vertices (if any). The second statement follows. $\square$

The original formulation of congestion games [30] allowed each player to use several resources at the same time. However, there are some restrictions: each player has a list of subsets of resources, and a player must choose one such subset, and be assigned to all resources contained in this subset. We now show that even a relaxation of this requirement, in which each player can be assigned only to a certain set of resources, is hard to approximate.

**Proposition B.4** *When each player is allowed to use only a specific set of resources, there is no $n^{\frac{1}{2}-\epsilon}$-approximation, for any $\epsilon > 0$, unless $P = NP$.*

**Proof:** The proof is similar to the proof of proposition B.3, the POS model. We are reducing again from independent set. Given a graph $G = (V, E)$, we build a congestion game with $|E|$ players and $V$ resources. Each resource represents a vertex. For each edge $e \in E$, we define the following valuation: $v_e(j,t) = \frac{1}{t}$, if $n(v) = t$, and $v_e(j,t) = \frac{1}{t}$ otherwise, where $n(v)$ denotes the number of neighbors of $v$. Allow the player corresponding to an edge $(u, v)$ to take only the resources $u$ and $v$. A similar reasoning to B.3 concludes the proof. □

## B.2 Minimizing Cost in Congestion Games

**Proof of Theorem 3.7:**

We prove the lower bound by a reduction from FACILITY-LOCATION. In this problem, we are given a set $K$ of cities and a set $F$ of facilities. We denote by $f_i$ be the opening cost of facility $i$, and for every facility $i$ and city $j$ we have a connection cost $d_{ij}$ between them. The goal is to open a subset $F' \subseteq F$, such that the cost of opening $F'$ and connecting each element to some facility is minimized. Unless $P = NP$, there is no better approximation to this problem than $O(\log n)$.

We construct the following congestion game: let $N$=$K$ and $M$=$F$. The cost of each player $i \in N$ for resource $j$ when total of $k$ players are using this resource will be $d_{ij} + \frac{f_i}{k}$. The cost incurred by a set $S$ of players that are using a resource $j$ is exactly $f_j + \Sigma_{i \in S} d_{ij}$. Notice that the valuations are cost monotonic. The lower bound immediately follows from this construction.

We now turn to describe the upper bound. We first construct the following SET-COVER instance: the elements are the set of players $N$, the subsets are all the possible subsets $S \subseteq N$. The weight of each subset $S$ is defined as the minimal cost incurred for the players in $S$ on any resource, i.e., $w(S) = \min_{j \in M} C_j(S)$. We denote the resource in which the bundle $S$ receives its minimal cost by $f(S)$, i.e., $f(S) = \arg\min_{j \in M} C_i(S)$. We will see how to handle the exponential number of subsets later.

Let $S_1, ..., S_t$ be a solution for this SET-COVER problem. Observe that, due to the cost-monotonicity assumption, we can assume that the sets $S_1, ..., S_t$ are disjoint[13]. Denote the value of the optimal solution of the SET-COVER problem by $OPT_{SC}$. We will now show an approximation preserving reduction from the congestion game to the SET-COVER problem.

Let $Q_1, ..., Q_m$ be some assignment in the congestion game. First, we now show how to get a solution $T_1, .., T_n$ to the SET-COVER with a cost of at most $\Sigma_j w(Q_j)$:

1. Assign each set $Q_i$ to the resource $f(Q_i)$.
2. If more than one set was assigned to some resource $j$, reassign their union $U = \cup_{i|f(Q_i)=j} Q_i$ to $f(U)$. Let $T_j$ be the set of players assigned to resource $j$ in the end of the process.

**Observation B.5** *By the definition of the weights for the SET-COVER problem, the allocation $T_1, ..., T_m$ is feasible for the set-cover problem.*

**Observation B.6** *At each iteration of the procedure above, the total cost cannot increase.*[14]

And thus: $\Sigma_{j=1}^m C_j(Q_j) \geq \Sigma_{j=1}^m C_j(T_j) = \Sigma_{j=1}^m w(T_j)$, where the leftmost inequality holds due to Observation 1, the right most inequality holds due to Observation 2, and the equality in between is due to the definition of the weights in the SET-COVER problem.

Next, we show that that every feasible solution for the SET-COVER problem is actually an assignment for the congestion game with at most the same cost. Given a solution $S_1, ..., S_m$ to the set cover problem, let the assignment $T_1, ..., T_n$ in the congestion game be $T_j = \cup_{i|f(S_i)=j} S_i$ for $1 \leq j \leq m$. Then, $\Sigma_{j=1}^t w(S_i) = \Sigma_{j=1}^t C_{f(S_j)}(S_i) \geq \Sigma_{j=1}^m C_j(T_j)$, where the inequality holds due to the positive externalities assumption.

What is left to be proved is that the greedy algorithm for weighted SET-COVER can be run in polynomial time despite the exponential number of subsets. For implementing the algorithm, we have to find in each iteration the subset $S$ of players that minimizes $\frac{w(S)}{|S|}$. This operation may be simulated in polynomial time with the following procedure: given some

[13]Otherwise, we can remove the elements from $S_2, ..., S_t$ (in order) that have been chosen more than once, and the cost will not increase, while all the elements remain allocated.

[14]*Since $\Sigma_{j=1}^m C_j(Q_j) \geq \Sigma_{j=1}^m C_{f(Q_j)}(Q_j) \geq \Sigma_{j=1}^m C_j(\cup_{i|f(Q_i)=j} Q_i)$, where the first inequality is due to the definition of $f(j)$ and the second is due to the positive externalities.*

fixed $k$, sort (for each resource $j$) the players according to $v_i(j,k)$, and pick the $k$ players with the lowest values – this is the $k$-player set with the minimal per-player cost. Performing this computation over all $1 \leq k \leq m$ outputs the desired bundle $S$.

# C Truthful Mechanisms for Congestion Games

**Proof of Proposition 4.3**

1. The dynamic programming algorithm by [6] actually works for unrestricted player-specific valuations. Note that in the POS model, all players will use the same resource.
2. Again, assign all players to the same resource. (Since the resources are symmetric, the players in the POS model will always prefer using the most loaded resource.)
3. A proof for a similar claim was given in [6] for the NEG model, but it is easy to see that it holds for unconstrained externalities.
4. The optimal allocation in combinatorial auction where the number of items is logarithmic in the number of bidders can be found in polynomial time with the known dynamic programming algorithm. All the algorithm requires is to find value of bundles ("value queries"), and this can be done in polynomial time for player-specific congestion games.
5. We denote the value of player $i$ for the item $j$ by the function $v_i(j,k) = a_i(j) \cdot k$. Again, the optimal solution is to assign all players to the same resource: let $OPT_1, ..., OPT_m$ be an optimal assignment, and let $t = \arg\max_{j \in M} \Sigma_{i \in OPT_j} a_i(j)$. Now we have that:

$$OPT = \Sigma_{j=1}^{m} |OPT_j| \Sigma_{i \in OPT_j} a_i(j) \leq \Sigma_{j=1}^{m} |OPT_j| \Sigma_{i \in OPT_t} a_i(t) \leq n \cdot \Sigma_{i \in N} a_i(t)$$