

Resource-Aware Protocols for Network Cost-Sharing Games

George Christodoulou^{*} Vasilis Gkatzelis[†] Mohamad Latifian[‡]
Alkmini Sgouritsa[§]

Abstract

We study the extent to which decentralized cost-sharing protocols can achieve good price of anarchy (PoA) bounds in network cost-sharing games with n agents. We focus on the model of resource-aware protocols, where the designer has prior access to the network structure and can also increase the total cost of an edge (overcharging), and we study classes of games with concave or convex cost functions. We first consider concave cost functions and our main result is a cost-sharing protocol for symmetric games on directed acyclic graphs that achieves a PoA of $2 + \varepsilon$ for some arbitrary small positive ε , which improves to $1 + \varepsilon$ for games with at least two players. We also achieve a PoA of 1 for series-parallel graphs and show that no protocol can achieve a PoA better than $\Omega(\sqrt{n})$ for multicast games. We then also consider convex cost functions and prove analogous results for series-parallel networks and multicast games, as well as a lower bound of $\Omega(n)$ for the PoA on directed acyclic graphs without the use of overcharging.

1 Introduction

In this paper we study the classic model of network cost-sharing games introduced by Anshelevich et al. [2] and Chen et al. [9], where a set of n agents need to use a path that connects a designated source to a sink in a given graph G . Each edge of the graph is characterized by a cost function, and the cost of each edge needs to be distributed among the agents that use it in their chosen paths. Anshelevich et al. [2] considered cost functions that are constant, i.e., independent of the number of agents that use the edge, and assumed that the cost of each edge is shared *equally* among its users (which is known as the *Shapley cost-sharing protocol*). This gave rise to a, now, very well-studied class of games where each agent strategically chooses a path that minimizes her share of the total cost, anticipating the strategic choices of the other agents. It is well-known that the Nash equilibria of this game can have a high social cost, and a lot of subsequent work has focused on evaluating exactly how much greater this cost can be compared to the optimal social cost.

Deviating from this line of work which assumed that the cost-sharing protocol for each edge is predetermined to be the Shapley protocol, Chen et al. [9] instead approached the problem from the perspective of a designer and asked a compelling question: *can we design a better cost-sharing protocol?* In particular, if the designer of a network wishes to minimize the social cost in equilibrium, what cost-sharing protocol should she choose? Focusing on graphs with constant

^{*}University of Liverpool, UK. Email: gchristo@liverpool.ac.uk

[†]Drexel University, USA. Email: gkatz@drexel.edu. This work was supported by NSF grant CCF-1755955.

[‡]Sharif University of Technology, Iran. Email: mohamad.latifian@gmail.com. Part of this work was done when the third and fourth authors were employed by the Max-Planck-Institute for Informatics, Saarland University Campus, Saarbrücken, Germany.

[§]University of Liverpool, UK. Email: alkmini@liverpool.ac.uk. This work was partially supported by the Lise Meitner Award Fellowship.

cost functions, they proposed a list of desirable properties that the cost-sharing protocols should satisfy, they provided characterizations of these protocols, and proved tight bounds regarding their performance. Following the same research agenda, von Falkenhausen and Harks [35] provided tight bounds for more general cost functions, but for more restricted network structures, namely for scheduling games. Our work is motivated by the same question and the main differences are two-fold: *i)* we consider more general classes of cost functions than [9] and more general network structures than [35], and *ii)* we focus on *resource-aware* protocols which were recently introduced by Christodoulou and Sgouritsa [12] and studied by Christodoulou et al. [13].

The class of resource-aware protocols provide a refinement regarding the amount of *information* available to the protocol and they provide a middle-ground between the two extreme information models that were studied in the past. Note that the amount of information that is available to the protocol plays a central role. For instance, can the decisions of the protocol regarding how to distribute the cost of an edge depend on the structure of the graph G beyond that edge? Can they depend on the set of agents that are using the network at any given time? The more information the protocol has access to regarding the state of the system, the greater the ability of the designer to reach efficient outcomes. Most of the prior work focused on either *omniscient* protocols that have *full knowledge* of the instance (the structure of the graph and its cost functions as well the set of agents that are using the system), or *oblivious* protocols that have *no knowledge* of the instance, except the set of users who use the corresponding edge. Previous work has also referred to these protocols as “non-uniform” and “uniform”, respectively.

These assumptions lie at two extremes of the information spectrum. The former applies to a very static, or centralized, system where each edge always has access to up-to-date information regarding the users and the graph G . The latter, on the other hand, is very pessimistic, assuming that the cost-sharing decisions of each edge need to be oblivious to the state of the system. The resource-aware model strikes a balance between these two extremes by assuming that the protocol of each edge knows graph G and the edge cost functions, but is aware only of the set of users who chose to utilize that particular edge. This model is more realistic than the oblivious one when the network does not change often, which is the case in many real systems. In that case, the structure of the graph can be used to inform the design of the cost-sharing protocol of each edge and this protocol needs only to be updated when the structure of the network is changed. Of course, the fact that a resource-aware protocol has additional information compared to an oblivious one is only interesting if this can lead to improved outcomes. Understanding the ways in which this information can be useful is one of the central goals of this paper.

In designing cost-sharing protocols, the goal is to ensure that the worst-case equilibria of the induced game approximate the optimal social cost within some small factor. We measure the performance of these protocols using the worst-case price of anarchy (PoA) measure, i.e., the ratio of the social cost in the worst equilibrium over that in the optimal solution. Essentially, once the designer selects the protocol on each edge, then an adversary chooses the requested subset of agents so that the PoA of the induced game is maximized.

Overcharging. In accordance with some of the recent work on resource-aware protocols [13], we also consider protocols that use overcharging. This is in contrast to some of the prior work (e.g., [9],[35] and [12]) which was restricted to *budget-balanced* protocols, i.e., protocols such that the costs of the users using some edge add up to *exactly* the cost of the edge. Our protocols can instead choose to charge the users additional costs in order to dissuade them from using some of the edges. It is important to note that, once we introduce increased costs in our protocols, we compare the performance of the equilibria in the induced game with the *increased* costs to the original optimal

solution with the *initial* cost functions. Therefore, using this technique requires a careful trade-off between the benefit of improved incentives and the drawback of penalties suffered.

1.1 Our Results

In this paper we provide several upper and lower bounds for the price of anarchy (PoA) of interesting classes of network cost-sharing games. We consider *symmetric* games, where the source and sink of every agent is the same, as well as *multicast* games, where the sink is the same for every agent, but the source may be different. In terms of the graphs that we consider, we go beyond parallel link networks and consider *directed acyclic graphs* as well as *series-parallel graphs*. Finally, we consider different types of functions that describe the costs of the network edges, and we group our results based on whether these functions are *concave* or *convex*.

It is for **concave cost functions** that we get the main result of the paper (see Section 4). That is, a resource-aware cost-sharing protocol, the Never-Walk-Alone protocol, which guarantees a PoA of $2 + \varepsilon$ for symmetric games on arbitrary directed acyclic graphs and some arbitrarily small constant $\varepsilon > 0$. In fact, for instances that involve more than one agent the PoA of this protocol drops to $1 + \varepsilon$ for an arbitrarily small constant ε . We show that every game that this protocol gives rise to is guaranteed to possess a pure Nash equilibrium, and *every* such equilibrium corresponds to an optimal assignment of agents to paths. The success of this protocol in such a large class of instances crucially depends on the fact that the cost-sharing decisions depend on the structure of the graph, and we leverage this power in a novel way, using careful overcharging in order to avoid suboptimal equilibria despite the decentralized nature of the protocol. Apart from this result, we also provide an alternative protocol that achieves a PoA of 1 for all symmetric instances with series-parallel graphs and strictly concave costs, without using any overcharging; with a minor amount of overcharging, this protocol can also be adapted to get a PoA of $1 + \varepsilon$ even for weakly concave costs (Section 5). We complement these positive results with two impossibility results for the class of (non-symmetric) multicast games, for which we show that no resource-aware protocol can achieve a PoA better than $\Omega(n)$ without overcharging and $\Omega(\sqrt{n})$ even with overcharging (Section 6).

We then transition to **convex cost functions** in Section 7. We show that for symmetric games a simple protocol can achieve a PoA of 1 for series-parallel graphs, but for directed acyclic graphs no budget-balanced protocol can achieve a PoA better than $\Omega(n)$. Also, for multicast games, we show that no protocol can achieve a PoA better than $\Omega(\sqrt{n})$, even if it uses overcharging.

2 Related Work

The papers that are most closely related to ours are that of Christodoulou and Sgouritsa [12] and Christodoulou et al. [13], which were the first ones to study the design of resource-aware protocols for network cost-sharing games. Christodoulou and Sgouritsa [12] considered constant cost functions, as in Chen et al. [9], and they showed that for outerplanar graphs there exists a resource-aware protocol that performs better than any oblivious one. On the negative side, they showed that for general graphs the best protocol performs asymptotically the same with the best oblivious one. Christodoulou et al. [13] considered more general classes of cost functions beyond constant ones, as in von Falkenhausen and Harks [35], but were restricted to parallel link networks. In this paper we combine the strengths of these two papers by studying general classes of cost functions for larger classes of networks beyond parallel links.

Apart from the aforementioned papers, there are several others that also focus on the design and analysis of cost-sharing protocols. Harks and von Falkenhausen [24] applied cost-sharing protocols to capacitated facility location games and, among other results, showed that omniscient

protocols can achieve efficiency for convex cost functions and general strategy spaces. Marden and Wierman [27] studied multiple cost-sharing protocols in a model of utility maximization instead of cost-minimization. Gopalakrishnan et al. [19] provide a complete characterization of the space of oblivious cost-sharing protocols that always guarantee the existence of a pure Nash equilibrium and showed that it corresponds to the set of generalized weighted Shapley value protocols. Leveraging this characterization, Gkatzelis et al. [18] analyzed this family of cost-sharing protocols and showed that the unweighted Shapley value achieves the optimal price of anarchy guarantees for network cost-sharing games. Also, Harks and Miller [23] studied the performance of several cost-sharing protocols in a slightly modified setting, where each player declares a different demand for each resource. Christodoulou et al. [11] studied network design games with constant cost functions under the Bayesian setting, where the position of the players' sources on the graph is drawn from a distribution over all vertices. They considered overcharging, where they could use any non-budget-balanced policy under the restriction of preserving budget-balance in all equilibria. More recently Harks et al. [20] studied cost-sharing in a model that imposes some constraints over the portions of the cost that can be shared among the agents.

There is also a very close connection between network cost-sharing games and the very well-studied class of congestion games (e.g., [34], [28], [29], [3], [17], [6], [21], [22]). In particular, a congestion game can be interpreted as the game that arises from a network cost-sharing setting if the costs of the resources are divided equally among the agents. For instance, a congestion game with cost functions that are polynomials of degree d can be cast as a network cost-sharing game with costs that are polynomials of degree $d + 1$ combined with the equal sharing protocol.

The impact of cost-sharing methods on the quality of equilibria has also been studied in other models: Moulin and Shenker [33] focused on participation games, while Moulin [32] and Mosk-Aoyama and Roughgarden [30] studied queueing games. Also, very closely related in spirit is previous work on coordination mechanisms, beginning with the work of Christodoulou et al. [10] and subsequently in the papers of [25, 4, 7, 1, 26, 15, 14, 5]. Most work on coordination mechanisms concerns scheduling games and how the price of anarchy varies with the choice of local machine policies (i.e., the order in which to process jobs assigned to the same machine). Some connections between cost-sharing policies and coordination mechanisms are also provided in [8].

3 Preliminaries

We study the performance of cost-sharing protocols on classes of network cost-sharing games. A *class* of network cost-sharing games, $\mathbf{\Gamma}$, is defined by a tuple $(\mathcal{N}, \mathcal{G}, \mathcal{C}, \div)$, which comprises a universe of players \mathcal{N} , a universe of graphs \mathcal{G} , a universe of cost functions \mathcal{C} , and a cost-sharing protocol Ξ (formally defined later on). An *instance* of a cost-sharing game $\Gamma \in \mathbf{\Gamma}$ then consists of a set of players $N \subseteq \mathcal{N}$, a graph $G = (V, E) \in \mathcal{G}$ where each edge $e \in E$ is assigned a cost function c_e drawn from \mathcal{C} , and a cost-sharing protocol Ξ . Each player $i \in N$ is associated with a source s_i and a sink t_i in the graph, and she needs to use a path in G that connects them. For each edge e , the cost function $c_e(\ell)$ is a non-decreasing function that satisfies $c_e(0) = 0$, and indicates the cost of this edge when the *load* on this edge, i.e., the number of players using it, is ℓ .

The outcome of a game is a strategy profile $\mathbf{S} = (S_1, S_2, \dots, S_n)$, where S_i is the path that player i chose to connect her source and sink. For each edge $e \in E$, let $\ell_e(\mathbf{S})$ be the number of players using this edge in their path, i.e., $\ell_e(\mathbf{S}) = |\{i \in N : e \in S_i\}|$. The cost of e in this strategy profile is $c_e(\ell_e(\mathbf{S}))$, and this cost needs to be covered by the players using this edge. In this paper we design *cost-sharing protocols*, i.e., protocols that decide how the cost of each edge will be distributed among its users. A cost-sharing protocol Ξ defines at each strategy profile \mathbf{S} a cost share $\xi_{ie}(\mathbf{S})$ for each

edge $e \in E$ and every agent $i \in N$ that uses this edge in \mathbf{S} , i.e., such that $e \in S_i$. For agents i with $e \notin S_i$ we have $\xi_{ie}(\mathbf{S}) = 0$, so only the agents using an edge are responsible for its costs. We denote the total cost-share of player i in \mathbf{S} as

$$\xi_i(\mathbf{S}) = \sum_{e \in E} \xi_{ie}(\mathbf{S}).$$

Pure Nash Equilibrium (PNE). The goal of every user is to minimize her cost-share, so different cost-sharing protocols lead to different classes of games and, hence, to possibly very different outcomes. The efficiency of a game, thus, crucially depends on the choice of the protocol. In evaluating the performance of a cost-sharing protocol, we measure the quality of the pure Nash equilibria that arise in the games that it induces. A strategy vector \mathbf{S} is a *pure Nash equilibrium* (PNE) of a game Γ if for every player $i \in N$, and every other path S'_i with the same source and sink as S_i , we have

$$\xi_i(\mathbf{S}) = \xi_i(S_i, \mathbf{S}_{-i}) \leq \xi_i(S'_i, \mathbf{S}_{-i}),$$

where \mathbf{S}_{-i} denotes the vector of strategies for all players other than i . In other words, in a PNE no player can decrease her cost share by unilaterally deviating from path S_i to S'_i if all the other players' choices remain fixed. In accordance with prior work, we restrict our attention to protocols that possess at least one PNE for every $\Gamma \in \mathbf{\Gamma}$, which are called *stable* protocols.

Price of Anarchy (PoA). To evaluate the efficiency of a strategy profile \mathbf{S} , we use the total cost $C(\mathbf{S}) = \sum_{e \in E} c_e(\ell_e(\mathbf{S}))$, and we quantify the performance of the cost-sharing protocol using the price of anarchy measure. Given a cost-sharing protocol Ξ , the *price of anarchy* (PoA) of the induced class of games $\mathbf{\Gamma} = (\mathcal{N}, \mathcal{G}, \mathcal{C}, \div)$ is the worst-case ratio of equilibrium cost to optimal cost over all games in $\mathbf{\Gamma}$. That is, if $Eq(\Gamma)$ is the set of pure Nash equilibria of the game Γ , and $F(\Gamma)$ the set of all strategy profiles of Γ , then

$$\text{PoA}(\mathbf{\Gamma}) = \sup_{\Gamma \in \mathbf{\Gamma}} \frac{\max_{\mathbf{S} \in Eq(\Gamma)} C(\mathbf{S})}{\min_{\mathbf{S}^* \in F(\Gamma)} C(\mathbf{S}^*)}.$$

Budget-balance and Overcharging. We say that a cost-sharing protocol is *budget-balanced* if for every edge e and profile \mathbf{S} we have $\sum_{i \in N} \xi_{ie}(\mathbf{S}) = c_e(\ell_e(\mathbf{S}))$, i.e., the cost shares that the protocol distributes to the players using an edge adds up to exactly the cost of that edge. Apart from budget-balanced mechanisms, we also consider mechanisms that may use overcharging. These mechanisms define for each edge e a cost function \hat{c} such that $\hat{c}_e(\ell) > c_e(\ell)$ for some values of ℓ . As a result, the social cost of a given strategy profile \mathbf{S} may be increased from $C(\mathbf{S})$ to $\hat{C}(\mathbf{S}) = \sum_{e \in E} \hat{c}_e(\ell_e(\mathbf{S}))$. In these mechanisms, we measure the quality of the equilibria using the new costs, but we compare their performance to the optimal solutions based on the original cost functions:

$$\text{PoA}(\mathbf{\Gamma}) = \sup_{\Gamma \in \mathbf{\Gamma}} \frac{\max_{\mathbf{S} \in Eq(\Gamma)} \hat{C}(\mathbf{S})}{\min_{\mathbf{S}^* \in F(\Gamma)} C(\mathbf{S}^*)}.$$

Classes of Games. We evaluate the price of anarchy for several classes of games that may differ in the strategies of the players, the structure of the graph, or the cost functions used for the edges. We call a cost-sharing game *symmetric* if all the players have the same source s and sink t , and hence the same set of strategies (all the paths from s to t). Most of this paper focuses on symmetric cost-sharing games, but we also consider games where players have different sources and the same sink, known as *multicast* cost-sharing games. In terms of classes of graphs, the two main ones that we consider are *directed acyclic graphs* (DAGs) and *series-parallel graphs*. The former include all directed graphs that do not have a cycle, while the latter are defined recursively using two

simple composition operations (see Section 5.1 for a formal definition). When it comes to classes of cost functions, the two main ones that we consider are convex functions (exhibiting non-decreasing marginal costs), or concave cost functions (exhibiting non-increasing marginal costs).

Informational Assumptions. Throughout this paper we focus on the design of *resource-aware* cost-sharing protocols. The decisions of these protocols regarding how to share the cost of an edge e can depend on the set of players using the edge, the structure of the network, and the cost functions of its machines, but it cannot depend on the set of agents that are participating in the game, but are not using edge e . This captures the intuition that these cost-sharing mechanisms are decentralized and are not aware of the set of players that is using the network at any given time. On the other hand, since the structure of the network and the characteristics of its edges are not expected to change radically, these protocols can be designed in a way that leverages this information. This is in contrast to the class of *oblivious* protocols that are unaware of anything other than the set of players that use the edge or the class of *omniscient* protocols that are assumed to have access to all the information regarding the instance. In this paper we design resource-aware protocols that take advantage of the additional information that they have access to (relative to oblivious ones) in order to achieve good price of anarchy bounds in general classes of instances.

Global Ordering. Some of our mechanisms, as well as many protocols in the related work (e.g., [31, 13]), decide how to distribute the cost among the agents for a class of games $\Gamma = (\mathcal{N}, \mathcal{G}, \mathcal{C}, \div)$ by using a global ordering over the universe \mathcal{N} of players. Then, for any instance $\Gamma \in \Gamma$ from this class with a set of players $N \subseteq \mathcal{N}$, this global ordering also implies an ordering of the agents in N that are actually participating. Based on a global ordering, we use $h_e(\mathbf{S})$ to denote the highest priority agent using edge e in strategy profile \mathbf{S} , and we refer to her as the *leader* on this edge.

4 Directed Acyclic Graphs with Concave Cost Functions

In this section we consider a directed acyclic graph G with two designated vertices s and t , and concave cost functions on the edges. We study symmetric n -player games where each player strives to establish a connection from s to t . In Section 4.2 we define a family of protocols, which we call *leader-based protocols*, and in Section 4.3 we present our main result which is a leader-based protocol that uses overcharging and achieves a PoA of almost 2.

We also provide some evidence that overcharging is needed in order to achieve constant PoA, if one focuses on a subset of leader-based protocols that we call *static-share leader-based protocols*; we show (in Section 4.4) that, even for strictly concave cost functions, no such budget-balanced protocol can achieve a PoA better than $\tilde{\Omega}(\sqrt{n})$. This is in contrast to the case of parallel links where prior work has shown that such a protocol can achieve a PoA of 1 for strictly concave cost functions [13]. We leave as an open question the existence of a budget-balanced leader-based (or any resource-aware) protocol with constant PoA.

4.1 Structure of the Optimal Network

We first observe a property of the optimal solution, i.e., the assignment of users to paths that minimizes the total cost, for symmetric games. Specifically, we show that when the cost functions are concave there always exists an optimal solution where all the users are assigned to the same path connecting the source to the sink.

Lemma 1. Let G be a graph with a designated source s and sink t , and with concave cost functions on the edges. In any symmetric instance where all the users need to connect from s to t , there exists an optimal solution which uses a single path from s to t for all the users.

Lemma 1 implies that for any graph G with source s and sink t , and any number of agents n , there exists at least one path from s to t such that assigning all n players to that path minimizes the total cost. The following definition provides us with a way to consistently choose one of these paths, which will be useful in defining our protocols.

Definition 2. ($\text{OPT}(\ell)$). For a given symmetric instance, let $\text{OPT}(\ell)$ be an optimal path from s to t when the total load, i.e., the total number of agents in the system, is ℓ . When there are multiple such paths, $\text{OPT}(\ell)$ breaks ties arbitrarily but consistently.

4.2 Leader-based Protocols

In this work we consider a class of protocols that we call *leader-based protocols*. The idea behind these protocols is that for each edge we (consistently) identify a single player to be the *leader* and charge that player some load-dependent amount $\psi_e(\ell)$, while the rest of the players share the remaining cost equally. We study resource-aware protocols, so the choice of the $\psi_e(\ell)$ values may depend on the structure of the graph and the cost functions of all the edges.

Definition 3. (Leader-based Protocol). Given a class of games $\mathbf{\Gamma}$, a *leader-based protocol* uses a predetermined priority ordering π over the universe of players \mathcal{N} . For any strategy profile \mathbf{S} , the protocol identifies a *leader* on each edge as the highest priority player (according to π) using edge e , denoted as $h_e(\mathbf{S})$. The protocol further defines a value $\psi_e(\ell) \leq c_e(\ell)$ for each edge e and load ℓ . Then, for each strategy profile \mathbf{S} , the cost-share of player i for using edge e is

$$\xi_{ie}(\mathbf{S}) = \begin{cases} \psi_e(\ell_e(\mathbf{S})) & \text{if } i = h_e(\mathbf{S}) \\ \frac{c_e(\ell_e(\mathbf{S})) - \psi_e(\ell_e(\mathbf{S}))}{\ell_e(\mathbf{S}) - 1} & \text{otherwise.} \end{cases}$$

We remark that the protocol used in [13] to achieve a PoA of 1 for parallel-link graphs lies in the class of leader-based protocols, with $\psi_e(\ell_e(\mathbf{S}))$ being a fixed value when $e \in \text{OPT}(\ell_e(\mathbf{S}))$ and $\psi_e(\ell_e(\mathbf{S})) = c_e(\ell_e(\mathbf{S}))$ when $e \notin \text{OPT}(\ell_e(\mathbf{S}))$. We call these protocols *static-share leader-based protocols* and show in Section 4.4 that their PoA is $\Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$ for DAGs. Note that $\text{OPT}(\ell_e(\mathbf{S}))$ corresponds to the path that the optimal solution would use if the total load in the system was $\ell_e(\mathbf{S})$, and a resource-aware protocol has access to this information since it depends only on the graph and the edge costs; not on the *actual* load in the system, which may be different than $\ell_e(\mathbf{S})$.

4.3 An Almost Efficient Leader-based Protocol with Overcharging

In this section we design a protocol that uses overcharging and has $\text{PoA} = 2 + \varepsilon$ for an arbitrarily small constant $\varepsilon > 0$; in fact, for instances with $n > 1$ users the PoA becomes $1 + \varepsilon$. This protocol lies in the family of leader-based protocols but it is not a static-share leader-based protocol. The main idea behind this protocol is that the cost share that the leader pays for each edge decreases as the cost shares of the other players increase. The cost share of the leader is infinitesimal when he is not the only one using an edge and therefore he has an incentive to never be alone at any edge; for this reason, we call this protocol the *Never-Walk-Alone* protocol.

Our protocol uses a property of DAGs, according to which we can assign weights on the edges such that for *any* two paths with the same endpoints the sum of their edges' weights are equal. Those weights are used in order to guarantee that the charges of the leader are positive, and the equality between the weights of alternative paths is used in our main theorem in order to compare the charges on those paths.

Lemma 4. Consider a directed acyclic graph $G = (V, E)$. We can assign to each edge $e \in E$ an integer weight $w_e > 0$, and to each vertex $v \in V$ a value $\omega_v \geq 0$, such that for any two vertices u and v that are connected through a path p in G , we have $\sum_{e \in p} w_e = \omega_v - \omega_u$.

Proof. Since G is a DAG, we can topologically sort the vertices in V so that all the edges are directed from left to right. We assign to each vertex v a value ω_v equal to its position in the sorted ordering, and to each edge $e = (u, v)$ we assign the integer weight $w_e = \omega_v - \omega_u$. Since the vertices are topologically sorted, for every edge (u, v) we have $\omega_u < \omega_v$, which implies that $w_e > 0$ for every edge e . In other words, the weight of each edge corresponds to the number of vertices ahead in the topological sorting it points to. It is then easy to verify that for any two vertices, the sum of the weights of the edges for any path that connects them will be exactly their “distance” in this topological ordering. \square

Never-Walk-Alone Protocol. Let $\text{OPT}(d)$ be an optimal path from s to t when the total load is d (according to Definition 2), let $C > 2 \sum_e c_e(|\mathcal{N}|)$ be an arbitrarily large number greater than the total cost that may ever appear, and let $\epsilon > 0$ be an arbitrarily small value satisfying $\epsilon < \min_{e'} \{c_{e'}(1)\} / \sum_e w_e$, where w_e is the weight of each edge e implied by Lemma 4. The Never-Walk-Alone protocol defines a new cost function \hat{c}_e for every edge e , satisfying $\hat{c}_e(\ell) \geq c_e(\ell)$ for any load ℓ , as follows:

$$\hat{c}_e(\ell) = \begin{cases} 2c_e(1) & \text{if } \ell = 1 \\ 2(\ell - 1)c_e(\ell) + \epsilon_e(2c_e(\ell)) & \text{if } e \notin \text{OPT}(\ell) \text{ and } \ell \neq 1, \\ c_e(\ell) + \epsilon_e\left(\frac{c_e(\ell)}{\ell - 1}\right) & \text{otherwise.} \end{cases} \quad \text{where } \epsilon_e(x) = \frac{w_e C - x}{C} \epsilon,$$

Then, in order to decide how to share these new (overcharged) costs, the protocol considers a global ordering π over the universe of players, and for any strategy profile \mathbf{S} it identifies a *leader* on each edge, as the highest priority player $h_e(\mathbf{S})$ using edge e . If there are at least two players using an edge, then the cost of the leader is the second term of the costs defined in cases 2 and 3 above, and the rest of the players equally share the remaining cost, i.e., the first term. In other words, for each strategy profile \mathbf{S} , each player i using edge e is charged:

$$\xi_{ie}(\mathbf{S}) = \begin{cases} \zeta_e(\ell_e(\mathbf{S})) & \text{if } i \neq h_e(\mathbf{S}) \text{ or } \ell_e(\mathbf{S}) = 1 \\ \epsilon_e(\zeta_e(\ell_e(\mathbf{S}))) & \text{otherwise,} \end{cases},$$

$$\text{where } \zeta_e(\ell) = \begin{cases} 2c_e(\ell) & \text{if } e \notin \text{OPT}(\ell) \text{ or } \ell = 1 \\ \frac{c_e(\ell)}{\ell - 1} & \text{otherwise} \end{cases}.$$

Note that C is used to guarantee that the leader does not pay a negative cost and ϵ is used to keep that overcharging arbitrarily small¹ and to guarantee that the highest priority player has an incentive not to be alone in any edge, unless the total number of players in the system is 1, as stated in the following lemma.

Lemma 5. For every game with concave cost functions and at least two players, when we use the Never-Walk-Alone protocol there is no Nash equilibrium \mathbf{S} where the player h with the highest priority according to π , among all active agents, is the only user of some edge.

¹We remark that for any load $\ell > 1$ and any $e \in \text{OPT}(\ell)$ we can avoid even this small overcharging of $\epsilon\left(\frac{c_e(\ell)}{\ell - 1}\right)$ by charging all players but the leader $\zeta_e(\ell) = \frac{c_e(\ell) - w_e \epsilon}{\ell - 1 - \epsilon/C}$ instead of $\frac{c_e(\ell)}{\ell - 1}$ and the leader is still charged $\epsilon_e(\zeta_e(\ell_e(\mathbf{S})))$. For the sake of simplicity we keep this small overcharging.

Proof. If player h chooses a path where he is alone in some edge, then he pays at least $\min_{e'}\{c_{e'}(1)\}$. On the other hand, if this player chooses a path where he shares every edge with some agent, he is charged at most $\sum_e \epsilon_e(0) \geq \sum_e \epsilon_e(x)$, for any $x \geq 0$. Given the way we restricted ϵ to satisfy $\epsilon < \min_{e'}\{c_{e'}(1)\} / \sum_e w_e$, we have $\min_{e'}\{c_{e'}(1)\} > m\epsilon_e(0)$, so h always prefers to share. Therefore, since he always has the option of sharing, e.g., by just following the path of some other player, a profile \mathbf{S} where he is not sharing will never be a Nash equilibrium. \square

A crucial subtlety in order to prove our main theorem is to guarantee that no ties appear in the charges of alternative paths. We assume that the mechanism first applies some arbitrarily small overcharging using the following lemma (whose proof can be found in the appendix) and then uses these cost functions in the protocol definition.

Lemma 6. We can always increase the cost functions by adding arbitrarily small constants in a way that for any two vertices u, v , any two paths p_1 and p_2 from u to v , and any strategy profiles $\mathbf{S}_1, \mathbf{S}_2$, we have

$$\sum_{e \in p_1} \zeta_e(\ell_e(\mathbf{S}_1)) \neq \sum_{e \in p_2} \zeta_e(\ell_e(\mathbf{S}_2)).$$

Lemma 7. The Never-Walk-Alone protocol is stable, i.e., it always induces games that possess a pure Nash equilibrium.

Proof. We show that for any number of agents, n , the optimal strategy profile \mathbf{S}^* in which all of the load $\ell = n$ is assigned to the optimal path $\text{OPT}(\ell)$ is a pure Nash equilibrium. Consider any other path p and let p' and OPT' be any two edge-disjoint sub-paths of p and $\text{OPT}(\ell)$, respectively, with the same endpoints u and v , as in Figure 1.

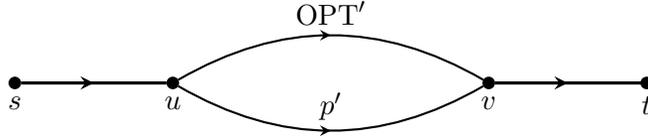


Figure 1: Network structure if players use one path.

Suppose that $\ell > 1$; the other case is trivial. If all the players use the $\text{OPT}(\ell)$ path, then the cost share of each of them (including the leader) for the $\text{OPT}'(\ell)$ portion of this path is at most

$$\sum_{e \in \text{OPT}'} \frac{c_e(\ell)}{\ell - 1} \leq \sum_{e \in p'} \frac{c_e(\ell)}{\ell - 1} \leq \sum_{e \in p'} \frac{c_e(1)\ell}{\ell - 1} \leq \sum_{e \in p'} 2c_e(1),$$

where the last term is the cost share that any player should pay if they unilaterally deviate to p' . The first inequality is due to the optimality of $\text{OPT}(\ell)$ and the second due to the concavity of the cost functions. Since this is true for any two edge-disjoint sub-paths of p and $\text{OPT}(\ell)$, we conclude that no player has an incentive to deviate to p . \square

Lemma 8. For any number of players n and any non-optimal strategy profile \mathbf{S} in which all the players use the same path p from s to t , \mathbf{S} is not a Nash equilibrium of the Never-Walk-Alone protocol.

Proof. Similar to the previous proof, let p' and OPT' be two edge-disjoint sub-paths of p and $\text{OPT}(\ell)$, respectively, with the same endpoints (Figure 1); if there are many such paths, pick the

one with total cost for load $\ell = n$ strictly greater than that of OPT' (there should be at least one, otherwise p would be an optimal path). Based on the protocol, there is at least one player paying $2c_e(\ell)$ for every edge e in p' . Note that

$$\sum_{e \in p'} 2c_e(\ell) < \sum_{e \in \text{OPT}'} 2c_e(\ell) \leq \sum_{e \in \text{OPT}'} 2c_e(1),$$

which is the cost share if a player deviates to OPT' . Hence, that player has an incentive to deviate from p' to OPT' . \square

Lemma 9. For any strategy profile \mathbf{S} in which players use at least two different paths, \mathbf{S} is not a Nash equilibrium of the Never-Walk-Alone protocol.

Proof. Aiming for a contradiction, assume that there exists some instance where such a strategy profile \mathbf{S} is a Nash equilibrium, and let a_1 denote the highest priority player in this instance. Let p_1 be the path that the highest priority player is using in \mathbf{S} , and let (v^+, u) be the first edge along p_1 such that some player, denoted as a^+ , uses edge (v^+, u) without having previously used *all* the previous edge in the p_1 path; in other words, a_1 and a^+ use paths that differ by at least one edge in getting from s to v^+ but they both use edge (v^+, u) . If no such edge exists, then let $v^+ = t$. Let p'_1 be the sub-path of p_1 from s to v^+ used by a_1 , and p^+ be the sub-path from s to v^+ used by a^+ . Since no player enters p'_1 before v^+ and due to Lemma 5, which states that a_1 is not alone in any edge under s , there must exist some other player $a_2 \neq a_1$ that uses all of the edges in p'_1 (Figure 2).

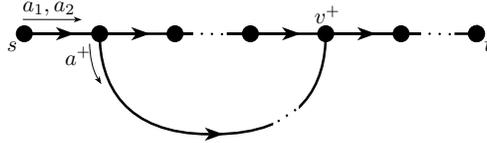


Figure 2: Network structure if players use at least two paths.

Since a_2 is not the leader in any of the edges of p'_1 (because a_1 is), he is charged $\sum_{e \in p'_1} \zeta_e(\ell_e(\mathbf{S}))$ for p'_1 . If he deviated to follow p^+ instead of p'_1 he would be charged at most $\sum_{e \in p^+} \zeta_e(\ell_e(\mathbf{S}) + 1)$. Since \mathbf{S} is a Nash equilibrium, a_2 cannot improve his cost by deviating this way, so

$$\sum_{e \in p'_1} \zeta_e(\ell_e(\mathbf{S})) < \sum_{e \in p^+} \zeta_e(\ell_e(\mathbf{S}) + 1), \quad (1)$$

where we have a strict inequality due to Lemma 6. However, since a_1 is the highest priority player among all players, his cost for p'_1 under \mathbf{S} is

$$\frac{(\omega_{v^+} - \omega_s)C - \sum_{e \in p'_1} \zeta_e(\ell_e(\mathbf{S}))}{C} \epsilon > \frac{(\omega_{v^+} - \omega_s)C - \sum_{e \in p^+} \zeta_e(\ell_e(\mathbf{S}) + 1)}{C} \epsilon,$$

where the inequality is due to (1) and the fact that $\sum_{e \in p'_1} w_e = \sum_{e \in p^+} w_e = \omega_{v^+} - \omega_s$. But, the right hand side of this inequality would correspond to the cost of a_1 if he deviates from p'_1 to p^+ . Therefore a_1 has a reason to deviate from p'_1 to p^+ which contradicts the assumption that \mathbf{S} is a Nash equilibrium. \square

Theorem 10. The PoA of the Never-Walk-Alone protocol for directed acyclic graphs with concave functions is at most $2 + \epsilon$, for an arbitrarily small value $\epsilon > 0$. If we assume $n > 1$, then $\text{PoA} = 1 + \epsilon$.

Proof. By Lemmas 7, 8 and 9 the only Nash equilibria are optimal paths so the PoA increases only due to overcharging.

Recall that in order to avoid ties on players' charges for alternative paths we increase the cost functions by arbitrarily small amounts; let ε_1 be the total such increment. Moreover, let $\varepsilon_2 = \sum_e \epsilon_e(0)$ be an upper bound on the share that the highest priority player may ever pay in any Nash equilibrium. Note that ε_2 depends on the choice of ϵ and therefore can be arbitrarily small.

For $n = 1$ the overcharging trivially results in $\text{PoA} \leq 2 + \varepsilon_1$. For $n > 1$, the total charges of all but the highest priority player sum up to the total cost occurred (after the increments for breaking ties on the cost shares) and therefore, $\text{PoA} \leq 1 + \varepsilon_1 + \varepsilon_2$. \square

4.4 Static-share Leader-based Protocols

In this section we argue that the approach of [13], which resulted in a budget-balanced protocol for parallel links with $\text{PoA} = 1$, cannot be successfully applied here. The protocol used in [13] is a leader-based protocol where, each edge essentially assumes that the set of agents currently using it are the only active users in the system. If, given that assumption, the particular edge should be used in the optimal solution, then the leader is charged some carefully chosen fixed cost, otherwise, the leader is charged the whole cost of the edge, incentivizing that agent to deviate to another edge. This is an instance of a class of protocols which we refer to as *static-share leader-based protocols*.

Definition 11. (Static-share Leader-based Protocol). A *static-share leader-based protocol* is a leader-based protocol with the following choice of $\psi_e \leq c_e(\ell)$:

$$\psi_e(\ell) = \begin{cases} \psi_e & \text{if } e \in \text{OPT}(\ell) \\ c_e(\ell) & \text{otherwise} \end{cases}$$

We next show that *no* static-share leader-based protocol can guarantee a constant PoA. In particular, the price of anarchy grows with the number of agents $n = |N|$ of the instance at hand. Our proof considers concave cost functions for simplicity, but we remark that the lower bound also holds for strictly concave cost functions as we describe in the proof. Due to limited space the proof of this theorem can be found in the appendix.

Theorem 12. Any static-share leader-based protocol has $\text{PoA} = \Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$ for directed acyclic graphs, even for the class of strictly concave cost functions.

5 Series-Parallel Graphs with Concave Cost Functions

In this section we study instances with series-parallel graphs and strictly concave cost functions and we are able to design a budget-balanced static-share leader-based protocol with $\text{PoA} = 1$, which generalizes the protocol used in [13] for parallel-link graphs. In fact, the same protocol extends to general concave cost functions results with $\text{PoA} = 1 + \epsilon$, for an arbitrarily small value $\epsilon > 0$, after applying minor overcharging on the cost functions in order to transform them to strictly concave functions, similar to [13].

5.1 Preliminaries in Series-Parallel Graphs

Here, we provide the basic definitions and properties of directed series-parallel graphs (SPGs), based on [16].

A series-parallel graph (SPG) can be constructed by performing *series* and *parallel* compositions of smaller SPGs, starting with copies of the basic SPG which is an edge. Each SPG has two designated vertices/terminals called source and sink; regarding the edge which is the basic SPG, one vertex serves as the source and the other as the sink. One can construct any SPG through a (not unique) sequence of the following operations:

- Create a new graph, consisting of a single edge directed from s to t .
- Given two SPGs C_1 and C_2 with terminals s_1, t_1, s_2, t_2 , form a new SPG C by merging s_1 and s_2 into one new terminal, s , and merging t_1 and t_2 into a new terminal, t . This is known as the *parallel composition* of C_1 and C_2 .
- Given two SPGs C_1 and C_2 , with terminals s_1, t_1, s_2, t_2 , form a new SPG C , with terminals s, t , by merging t_1 and s_2 and identifying $s = s_1$ and $t = t_2$ as the new terminals. This is known as the *series composition* of C_1 and C_2 .

For the rest of the section we will consider a SPG $G = (V, M)$, with source s and sink t , and C_1, \dots, C_r is the sequence of SPGs that are constructed in the process of generating G (with $C_r = G$). Let $\mathcal{C} = \{C_1, \dots, C_r\}$ and s_C, t_C be the terminals of any component $C \in \mathcal{C}$.

5.2 An Efficient Leader-based Protocol

Here, we study symmetric games where each edge of G has a *strictly* concave cost function, and we present a budget-balanced cost-sharing protocol with $\text{PoA} = 1$. In fact, as the following remark argues, this also implies that we can achieve a PoA of $1 + \epsilon$ for *weakly* concave cost functions for an arbitrarily small constant $\epsilon > 0$ if we use a very small amount of overcharging.

Remark 13. Using similar overcharging arguments like the one used in [13], one can transform any concave cost function into a *strictly* concave one by replacing any linear portion of the function with a concave portion that is never more than ϵ times greater. This will achieve a PoA equal to $1 + \epsilon$, where $\epsilon > 0$ is an arbitrarily small constant.

Protocol Definition. Like all leader-based protocols, this protocol uses a global order π to identify the leader $h_e(\mathbf{S})$ of each edge for each profile \mathbf{S} . After carefully assigning values $\psi_e \leq c_e(\ell_e(\mathbf{S}))$ on the edges of G , the cost share of player i for using edge e is defined as

$$\xi_{ie}(\mathbf{S}) = \begin{cases} c_e(\ell_e(\mathbf{S})) & \text{if } e \notin \text{OPT}(\ell_e(\mathbf{S})) \text{ and } i = h_e(\mathbf{S}) \\ 0 & \text{if } e \notin \text{OPT}(\ell_e(\mathbf{S})) \text{ and } i \neq h_e(\mathbf{S}) \\ \psi_e & \text{if } e \in \text{OPT}(\ell_e(\mathbf{S})) \text{ and } i = h_e(\mathbf{S}) \\ \frac{c_e(\ell_e(\mathbf{S})) - \psi_e}{\ell_e(\mathbf{S}) - 1} & \text{if } e \in \text{OPT}(\ell_e(\mathbf{S})) \text{ and } i \neq h_e(\mathbf{S}) \end{cases}$$

In order to complete the definition of the cost-sharing protocol we need to define the ψ_e 's. In fact we define a value ψ_C for each component $C \in \mathcal{C} = \{C_1, \dots, C_r\}$. This will be done iteratively starting from G and following its decomposition in reverse order.

To capture essential information of the network structure, we first define the minimum load ℓ for which some edge of component $C \in \mathcal{C}$ is used in the optimal path $\text{OPT}(\ell)$.

Definition 14. (ℓ_C^*). Given a component $C \in \mathcal{C}$ of a SPG, we define ℓ_C^* to be the minimum total load such that C is used in the optimal path $\text{OPT}(\ell_C^*)$, i.e., if E_C is the set of the edges of C , $\ell_C^* = \min\{\ell \mid E_C \cap \text{OPT}(\ell) \neq \emptyset\}$.

Next we define the minimum cost required for connecting ℓ through C .

Definition 15. (ϕ_C). We define $\phi_C(\ell)$ to be the minimum cost for establishing a path that connects a load of ℓ from the source, s_C , of C to its sink, t_C . Note that ϕ_C is a strictly concave function as the minimum of strictly concave functions². Trivially, for each edge e , $\phi_e = c_e$.

Remark 16. If C is constructed by the composition of C_1 and C_2 , then if this is a

- parallel composition, $\ell_C^* \leq \ell_{C_1}^*$, $\ell_C^* \leq \ell_{C_2}^*$, $\phi_C(\ell) \leq \phi_{C_1}(\ell)$, and $\phi_C(\ell) \leq \phi_{C_2}(\ell)$ for any ℓ .
- series composition, $\ell_C^* = \ell_{C_1}^* = \ell_{C_2}^*$ and $\phi_C(\ell) = \phi_{C_1}(\ell) + \phi_{C_2}(\ell)$ for any ℓ .

We now define ψ_C for each component $C \in \mathcal{C}$ so as to guarantee the following properties:

1. any two components that participate in a parallel composition have the same ψ value.
2. the ψ values of any two components that participate in a series composition sum up to the ψ value of the composed component.
3. $\psi_C > \frac{\phi_C(\ell)}{\ell}$ for any total load ℓ such that C is used in the optimal outcome $\text{OPT}(\ell)$.
4. $\psi_C \leq \phi_C(1)$.

The first two properties ensure that for each component C , we have $\psi_C = \sum_{e \in p} \psi_e$ for any (acyclic) path p from s_C to t_C . The third property guarantees that the highest priority player always pays more than the rest of the players. Combined with the fact that all alternative paths have the same ψ , implied by the first two properties, this allows us to show that no two paths can be used in any Nash equilibrium (Theorem 23). Finally, the fourth property is crucial in ensuring that the optimal path is a Nash equilibrium; if this inequality did not hold, the highest priority player could have an incentive to deviate to the path with unit cost $\phi_C(1)$.

Definition 17. (ψ_C). We define the ψ values iteratively, starting from C_r (recall $C_r = G$) and following its decomposition. We set ψ_{C_r} to be equal to the minimum total cost when a unit load appears in G , i.e. $\psi_{C_r} = \phi_{C_r}(1)$. Then, after defining a value ψ_C for a component C , we move on to define the ψ_{C_1} and ψ_{C_2} values of the two components C_1, C_2 whose composition led to C , as follows:

- If C is constructed by a parallel composition of C_1 and C_2 , $\psi_{C_1} = \psi_{C_2} = \psi_C$.
- If C is constructed by a series composition of C_1 and C_2 , then we consider two sub-cases:
 - if $\phi_{C_1}(1) < \psi_C \frac{\phi_{C_1}(\ell_C^*)}{\phi_C(\ell_C^*)}$, then we let $\psi_{C_1} = \phi_{C_1}(1)$ and $\psi_{C_2} = \psi_C - \psi_{C_1}$.
 - if $\phi_{C_2}(1) < \psi_C \frac{\phi_{C_2}(\ell_C^*)}{\phi_C(\ell_C^*)}$, then we let $\psi_{C_2} = \phi_{C_2}(1)$ and $\psi_{C_1} = \psi_C - \psi_{C_2}$.
 - if $\phi_{C_1}(1) \geq \psi_C \frac{\phi_{C_1}(\ell_C^*)}{\phi_C(\ell_C^*)}$ and $\phi_{C_2}(1) \geq \psi_C \frac{\phi_{C_2}(\ell_C^*)}{\phi_C(\ell_C^*)}$, then $\psi_{C_1} = \psi_C \frac{\phi_{C_1}(\ell_C^*)}{\phi_C(\ell_C^*)}$ and $\psi_{C_2} = \psi_C \frac{\phi_{C_2}(\ell_C^*)}{\phi_C(\ell_C^*)}$.

Note that the ψ_{C_1} and ψ_{C_2} values as defined above always guarantee that $\psi_{C_1} + \psi_{C_2} = \psi_C$. To see this note that $\phi_{C_1}(\ell) + \phi_{C_2}(\ell) = \phi_C(\ell)$ in a series composition (Remark 16).

In the next lemma (whose proof can be found in the appendix) we show that $\psi_C \leq \phi_C(1)$ which, as we mentioned above, is crucial for the optimal path to be a Nash equilibrium.

²The cost of each path from s to t is a concave function as it is the summation of concave functions. Due to Lemma 1, the optimal total cost for each ℓ equals the minimum cost among those paths.

Lemma 18. For any $C \in \mathcal{C}$, $\psi_C \leq \phi_C(1)$.

In the next lemma (whose proof can be found in the appendix) we show that $\psi_C > \frac{\phi_C(\ell)}{\ell}$ for any ℓ that C is used in the optimal path.

Lemma 19. For any $C \in \mathcal{C}$ and any $\ell \geq \ell_C^*$ with $\ell > 1$, $\psi_C > \frac{\phi_C(\ell)}{\ell}$; if $\ell_C^* = 1$ then $\psi_C = \phi_C(1)$.

In the next lemma (whose proof can be found in the appendix) we formally show that for each edge e , the highest priority player pays at least ψ_e and the rest of the players pay *strictly* less than ψ_e . This lemma will be used in a key argument in our main theorem (Theorem 23) in order to show that no two paths can be used in any Nash equilibrium.

Lemma 20. For any strategy profile \mathbf{S} and any edge e , $\xi_{ie}(\mathbf{S}) \geq \psi_e$ if $i = h_e(\mathbf{S})$ and $\xi_{ie}(\mathbf{S}) < \psi_e$ otherwise.

In the next two lemmas we show that the protocol belongs in the class of budget-balanced and stable protocols.

Lemma 21. This static-share leader-based cost-sharing protocol is budget-balanced.

Proof. Consider any strategy profile \mathbf{S} and any edge e . If $e \notin \text{OPT}(\ell_e(\mathbf{S}))$ then the leader, $h_e(\mathbf{S})$, covers the whole cost and the rest of the shares are 0; therefore the protocol is budget-balanced for those edges. If $e \in \text{OPT}(\ell_e(\mathbf{S}))$ and $\ell_e(\mathbf{S}) > 1$, by Lemma 18, the leader covers at most $c_e(1)$ (recall that $\phi_e = c_e$) and the rest of the players cover the rest of the cost by definition. If $e \in \text{OPT}(\ell_e(\mathbf{S}))$ and $\ell_e(\mathbf{S}) = 1$, by Lemma 19, the player is charged exactly $c_e(1)$. \square

Lemma 22. This static-share leader-based cost-sharing protocol is stable.

Proof. We show that for any load ℓ , the profile where all players choose $\text{OPT}(\ell)$ is a pure Nash equilibrium. If $\ell = 1$ then the optimal path has the minimum total cost and therefore that player has no incentive to deviate. For the rest of the proof we assume that $\ell > 1$.

Consider any alternative path p and let $C \in \mathcal{C}$ be any component composed by a parallel composition of C_1 and C_2 such that $\text{OPT}(\ell)$ uses C_1 and p uses C_2 ; let $\text{OPT}_C(\ell)$ and p_C be the subpaths of $\text{OPT}(\ell)$ and p inside C , respectively³. If any player i deviates to p , she would be alone at p_C and her cost share would be at least $\phi_C(1) \geq \psi_C$ (by Lemma 18). However, by Lemma 20, everybody pays at most $\sum_{e \in \text{OPT}_C(\ell)} \psi_e = \psi_C$. If we aggregate the cost shares over all such C (those components should be disjoint) we conclude that no player can improve their cost share by choosing p instead of $\text{OPT}(\ell)$. \square

Theorem 23. This static-share leader-based cost-sharing protocol has $\text{PoA} = 1$ for series-parallel graphs with concave cost functions.

Proof. We will show that the only possible equilibrium is an optimal path. Consider any strategy profile \mathbf{S} of total load ℓ that is not an optimum and for the sake of contradiction assume that it is a Nash equilibrium.

\mathbf{S} is a single path. If \mathbf{S} is a single path p , there should be some component $C \in \mathcal{C}$ composed by a parallel composition of C_1 and C_2 such that $\text{OPT}(\ell)$ uses C_1 and p uses C_2 ; let $\text{OPT}_C(\ell)$ and p_C be the subpaths of $\text{OPT}(\ell)$ and p inside C , respectively. If there are more than one such components, we choose one with $\sum_{e \in p_C} c_e(\ell) > \phi_C(\ell)$; there should be such a component, otherwise \mathbf{S} is optimum.

³There should be such a component with source vertex the first vertex that the two paths split and sink vertex the first vertex that they merge again.

Let h be the highest priority player overall. For every $e \in p_C$, it holds that $e \notin \text{OPT}(\ell)$ and therefore the share of h equals to $\sum_{e \in p_C} c_e(\ell) > \phi_C(\ell) = \sum_{e \in p_C^*} c_e(\ell) \geq \sum_{e \in \text{OPT}_C(\ell)} c_e(1)$. So, player h has an incentive to deviate from p_C to $\text{OPT}_C(\ell)$ and hence \mathbf{S} could not be a Nash equilibrium.

S is not a single path. Suppose now that \mathbf{S} is not a single path. There should be some component C composed by a parallel composition of C_1 and C_2 such that some nonzero load uses both C_1 and C_2 and it uses only single paths p_1 and p_2 respectively⁴. Let ℓ_1 and ℓ_2 be the loads using p_1 and p_2 respectively and let h_1 and h_2 be the highest priority players in p_1 and p_2 , respectively. W.l.o.g. suppose that $h_1 > h_2$.

By Lemma 20 h_2 is charged at least ψ_e for edge e and therefore, he is charged at least ψ_C for the whole path p_2 . If h_2 deviates to p_1 he will not be the highest priority player and, by Lemma 20, for any edge $e \in p_1$ he would be charged strictly less than ψ_e and therefore strictly less than ψ_C for the whole path p_1 . Hence, h_2 has an incentive to deviate from p_2 to p_1 , meaning that \mathbf{S} could not be a Nash equilibrium.

Overall, we ended up with a contradiction showing that \mathbf{S} cannot be a Nash equilibrium if it is not an optimum. Hence, $\text{PoA} = 1$. \square

6 Multicast Games with Concave Cost Functions

In this section, we show that our positive results for symmetric games with concave cost functions cannot be extended to the non-symmetric case of multicast network games, i.e., games where each agent i may have a different source s_i that she needs to connect to the designated sink t . In particular we show that even for constant cost functions, which are a special class of concave cost functions, resource-aware protocols cannot achieve a constant price of anarchy. Specifically the price of anarchy of any resource-aware protocol grows with the number of agents n in the instance: we first show a linear lower bound for budget-balanced resource-aware protocols, and then we extend it to protocols that allow overcharging, obtaining a lower bound of \sqrt{n} .

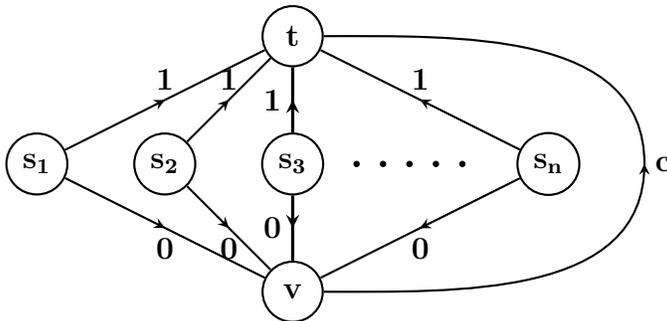


Figure 3: A simple acyclic graph with a sink t and sources s_1, s_2, \dots, s_n for each agent. All of the cost functions are constant, i.e., the cost of the edge is zero if it is not used and some constant if it is used. The direct edges from the sources to the sink have a constant cost of 1 and those from the sources to v have a constant cost of 0. The cost of the edge connecting v to the sink is some constant c , which we define appropriately for the proofs of Theorems 24 and 25.

Theorem 24. There is no resource-aware budget-balanced protocol that can achieve a PoA better than n for the case of multicast networks with constant cost functions.

⁴If for any C_1 or C_2 the load doesn't use a single path we look at the subcomponent where the structure splits until we find the required component.

The proof follows directly from the instance of Proposition 4.12 of [9] and the instance is presented in Figure 3 where c should be 1. We give the complete proof in the appendix.

Theorem 25. There is no resource-aware protocol even with overcharging that can achieve a PoA better than \sqrt{n} for the case of multicast networks with constant cost functions.

Proof. We again consider the graph of $n + 2$ vertices from Figure 3, as in the previous proof, with the only difference that the constant cost c of the (v, t) edge is set to be equal to \sqrt{n} . Note that this cost function is constant and does not depend on the number of agents using it (its cost is zero if no agents use it and equal to \sqrt{n} if at least one agent uses it). If the resource-aware protocol uses overcharging but the new costs of the (s_i, t) edges after the overcharging are no more than \sqrt{n} for all i , then we consider the instance with n agents with sources s_1, s_2, \dots, s_n and claim that the strategy profile where every one of these agent uses their (s_i, t) edge is an equilibrium: a unilateral deviation through edge (v, t) would cost at least \sqrt{n} due to the cost of (v, t) , while every agent's cost including the overcharging is at most \sqrt{n} . On the other hand, if the resource-aware protocol increases the cost of an (s_i, t) edge above \sqrt{n} for some i , then consider the alternative problem instance with the same graph, but with only one player, player i , in the system. Since the resource-aware protocol that decides the cost of edge (s_i, t) cannot tell the difference between these two instances, it would still increase the cost of that edge to above \sqrt{n} for this instance and force the agent to pay a cost of \sqrt{n} . On the other hand, in the optimal solution for that instance agent i should just use the (s_i, t) edge and pay a cost of 1. \square

7 Convex Cost Functions

For the rest of the paper, we now consider networks convex cost functions rather than concave ones. We first prove that for convex cost functions, just like in the case of concave cost functions, there exists a budget-balanced protocol with a PoA of 1 for series-parallel graphs. This protocol is the *incremental cost-sharing protocol* introduced by Moulin [31] and it is the same protocol used in [13] for parallel links. We note that this is an oblivious protocol, meaning that it requires no knowledge of the instance other the number of agents using the edge at hand.

We complement this positive result by showing that in symmetric games on directed acyclic graphs any resource-aware budget-balanced protocol has $\text{PoA} = \Omega(n)$. For protocols that use overcharging, we show that optimality cannot be achieved, by showing a lower bound of 1.18, and we leave as an open question the existence (or not) of a protocol that achieves a constant PoA with the use of overcharging. However, if we consider games beyond ones that are symmetric, such as multicast network games, we provide a lower bound of \sqrt{n} for all resource-aware protocols even with use of overcharging.

7.1 Series-Parallel Graphs

In this section we show that in series-parallel graphs (SPGs) with convex cost functions the *incremental cost-sharing protocol* proposed by Moulin [31] has $\text{PoA} = 1$. The incremental cost-sharing protocol considers a global order π of the players and defines the cost-share of each player i for using edge e to be its marginal contribution if only players preceding him in π were using e . For simplicity, for the rest of the section we name the players based on the order π .

Definition 26. (Prior Load $\ell_e^{<i}(\mathbf{S})$). Given a strategy profile \mathbf{S} and an edge e , we define the *prior load*, $\ell_e^{<i}(\mathbf{S})$, for a player i using e to be the load on e due to players preceding i according to π , i.e.,

$$\ell_e^{<i}(\mathbf{S}) = |\{k < i : e \in S_k\}|.$$

Based on the definition of the prior load we can formally define the incremental cost-sharing protocol as follows.

Definition 27. (Incremental Cost-Sharing Protocol). Given a strategy profile \mathbf{S} , the cost share of player i for using edge e is

$$\xi_{ie}(\mathbf{S}) = c_e(\ell_e^{<i}(\mathbf{S}) + 1) - c_e(\ell_e^{<i}(\mathbf{S})).$$

The following lemma is a key lemma in order to show that the incremental cost-sharing protocol has $\text{PoA} = 1$. Due to space limitations, the proof of this lemma can be found in the appendix.

Lemma 28. Given any Nash equilibrium \mathbf{S} , there exists an optimal assignment \mathbf{S}^* such that for any player i and any edge $e \in S_i^*$, it holds that $\ell_e^{<i}(\mathbf{S}) \leq \ell_e^{<i}(\mathbf{S}^*)$.

Next we give a technical lemma to be used in our main theorem.

Lemma 29. Let $\mathbf{S} = (S_1, \dots, S_n)$ be any Nash equilibrium under the incremental cost-sharing protocol and $\mathbf{S}^* = (S_1^*, \dots, S_n^*)$ be some optimal assignment satisfying Lemma 28. Then for any player i and any $e \in S_i^*$ it holds that

$$\xi_{ie}(S_i^*, \mathbf{S}_{-i}) \leq \xi_{ie}(\mathbf{S}^*).$$

Proof. By the definition of the incremental cost-sharing protocol

$$\begin{aligned} \xi_{ie}(S_i^*, \mathbf{S}_{-i}) &= c_e(\ell_e^{<i}(S_i^*, \mathbf{S}_{-i}) + 1) - c_e(\ell_e^{<i}(S_i^*, \mathbf{S}_{-i})) \\ &= c_e(\ell_e^{<i}(\mathbf{S}) + 1) - c_e(\ell_e^{<i}(\mathbf{S})) && \text{(by the definition of } \ell_e^{<i}(\mathbf{S}) \text{)} \\ &\leq c_e(\ell_e^{<i}(\mathbf{S}^*) + 1) - c_e(\ell_e^{<i}(\mathbf{S}^*)) && \text{(due to Lemma 28 and convexity)} \\ &= \xi_{ie}(\mathbf{S}^*). \end{aligned}$$

□

Theorem 30. The incremental cost-sharing protocol has $\text{PoA} = 1$ in series-parallel graphs with convex cost function.

Proof. Let $\mathbf{S} = (S_1, \dots, S_n)$ be any Nash equilibrium under this protocol and let $\mathbf{S}^* = (S_1^*, \dots, S_n^*)$ be some optimal assignment satisfying Lemma 28. Then,

$$C(\mathbf{S}) = \sum_i \sum_{e \in S_i} \xi_{ie}(\mathbf{S}) \leq \sum_i \sum_{e \in S_i^*} \xi_{ie}(S_i^*, \mathbf{S}_{-i}) \leq \sum_i \sum_{e \in S_i^*} \xi_{ie}(\mathbf{S}^*) = C(\mathbf{S}^*),$$

where the first inequality is due to the fact that \mathbf{S} is a Nash equilibrium and the second inequality comes from Lemma 29. □

7.2 Directed Acyclic Graphs

In this section we consider symmetric games on directed acyclic graphs and prove a lower-bound of $\Omega(n)$ on the PoA of all budget balanced resource-aware protocols. Then, we obtain a lower-bound of 1.18 for protocols that allow overcharging.

Theorem 31. Any stable budget-balanced resource-aware cost-sharing protocol has $\text{PoA} = \Omega(n)$ for directed acyclic graphs with convex cost functions.

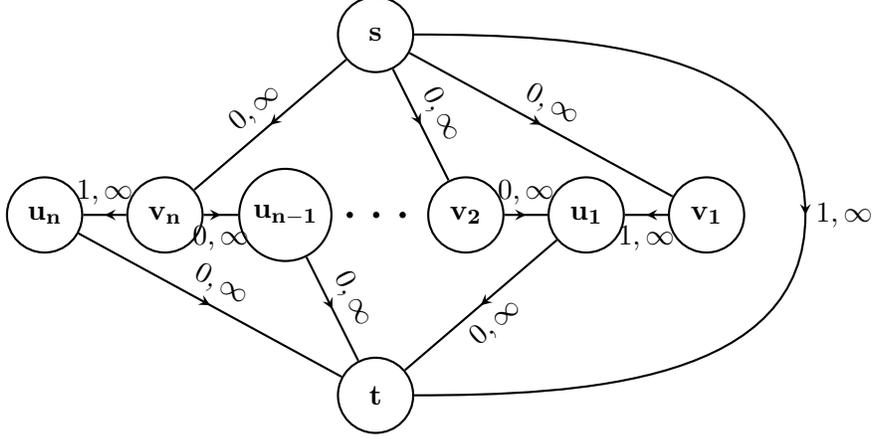


Figure 4: An acyclic graph with a single source s and sink t . All of the cost functions are infinity for load $\ell \geq 2$. This graph is used to prove Theorem 31.

Proof. We use the DAG shown in Figure 4 to prove the statement. In this graph all the players want to go from s to t and the cost of each edge is either 0 or 1 if a single player uses the edge and infinity if more than one players use it.

Considering $n + 1$ players, in the optimum, one player uses the (s, t) edge, and the others use sv_iu_it paths. The total cost of the optimum is $n + 1$. If any other strategy profile was an equilibrium it would result in an unbounded total cost and give an unbounded PoA. Therefore, there should exist an optimum strategy profile \mathbf{S}^* that is a Nash equilibrium. Assume w.l.o.g. that in \mathbf{S}^* , player $n + 1$ uses the (s, t) edge, and for each $i \leq n$, player i uses the sv_iu_it path.

If we now consider the instance in which only the first n players use the network, the strategy profile in which player i uses the sv_iu_it path should also be an equilibrium. The reason is that since \mathbf{S}^* is a Nash equilibrium, no player has an incentive to deviate to any other path but the (s, t) edge, and obviously, no player would prefer the (s, t) edge because she pays the same cost. The cost of this strategy profile is n . However, the optimum with n players is for one player to use the (s, t) edge, and the others to use the $sv_iu_{i-1}t$, for $i \geq 2$ with total cost 1. This results in $\text{PoA} = \Omega(n)$. \square

In the following theorem we give a lower-bound of 1.18 for protocols that allow overcharging. The proof can be found in the appendix.

Theorem 32. There is no stable resource-aware cost-sharing mechanism with $\text{PoA} < \frac{\sqrt{33}-1}{4} \simeq 1.18$, for directed acyclic graphs with convex cost functions even with overcharging.

7.3 Multicast Games

Here we extend our lower-bounds to multicast games. We prove a lower-bound of \sqrt{n} even for protocols that allow overcharging.

Theorem 33. There is no resource-aware protocol even with overcharging that can achieve a PoA better than \sqrt{n} for the case of multicast networks with convex cost functions.

This lower bound is inspired by the lower bound we gave in Theorem 31. The difference is that the cost of the (s, t) edge is \sqrt{n} for a single player and we use two sources s_1 and s_2 , where from s_2 there are only the alternative paths, namely (s_2, t) and $s_2v_1u_1t$ and all other paths start from

s_1 . Then the costly equilibria appeared in Theorem 31 can only be avoided by charging the player using path $\mathbf{s}_2 \mathbf{v}_1 \mathbf{u}_1 \mathbf{t}$ at least \sqrt{n} . Then if only that player appears in the system, he cannot avoid the \sqrt{n} charge whereas the original cost of the optimum was 1. We give the complete proof in the appendix.

References

- [1] Fidaa Abed and Chien-Chung Huang. Preemptive coordination mechanisms for unrelated machines. In *European Symposium on Algorithms*, pages 12–23. Springer, 2012.
- [2] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.
- [3] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *ACM Symposium on Theory of Computing*, pages 57–66. ACM, 2005.
- [4] Yossi Azar, Lisa Fleischer, Kamal Jain, Vahab S. Mirrokni, and Zoya Svitkina. Optimal coordination mechanisms for unrelated machine scheduling. *Operations Research*, 63(3):489–500, 2015.
- [5] Sayan Bhattacharya, Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Coordination mechanisms from (almost) all scheduling policies. In *5th conference on Innovations in theoretical computer science*, pages 121–134. ACM, 2014.
- [6] Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. *ACM Transactions on Economics and Computation*, 2(4):14, 2014.
- [7] Ioannis Caragiannis. Efficient coordination mechanisms for unrelated machine scheduling. *Algorithmica*, 66(3):512–540, 2013.
- [8] Ioannis Caragiannis, Vasilis Gkatzelis, and Cosimo Vinci. Coordination mechanisms, cost-sharing, and approximation algorithms for scheduling. In Nikhil R. Devanur and Pinyan Lu, editors, *Web and Internet Economics - 13th International Conference, WINE 2017, Bangalore, India, December 17-20, 2017, Proceedings*, volume 10660 of *Lecture Notes in Computer Science*, pages 74–87. Springer, 2017.
- [9] Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. Designing network protocols for good equilibria. *SIAM Journal on Computing*, 39(5):1799–1832, 2010.
- [10] George Christodoulou, Elias Koutsoupias, and Akash Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.
- [11] George Christodoulou, Stefano Leonardi, and Alkmini Sgouritsa. Designing cost-sharing methods for bayesian games. In *Algorithmic Game Theory - 9th International Symposium, SAGT 2016*, pages 327–339, 2016.
- [12] George Christodoulou and Alkmini Sgouritsa. Designing networks with good equilibria under uncertainty. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 72–89, 2016.

- [13] Giorgos Christodoulou, Vasilis Gkatzelis, and Alkmini Sgouritsa. Cost-sharing methods for scheduling games under uncertainty. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 441–458, 2017.
- [14] Giorgos Christodoulou, Kurt Mehlhorn, and Evangelia Pyrga. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica*, 69(3):619–640, 2014.
- [15] Richard Cole, José R. Correa, Vasilis Gkatzelis, Vahab S. Mirrokni, and Neil Olver. Decentralized utilitarian mechanisms for scheduling games. *Games and Economic Behavior*, 92:306–326, 2015.
- [16] David Eppstein. Parallel recognition of series-parallel graphs. *Inf. Comput.*, 98(1):41–55, 1992.
- [17] Martin Gairing and Florian Schoppmann. Total latency in singleton congestion games. In *Internet and Network Economics*, pages 381–387. Springer, 2007.
- [18] Vasilis Gkatzelis, Konstantinos Kollias, and Tim Roughgarden. Optimal cost-sharing in general resource selection games. *Operations Research*, 64(6):1230–1238, 2016.
- [19] Ragavendran Gopalakrishnan, Jason R. Marden, and Adam Wierman. Potential games are necessary to ensure pure nash equilibria in cost sharing games. *Mathematics of Operations Research*, 2014.
- [20] Tobias Harks, Martin Hoefer, Anja Huber, and Manuel Surek. Efficient black-box reductions for separable cost sharing. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, volume 107 of *LIPICs*, pages 154:1–154:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [21] Tobias Harks and Max Klimm. On the existence of pure nash equilibria in weighted congestion games. *Mathematics of Operations Research*, 37(3):419–436, 2012.
- [22] Tobias Harks, Max Klimm, and Rolf H. Möhring. Characterizing the existence of potential functions in weighted congestion games. *Theory of Computing Systems*, 49(1):46–70, 2011.
- [23] Tobias Harks and Konstantin Miller. The worst-case efficiency of cost sharing methods in resource allocation games. *Operations Research*, 59(6):1491–1503, 2011.
- [24] Tobias Harks and Philipp von Falkenhausen. Optimal cost sharing for capacitated facility location games. *Eur. J. Oper. Res.*, 239(1):187–198, 2014.
- [25] Nicole Immorlica, Li Erran Li, Vahab S. Mirrokni, and Andreas S. Schulz. Coordination mechanisms for selfish scheduling. *Theoretical Computer Science*, 410(17):1589–1598, 2009.
- [26] Konstantinos Kollias. Nonpreemptive coordination mechanisms for identical machines. *Theory of Computing Systems*, 53(3):424–440, 2013.
- [27] Jason R. Marden and Adam Wierman. Distributed welfare games. *Operations Research*, 61(1):155–168, 2013.
- [28] Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

- [29] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [30] Damon Mosk-Aoyama and Tim Roughgarden. Worst-case efficiency analysis of queueing disciplines. In *International Colloquium on Automata, Languages and Programming*, pages 546–557. Springer, 2009.
- [31] Hervé Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.
- [32] Hervé Moulin. The price of anarchy of serial, average and incremental cost sharing. *Economic Theory*, 36(3):379–405, 2008.
- [33] Hervé Moulin and Scott J. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
- [34] Robert W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.
- [35] Philipp von Falkenhausen and Tobias Harks. Optimal cost sharing protocols for scheduling games. *Mathematics of Operations Research*, 38(1):184–208, 2013.

A Proofs Deferred from the Main Body of the Paper

A.1 Proof of Lemma 1

Proof. Consider any optimum allocation with at least two paths p_1, p_2 from s to t ; if there is no such allocation then all optimal allocations use single paths and the lemma follows. We next show that shifting one player either from p_1 to p_2 or the other way around does not increase the total cost.

Let ℓ_e be the load at each edge e under the assumed optimum allocation. Suppose w.l.o.g. that

$$\sum_{e \in p_1} (c_e(\ell_e) - c_e(\ell_e - 1)) \geq \sum_{e \in p_2} (c_e(\ell_e) - c_e(\ell_e - 1)). \quad (2)$$

By cancelling out common edges and using the concavity of the cost functions we get,

$$\sum_{e \in p_1 \setminus p_2} (c_e(\ell_e) - c_e(\ell_e - 1)) \geq \sum_{e \in p_2 \setminus p_1} (c_e(\ell_e + 1) - c_e(\ell_e)). \quad (3)$$

Next we will show that shifting one player from p_1 to p_2 will not increase the total cost of the edges belonging to $p_1 \cup p_2$, hence the overall cost will not increase since the load on the rest of the edges remains intact. The total cost of the edges in $p_1 \cup p_2$ after a player shifts from p_1 to p_2 is equal to

$$\begin{aligned} & \sum_{e \in p_1 \setminus p_2} c_e(\ell_e - 1) + \sum_{e \in p_1 \cap p_2} c_e(\ell_e) + \sum_{e \in p_2 \setminus p_1} c_e(\ell_e + 1) \\ \leq & \sum_{e \in p_1 \setminus p_2} c_e(\ell_e) + \sum_{e \in p_2 \setminus p_1} c_e(\ell_e) + \sum_{e \in p_1 \cap p_2} c_e(\ell_e) = \sum_{e \in p_1 \cup p_2} c_e(\ell_e), \quad (\text{by using (3)}) \end{aligned}$$

where $\sum_{e \in p_1 \cup p_2} c_e(\ell_e)$ is the current total cost of the edges in $p_1 \cup p_2$. Next we need to show that (2) still holds for the new allocation. This implies that if all the players are shifted from p_1 to p_2 the cost will not increase and therefore p_2 would be an optimal allocation.

Let ℓ'_e be the new load on edge e after shifting a player from p_1 to p_2 . Obviously, $\ell'_e \leq \ell_e$ for any $e \in p_1$ and $\ell'_e \geq \ell_e$ for any $e \in p_2$. Then, due to concavity,

$$\sum_{e \in p_1} (c_e(\ell'_e) - c_e(\ell'_e - 1)) \geq \sum_{e \in p_1} (c_e(\ell_e) - c_e(\ell_e - 1)) \geq \sum_{e \in p_2} (c_e(\ell_e) - c_e(\ell_e - 1)) \geq \sum_{e \in p_2} (c_e(\ell'_e) - c_e(\ell'_e - 1)).$$

Therefore, (2) holds for the new allocation and the lemma follows. \square

A.2 Proof of Lemma 6

Proof. Let e_1, \dots, e_k be all the edges and r be an arbitrarily big constant. At first, we round up all the costs to r decimal places, i.e., for each edge e and load ℓ , we round up $c_e(\ell)$ to the closest multiple of 10^{-r} . Then we increase $c_{e_i}(\ell)$ by $\mathcal{N}! \infty^{r - (\nabla + \mathcal{N}!)}$, where \mathcal{N} is the maximum number of players ever appears. After these increments, for any e_i , any cost share $\zeta_e(\ell)$ has some non-zero values at decimal places from $r + \mathcal{N}!(\cdot - \infty) + \infty$ to $r + \mathcal{N}!$ and the rest of the decimal places greater than r have zero values. Since each edge appears at most once in a path, the cost shares of two different paths are equal if and only if they have the same set of edges.

Note that the increment at each edge is upper bounded by 10^r and therefore can be arbitrarily small as it only depends on the choice of r . \square

A.3 Proof of Theorem 12

Proof. In order to show this lower bound we construct a graph with appropriate edge cost functions. Consider the graph of Figure 5 which is basically a Braess's network with edge (s, u) replaced by multiple edges e_1, \dots, e_r . The edges (v, u) and (u, t) have constant cost of 0 and $2k$, respectively, for some integer $k \geq 6$. The cost of the edge (v, t) equals the load on that edge. For the rest of the edges we define the cost functions as follows:

$$c_0(\ell) = \begin{cases} \ell & \text{if } \ell \leq k \\ k & \text{otherwise} \end{cases}, \quad c_j(\ell) = \frac{\ell}{jk^2} + \frac{H_{j-1}}{k} + \varepsilon_j, \quad \text{for } 1 \leq j \leq r,$$

where H_j is the harmonic number, i.e. $H_j = \sum_{i=1}^j 1/i$, for $j \geq 1$ and $H_0 = 0$, and $0 < \varepsilon_1 < \varepsilon_2 < \dots < \varepsilon_r < \frac{1}{r^2 k^2}$ are used to guarantee a single optimum for each ℓ . r should be some integer such that $c_r(1) < 1$; we set $r = 2^k$ that satisfies this inequality if someone notice that $H_{2^{k-1}} \leq k - 1$ for $k \geq 6$.⁵

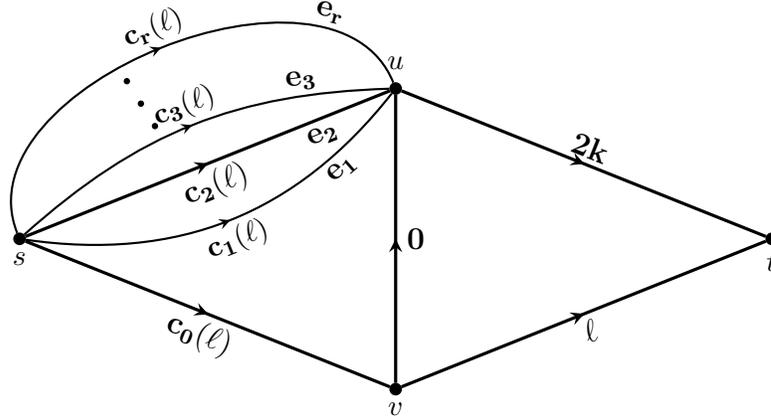


Figure 5: Graph where $\text{PoA} = \Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$ for all static-share leader-based protocol.

For simplicity we denote by sv , vu , ut and vt , the edges (s, v) , (v, u) , (u, t) and (v, t) . Moreover, by svt we denote the path using edges sv and vt and similarly for the rest of the paths.

First note that, for $\ell \leq k$, $c_0(\ell) + \ell \leq 2k$ meaning that the only optimal path for that load is the svt . It is not hard to verify that

$$c_0(\ell) + \ell > c_1(\ell) + 2k, \quad \text{for } \ell \geq k + 1, \quad (4)$$

$$c_j(\ell) < c_{j+1}(\ell), \quad \text{for } \ell \leq (j + 1)k, \quad (5)$$

$$c_{j-1}(\ell) > c_j(\ell), \quad \text{for } \ell \geq jk + 1, \quad (6)$$

for all $j \geq 1$. Inequality (4) indicates that for $\ell \geq k + 1$ the path svt cannot be the optimum. Inequalities (5) and (6) indicate that among all edges connecting s to v , e_j has the minimum cost for $jk + 1 \leq \ell \leq (j + 1)k$. Let ℓ^* be the minimum load such that $c_r(\ell^*) > k$ meaning that for $\ell \geq \ell^*$ the only optimal path is the $svut$; note that $\ell^* > rk + 1$. In Table 1 we summarize the optimal path for each load.

⁵We remark that our result holds for strictly concave functions as well. To see this, one can add $\epsilon - \epsilon/\ell$ to all cost functions and verify that all arguments still hold for sufficiently small $\epsilon > 0$.

ℓ	OPT(ℓ)
$[1, k]$	svt
$[jk + 1, (j + 1)k], 1 \leq j < r$	sut through e_j
$[rk + 1, \ell^* - 1]$	sut through e_r
$[\ell^*, \infty]$	$svut$

Table 1: Optimal paths for the network of Figure 5.

Consider now $n = r^2k^2$ players and any Nash equilibrium \mathbf{S} of any static-share leader-based protocol. We show that $\text{PoA} = \Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$.

First note that if more than $n/2$ players use some of the edges connecting s to u , due to concavity the minimum cost occurred if all those players use the same edge (Lemma 1) and due to inequality (6) this edge is e_r . Therefore, the total cost is at least $c_r(n/2) > r/2$. It is easy to check that $\ell^* \leq rk^2 < n$ (recall that ℓ^* is the minimum load such that $c_r(\ell^*) > k$) and hence the optimum is for all players to use the path $svut$ with total cost $3k$. So, $\text{PoA} > \frac{r}{6k} = \Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$.⁶

Now for the rest of the proof suppose that at least $n/2$ players use sv under \mathbf{S} . We distinguish between three cases based on which path the highest priority player overall, h , uses in \mathbf{S} .

h uses svt : If any edge connecting s to u is used under \mathbf{S} , definitely one of its users is charged by at least $1/rk^2$. If that player deviates from that edge to svu , he would pay at most $2/r^2k$ as he wouldn't be the highest priority player in sv ; h is. But, $2/r^2k < 1/rk^2$, for $k \geq 6$, meaning that \mathbf{S} couldn't be a Nash equilibrium.

Therefore, only the paths $svut$ and svt can be used in \mathbf{S} . If $\ell_{ut}(\mathbf{S}) \leq k$, then $\ell_{vt}(\mathbf{S}) \geq n - k$ resulting in $\text{PoA} > n/3k = \Omega(n/\log n)$. So we consider at last the case that $\ell_{ut}(\mathbf{S}) > k$ meaning that for any $\ell \geq \ell_{ut}(\mathbf{S})$, $ut \in \text{OPT}(\ell)$. First note that any player but h using vt pays either 1 or 0. The reason is that since $vt \in \text{OPT}(1)$, the ψ_{vt} defined by the protocol should be 1 so that it is budget-balanced. Similarly, $\psi_{sv} = 1$. This further means that h is charged at least 2 for using svt . We take two cases on the value of ψ_{ut} .

- $\psi_{ut} > 1$: player h_{ut} pays more than 1 for the path vut and if he deviates from vut to vt he would pay at most 1, as we mentioned above.
- $\psi_{ut} \leq 1$: player h currently pays at least 2 and if he deviates from svt to sut he would pay strictly less than 2, since the unit cost of any edge connecting s to u is less than 1 and the charge for using ut would be ψ_{ut} since $ut \in \text{OPT}(\ell_{e_{ut}}(\mathbf{S}) + 1)$.

Hence, \mathbf{S} cannot be a Nash equilibrium in this case. Overall, if h uses svt , $\text{PoA} = \Omega(n/\log n)$.

h uses sut : Similarly as above, if $\ell_{ut}(\mathbf{S}) \leq 2k + 1$, $\text{PoA} = \Omega(n/\log n)$, so we assume that $\ell_{ut}(\mathbf{S}) \geq 2k + 2$. If vt is used under \mathbf{S} , then h_{vt} pays at least 1 for vt whereas if he deviates from vt to vut he would pay less than 1 since he is not the highest priority player, h is. Therefore, vt cannot be used in \mathbf{S} , meaning that h_{sv} follows the path $svut$. It is easy to argue now that every edge connecting s to u is used under \mathbf{S} , otherwise h_{sv} would have an incentive to deviate to the empty edge. This is because h_{sv} pays at least 1 for sv and the unit cost of any su edge is less than 1. By Claim 34 $\text{PoA} = \Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$.

⁶This is true because $\frac{\sqrt{n}}{\log^2 n} = \frac{rk}{(\log 2^{2k}k^2)^2} < \frac{rk}{(\log 2^k)^2} = \frac{r}{k}$.

h uses svut : In this case we can argue right away that every su edge is used under \mathbf{S} , otherwise h would have an incentive to deviate to the empty edge. Then again by Claim 34 $\text{PoA} = \Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$.

Claim 34. If $n = r^2 k^2$ players appear and every su edge is used in some Nash equilibrium \mathbf{S} , then $\text{PoA} = \Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$.

Proof. We first argue that each edge e_j , for $1 \leq j \leq r$, is used by at least jk players, apart maybe from one. To show this we will prove that there are no two edges $e_j, e_{j'}$ such that both $e_j \notin \text{OPT}(\ell_{e_j}(\mathbf{S}) + 1)$ and $e_{j'} \notin \text{OPT}(\ell_{e_{j'}}(\mathbf{S}) + 1)$. Suppose on the contrary that there exist two such edges and w.l.o.g. let $h_{e_j} > h_{e_{j'}}$. Player $h_{e_{j'}}$ pays for $e_{j'}$ at least $\psi_{e_{j'}} > 0$ and if he deviates to e_j he would pay zero for that edge as he would not be the highest priority player in e_j and $e_j \notin \text{OPT}(\ell_{e_j}(\mathbf{S}) + 1)$. Therefore, since \mathbf{S} is a Nash equilibrium, all but one edges e_j should be used by at least jk players.

Now it is easy to compute that the total cost of \mathbf{S} is at least

$$\sum_{j=1}^{r-1} c_j(jk) \geq \frac{1}{k} \sum_{j=0}^{r-2} H_j = \frac{1}{\log r} ((r-1)H_{r-2} - (r-2)) = \Theta(r).^7$$

The optimal total cost is $3k$, which gives $\text{PoA} = \Omega\left(\frac{\sqrt{n}}{\log^2 n}\right)$. □

□

A.4 Proof of Lemma 18

Proof. We are going to prove this lemma by induction starting from G . For G the statement holds by definition.

Suppose that $\psi_C \leq \phi_C(1)$ for some $C \in \mathcal{C}$. We will show that the statement also holds for its components C_1, C_2 as well.

- If C is constructed by a parallel composition, for both $i \in \{1, 2\}$, it holds by Remark 16 that

$$\psi_{C_i} = \psi_C \leq \phi_C(1) \leq \phi_{C_i}(1).$$

- If C is constructed by a series composition, if the first case holds in the definition of ψ_{C_i} , then the statement trivially holds. If the second case holds then

$$\psi_{C_{i+1}} = \psi_C - \psi_{C_i} \leq \phi_C(1) - \phi_{C_i}(1) = \phi_{C_{i+1}}(1),$$

where the inequality holds by assumption and the last equality is due to Remark 16.

□

⁷We use the property $\sum_{j=1}^n H_j = (n+1)H_n - n$.

A.5 Proof of Lemma 19

Proof. We also prove this lemma by induction starting from G . For G by definition $\psi_{C_r} = \phi_{C_r}(1)$ and since ϕ_{C_r} is a strictly concave function $\psi_{C_r} > \frac{\phi_{C_r}(\ell)}{\ell}$ for any $\ell > 1$.

Suppose that the statement holds for some $C \in \mathcal{C}$. We show that the statement also holds for its components C_1, C_2 as well.

- If C is constructed by a parallel composition, for any $i \in \{1, 2\}$, we consider two cases:
 - $\ell_{C_i}^* = 1$: by Remark 16, $\ell_C^* = 1$ and further $\phi_C(1) = \phi_{C_i}(1)$. By the definition of ψ_{C_i} and the assumption that C satisfies the lemma's statement,

$$\psi_{C_i} = \psi_C = \phi_C(1) = \phi_{C_i}(1).$$

- $\ell_{C_i}^* > 1$: for any $\ell \geq \ell_{C_i}^*$,

$$\psi_{C_i} = \psi_C > \frac{\phi_C(\ell_{C_i}^*)}{\ell_{C_i}^*} = \frac{\phi_{C_i}(\ell_{C_i}^*)}{\ell_{C_i}^*} \geq \frac{\phi_{C_i}(\ell)}{\ell},$$

where the first inequality is because C satisfies the lemma's statement for $\ell_{C_i}^* \geq \ell_C^*$ and $\ell_{C_i}^* > 1$ ($\ell_{C_i}^* \geq \ell_C^*$ is true due to Remark 16). For the second equality note that if $\phi_C(\ell_{C_i}^*) < \phi_{C_i}(\ell_{C_i}^*)$ the optimal path wouldn't go through C_i , which contradicts the definition of $\ell_{C_i}^*$.

- If C is constructed by a series composition, then for any $i \in \{1, 2\}$ we consider again two cases:
 - $\ell_C^* = 1$: it holds by assumption that $\psi_C = \phi_C(1)$ and by Remark 16 that $\ell_{C_i}^* = 1$. Therefore, by definition, $\psi_{C_i} = \psi_C \frac{\phi_{C_i}(1)}{\phi_C(1)} = \phi_{C_i}(1)$, which is what is required for $\ell_{C_i}^* = 1$.
 - $\ell_C^* > 1$: here it is either $\psi_{C_i} = \phi_{C_i}(1) > \frac{\phi_{C_i}(\ell)}{\ell}$ for any $\ell > 1$, or $\psi_{C_i} \geq \psi_C \frac{\phi_{C_i}(\ell_C^*)}{\phi_C(\ell_C^*)}$. In the second case:

$$\psi_{C_i} \geq \psi_C \frac{\phi_{C_i}(\ell_C^*)}{\phi_C(\ell_C^*)} > \frac{\phi_C(\ell_C^*)}{\ell_C^*} \frac{\phi_{C_i}(\ell_C^*)}{\phi_C(\ell_C^*)} = \frac{\phi_{C_i}(\ell_C^*)}{\ell_C^*} \geq \frac{\phi_{C_i}(\ell)}{\ell}.$$

□

A.6 Proof of Lemma 20

Proof. We distinguish between two cases:

- $e \notin \text{OPT}(\ell_e(\mathbf{S}))$: for $i = h_e(\mathbf{S})$, $\xi_{ie}(\mathbf{S}) = c_e(\ell_e(\mathbf{S})) \geq c_e(1) = \phi_e(1) \geq \psi_e$ by Lemma 18. The rest of the players are charged 0.
- $e \in \text{OPT}(\ell_e(\mathbf{S}))$: for $i = h_e(\mathbf{S})$ then $\xi_{ie}(\mathbf{S}) = \psi_e$ by definition. If it exists $i \neq h_e(s^*)$ then $\ell_e(\mathbf{S}) > 1$. Additionally, since $e \in \text{OPT}(\ell_e(\mathbf{S}))$, $\ell_e(\mathbf{S}) \geq \ell_e^*$. By using Lemma 19 twice,

$$\xi_{ie}(s^*) = \frac{c_e(\ell_e(\mathbf{S})) - \psi_e}{\ell_e(\mathbf{S}) - 1} < \frac{c_e(\ell_e(\mathbf{S})) - \frac{c_e(\ell_e(\mathbf{S}))}{\ell_e(\mathbf{S})}}{\ell_e(\mathbf{S}) - 1} = \frac{c_e(\ell_e(\mathbf{S}))}{\ell_e(\mathbf{S})} < \psi_e.$$

□

A.7 Proof of Theorem 24

Proof. This follows directly from the instance of Proposition 4.12 of [9] and we include it only for completeness. Consider an instance with n players, each one with a source s_i and a common sink t as in Figure 3, and the constant edge costs are as they appear in the figure with the cost of the (v, t) edge being equal to 1. It is easy to verify that the strategy profile where every player i uses the direct edge (s_i, t) to get to the sink is an equilibrium for any budget-balanced resource-aware protocol. In particular, since there is only a single player that uses each (s_i, t) edge, every budget-balanced protocol will need to charge each of these agents a cost of 1. Any unilateral deviation would also cost them 1, however, due to the cost of (v, t) . In the optimal solution all agents share the (v, t) edge and the social cost is 1 instead of n . \square

A.8 Proof of Lemma 28

Proof. Starting from any optimum \mathbf{S}^* we swap subpaths of carefully selected players and end up to some optimum satisfying the lemma's statement. The swaps appear in rounds in each component starting from G and following its decomposition (C_r, \dots, C_1) (recall that $G = C_r$). We say that a strategy S_i passes through component C if it uses some of its edges.

For each component C (starting from C_r and reducing the index) we reallocate the optimum inside C by using the following reallocation procedure:

- If C is constructed by a series composition, we do not do any reallocation and proceed to the next component.
- If C is constructed by a parallel composition of C_1 and C_2 , for players i from 1 to n , if
 - both S_i and S_i^* pass through C ,
 - w.l.o.g. S_i passes through C_1 and S_i^* passes through C_2 and
 - there exists $j > i$ such that S_j^* passes through C_1

switch the subpaths of S_i^* and S_j^* inside C

Note that after the above procedure, the new \mathbf{S}^* has the same total cost with the initial profile and therefore is still an optimum. For the final optimum \mathbf{S}^* (after considering all components) we prove the following claim.

Claim 35. For any player j and any component C that S_j^* passes through, it holds that for any player $i < j$, if S_i passes through C so does S_i^* .

Proof. We prove the claim by induction starting from G and following its decomposition. The claim is trivially satisfied for G .

Suppose that component C satisfies the claim; we will show that also C_1 and C_2 that compose C satisfy the claim. We distinguish between the two possible cases of how C was composed.

C is constructed by a series composition of C_1, C_2 . It is easy to see that any path passes through C should also pass through both C_1 and C_2 and therefore both C_1, C_2 satisfy the statement of the claim.

C is constructed by a parallel composition of C_1, C_2 . For the sake of contradiction assume that C_1 doesn't satisfy the statement and so there exists j and i with $i < j$ and such that S_i and S_j^* pass through C_1 but not S_i^* ; let i be the minimum such index. Since C satisfies the claim's statement, S_i^* should pass through C and therefore through C_2 . Note though that while processing i during the round considering C in the above reallocation procedure, there shouldn't exist any S_j^*

with $i < j'$ passes through C_1 otherwise S_i^* should have been swapped with that; therefore there shouldn't exist S_j^* . However, since we assumed that S_j^* passes through C_1 ultimately, it should have been swapped afterwards with some other optimal subpath, but this proves the existence of some S_j^* , passing through C_1 with $i < j' < j$ while processing i .⁸ This is a contradiction meaning that C_1 should satisfy the statement as well. \square

Claim 35 holds for every edge $e \in S_i^*$ and trivially we get $\ell_e^{<i}(\mathbf{S}) \leq \ell_e^{<i}(\mathbf{S}^*)$. \square

A.9 Proof of Theorem 32

Proof. We use the DAG shown in Figure 6 to prove the statement. In this graph all the players want to go from s to t and the cost of each edge is as indicated in the figure if a single player uses the edge and infinity if more than one players use it.

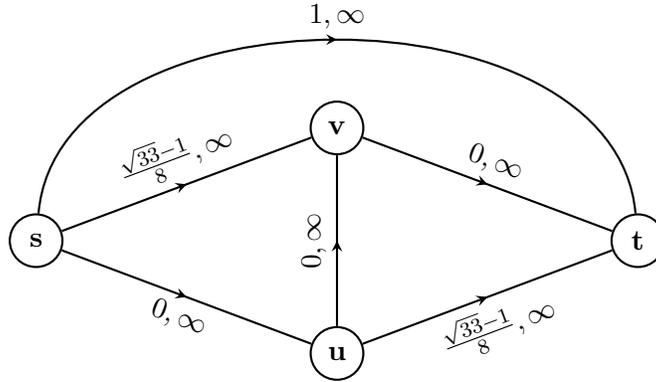


Figure 6: A simple acyclic graph with a single source s and sink t . All of the cost functions are infinity for load $\ell \geq 2$. This graph is used to prove Theorem 32.

Considering three players a, b and c , in the optimum one player uses the upper edge, and the other two the **svt** and the **sut** paths. The total cost of the optimum is $\frac{\sqrt{33}-1}{4} + 1$. If any other strategy profile was an equilibrium it would result in an unbounded total cost and give an unbounded PoA. Therefore, there should exist an optimum strategy profile \mathbf{S}^* that it is a Nash equilibrium. Assume w.l.o.g. that in \mathbf{S}^* , a uses the upper edge, b uses the **svt** path, and c uses the **sut** path.

If we consider now the instance in which only b and c use the network, either the strategy profile in which b uses the **svt** path, and c uses the **sut** path is an equilibrium, or one of b or c pays at least 1 in \mathbf{S}^* . The reason is that since \mathbf{S}^* is a Nash equilibrium, no player has an incentive to deviate to any other path but the upper one, and if neither b , nor c pays more than 1, obviously none of them would prefer the upper edge.

For the first case the cost of this strategy profile is $\frac{\sqrt{33}-1}{4}$, and the optimum with two players is one player to use the zig-zag path (**sutv**), and the other to use the upper edge with total cost 1. This results in $\text{PoA} \geq \frac{\sqrt{33}-1}{4}$.

For the second case if one of b or c pays at least 1 in \mathbf{S}^* the cost of \mathbf{S}^* becomes at least $\frac{\sqrt{33}-1}{8} + 2$

⁸Note that the order we process C 's in the reallocation procedure guarantees that any reallocation after processing C does not affect the input of C , so S_j^* cannot be swapped to pass through C_1 after processing C .

and we have

$$\text{PoA} \geq \frac{\frac{\sqrt{33}-1}{8} + 2}{\frac{\sqrt{33}-1}{4} + 1} = \frac{\sqrt{33}-1}{4}.$$

□

A.10 Proof of Theorem 33

Proof. We use the graph shown in Figure 7 to prove the statement.

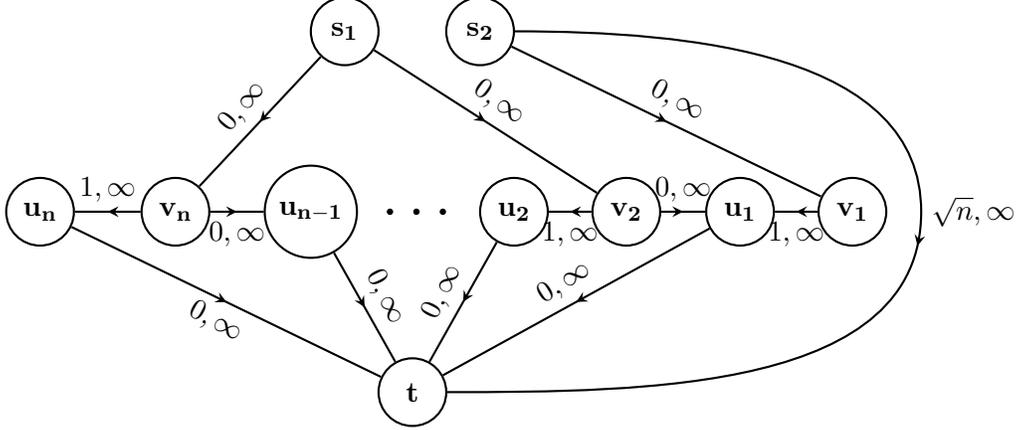


Figure 7: An acyclic graph with a two sources s_1 and s_2 , and sink t . All of the cost functions are infinity for load $\ell \geq 2$. This graph is used to prove Theorem 33.

Considering $n+1$ players where the first and the last players want to connect s_2 to t , and others want to connect s_1 to t . In the optimum, one of the players with source s_2 uses the right edge, the other uses the $s_2 v_1 u_1 t$ path, and other players use $s_1 v_i u_i t$ paths for $2 \leq i \leq n$. The total cost of the optimum is $n + \sqrt{n}$. If any other strategy profile was an equilibrium it would result in an unbounded total cost and give an unbounded PoA. Therefore, there should exist an optimum strategy profile \mathbf{S}^* that is a Nash equilibrium. Assume w.l.o.g. that in \mathbf{S}^* , player $n+1$ uses the right edge, player 1 uses the $s_2 v_1 u_1 t$ path, and for each $2 \leq i \leq n$, player i uses the $s_1 v_i u_i t$ path.

Now we consider two cases:

Case 1: Player 1 is charged at most \sqrt{n} for the $s_2 v_1 u_1 t$ path.

In this case we consider the instance in which only the first n players use the network. The strategy profile in which player 1 uses the $s_2 v_1 u_1 t$ path, and for each $2 \leq i \leq n$, player i uses the $s_1 v_i u_i t$ path should be an equilibrium. The reason is that since \mathbf{S}^* is a Nash equilibrium, no player has an incentive to deviate to any other path but the right one, and the only player that might have the incentive to deviate to the right edge is player 1. Since player 1 pays less than \sqrt{n} in $s_1 v_i u_i t$, she does not have the incentive to deviate either, and hence this strategy profile is an equilibrium with the total cost of at least n . However, the optimal strategy for these players is for player 1 to use the right edge and others to use the $s_1 v_i u_{i-1} t$ paths, with total cost \sqrt{n} . Therefore, we have $\text{PoA} \geq \sqrt{n}$ for this case.

Case 2: Player 1 is charged more than \sqrt{n} for the $s_2 v_1 u_1 t$ path.

In this case we consider a single player who wants to connect s_2 to t . The optimal strategy for such player costs 1, but with the assumption that she should pay at least \sqrt{n} , and hence we have $\text{PoA} \geq \sqrt{n}$ in this case too. □