

# Designing Networks with Good Equilibria under Uncertainty\*

George Christodoulou<sup>†</sup>      Alkmini Sgouritsa<sup>‡</sup>

June 21, 2019

## Abstract

We consider the problem of designing network cost-sharing protocols with good equilibria under uncertainty. The underlying game is a *multicast game* in a rooted undirected graph with nonnegative edge costs. A set of  $k$  *terminal vertices* or *players* need to establish connectivity with the root. The social optimum is the Minimum Steiner Tree. We study situations where the designer has incomplete information about the input. We propose two different models, the *adversarial* and the *stochastic*. In both models, the designer has prior knowledge of the underlying graph metric, but the requested subset of the players is *not known* and is activated either in an adversarial manner (adversarial model) or is drawn from a known probability distribution (stochastic model).

In the adversarial model, the goal of the designer is to choose a single, *universal* cost-sharing protocol that has low Price of Anarchy (PoA) for *all possible* requested subsets of players. The main question we address is: *to what extent can prior knowledge of the underlying graph metric help in the design?* We first demonstrate that there exist classes of graphs where knowledge of the underlying graph metric can dramatically improve the performance of good network cost-sharing design. For *outerplanar* graph metrics, we provide a universal cost-sharing protocol with constant PoA, in contrast to protocols that, by ignoring the graph metric, cannot achieve PoA better than  $\Omega(\log k)$ . Then, in our main technical result, we show that there exist graph metrics, for which knowing the underlying graph metric does not help and *any universal* protocol has PoA of  $\Omega(\log k)$ , which is tight. We attack this problem by developing new techniques that employ powerful tools from extremal combinatorics, and more specifically Ramsey Theory in high dimensional hypercubes.

Then we switch to the stochastic model, where the players are activated according to some probability distribution that is known to the designer. We show that there exists a *randomized* ordered protocol that achieves constant PoA. If further each player is activated independently with some probability, by using standard derandomization techniques, we produce a *deterministic* ordered protocol that achieves constant PoA. We remark, that the first result holds also for the *black-box* model, where the probabilities are not known to the designer, but she is allowed to draw independent (polynomially many) samples.

## 1 Introduction

**Network Cost-Sharing Games** We study a *multicast game* in a rooted undirected graph  $G = (V, E)$  with a nonnegative cost  $c_e$  on each edge  $e \in E$ . A set of  $k$  *terminal vertices* or *players*

---

\*An extended abstract of this paper appeared in the Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016 [28].

<sup>†</sup>Department of Computer Science, University of Liverpool (Email:{gchristo@liv.ac.uk}). This author was supported by EP/M008118/1 and LT140046.

<sup>‡</sup>Department of Algorithms and Complexity, Max Planck Institute for Informatics, (Email:{sgouritsa@gmail.com}). This author was supported by the Lise Meitner Award Fellowship. Part of this work has been done at the University of Liverpool.

$s_1, \dots, s_k$  need to establish connectivity with the root  $t$ . Each player selects a path  $P_i$  and the outcome produced is the graph  $H = \cup_i P_i$ . The global objective is to minimize the cost,  $\sum_{e \in H} c_e$ , of this graph, which is the *Minimum Steiner Tree*.

The cost of an edge may represent infrastructure cost for establishing connectivity or renting expense, and needs to be covered by the players that use that edge in the solution. There are several ways to split the edge costs among the users and this is dictated by a *cost-sharing protocol*. Naturally, it is in the players' best interest to choose paths that charge them with small cost, and therefore the solution will be a Nash equilibrium (NE). Algorithmic Game Theory provides tools to analyse the quality of the equilibrium solutions; this can be measured with the Price of Anarchy (PoA) [48] (or Price of Stability (PoS) [5]) that compares the worst-case (or the best-case) cost in a NE with the cost of the minimum Steiner tree. This is a fundamental network design game that was originated by Anshelevich et al. [5] and has been extensively studied since. [5] studied the *Shapley* cost-sharing protocol, where the cost of each edge is equally split among its users. They showed that the quality of equilibria can be really poor<sup>1</sup>.

**Cost-Sharing Protocol Design** Different cost-sharing protocols result in different quality of equilibria. In this work, we are interested in the design of protocols that induce good equilibrium solutions in the *worst-case*, therefore we focus on protocols that guarantee low PoA. Chen, Roughgarden and Valiant [23] were the first to address design questions for network cost-sharing games. They gave a characterization of protocols that satisfy some natural axioms and they thoroughly studied their PoA for the following two classes of protocols, that use different informational assumptions from the perspective of the designer.

*Non-uniform protocols.* The designer has full knowledge of the instance, that is, she knows both the network topology given by  $G$  and the costs  $c_e$ , and in addition the set of players' requests  $s_1, \dots, s_k$ . They showed that a simple priority protocol (see Example 1) has a constant PoA; the NE induced by the protocol simulate Prim's algorithm for the Minimum Spanning Tree (MST) problem, and therefore achieve constant approximation.

*Uniform protocols.* The designer needs to decide how to split the edge cost among the users *without knowledge of the underlying graph*. They showed that the PoA is  $\Theta(\log k)$ ; both upper and lower bound comes from the analysis of the Greedy Algorithm for the Online Steiner Tree problem.

**Cost-Sharing Design under Uncertainty** Arguably, there are situations where the former assumption is too optimistic while the latter is too pessimistic. We propose a model that lies in the middle-ground as a framework to design network cost-sharing protocols with good equilibria, when the designer has *incomplete information*.

We assume that the designer has prior knowledge of the underlying graph metric, (given by the graph  $G$  and the shortest path metric induced by the costs  $c_e$ ), but is *uncertain* about the requested subset of players. We consider two different models, the *adversarial model* and the *stochastic model*. In the former, the designer *knows nothing* about the number or the positions of the  $s_i$ 's and has as goal to process the graph and choose a single, *universal* cost-sharing protocol that has low PoA against *all possible* requested subsets. Here, no distributional assumptions are made about arrivals

---

<sup>1</sup>Even for simple networks the PoA grows linearly with the number of players,  $k$ . The PoS is not well-understood. It is a big open question to determine its exact value that is between constant and  $O(\log k / \log \log k)$  [50].




Figure 1: Figure (a) illustrates the worst-case graph for the Online Steiner Tree problem. Figure (b) is a variation of (a) that serves as the worst-case graph for the universal protocols network cost-sharing game. In edges with no written cost, we consider the unit cost; note that both graphs can be easily generalized so that the number of vertices is arbitrarily big (for (a) see also [43] for more details). In (a) and (b) we assume two orders on the vertices, denoted by  $q_i$  or  $p_i$ . The  $q$ -order is adversarially chosen and results to high PoA of  $\Omega(\log k)$ . The  $p$ -order results to constant PoA. Figure (c) shows an example where both the *best* ordered protocol and the Shapley protocol have  $\text{PoA} \geq 5/4$ , whereas there is an *intermediate* protocol with  $\text{PoA}=1$ ; we set  $\varepsilon > 0$  arbitrarily small.

of players, instead the worst-case approach is used similarly to Competitive Analysis. Once the designer selects the protocol, then an adversary will choose the requested subset of players and their positions in the graph (the  $s_i$ 's), in a way that *maximizes* the PoA of the induced game. In the stochastic model, the players/vertices are activated according to some probability distribution which is given to the designer. The goal is now to choose a universal protocol where the expected worst-case cost in the NE is not far from the expected optimal cost.

The following two examples demonstrate how the knowledge of the underlying graph metric can help in the protocol design resulting in improvements on the PoA.

**Example 1. (Ordered protocols).** In this example, we restrict ourselves to a specific class of protocols, called ordered protocols, and show that knowing the underlying graph metric can dramatically improve the PoA even for this strict class of protocols. *Ordered protocols* consist an important special class with interesting properties. The designer decides a total order of the users, and when a subset of players uses some edge, the full cost is covered by the player who comes first in the order. Any NE of the induced game corresponds to the solution produced by the Greedy Algorithm for the MST: each player is connected, via a shortest path, with the component of the players that come before him in the order. The analysis of the PoA in the uniform model boils down to the analysis of the Greedy Algorithm for the Online Steiner Tree problem, where the worst-case order is considered.

The following instance demonstrates that even this special class of ordered protocols becomes very rich, once the designer has prior knowledge of the underlying metric space. Uniform protocols throw away this crucial component, the structure of the underlying graph metric, that universal protocols can use in their favour in order to come up with better PoA guarantees.

Uniform protocols. The designer chooses an order of the players  $1, \dots, k$  without prior knowledge of the graph. The adversary constructs a worst-case graph by simulating the adversary for the Greedy Algorithm of the Online Steiner Tree problem [43] and places the players accordingly. See for example Figures 1(a),(b), the  $q$  labels. There is a Nash equilibrium that is formed by the bold edges, whereas the optimum solution is the path  $(t, q_3, q_2, q_4, q_1)$  for 1(a) and the path  $(t, q_1, q_6, q_4, q_7, q_3, q_8, q_5, q_9, q_2)$  for 1(b). Therefore, the PoA of uniform ordered protocol is  $\Omega(\log k)$  [23].

Universal protocols. The designer takes into account the graph; consider again the graphs of Figures 1(a),(b). For the graph of Figure 1(a), order the vertices according to their distance from  $t$  ( $p$  labels). For the graph of Figure 1(b), choose the linear order dictated from the path  $p_1, \dots, p_9$  (say from left to right). The adversary will choose  $k$  and the positions of the players  $(s_1, \dots, s_k)$ . In both cases, it is not hard to see that, *no matter which subset of players the adversary chooses*, the PoA remains constant as  $k$  grows.

**Example 2. (Generalized weighted Shapley).** In [23], it was shown that ordered protocols are essentially optimal among uniform protocols. Optimality of ordered protocols is no longer true in the case where the underlying graph metric is known in advance. Figure 1(c) shows an instance where Shapley cost-sharing protocol and ordered protocols have PoA at least  $5/4$ , while there exists a (generalized weighted Shapley) protocol that achieves PoA=1.

By using Shapley cost-sharing the adversary can choose to activate  $\{v_1, v_2, v_3\}$  and it is a NE if  $v_1, v_3$  connect directly to  $t$  and  $v_2$  connects through  $v_1$ . Regarding *any* ordered protocol, the square defined by the  $v_i$ 's contains a path of length 2 where the middle vertex comes *last* in the order. The adversary will select this triplet of players, say  $v_1, v_2, v_3$ . In the NE,  $v_1$  connects directly to  $t$ ,  $v_3$  and  $v_2$  connect through  $v_1$ . In both cases, (by ignoring  $\varepsilon$ ) the cost of the NE is 5 and the minimum Steiner tree that connects those vertices with  $t$  has cost 4 and therefore,  $\text{PoA} \geq 5/4$ .

However, the following (generalized weighted Shapley) protocol, has PoA=1. Partition the players into two sets  $S_1 = \{v_1, v_2\}, S_2 = \{v_3, v_4\}$ . If players from both partitions appear on some edge, then the cost is charged only to players from  $S_1$ . Players that belong to the same partition share the cost equally. One can verify that for all possible subsets of players this protocol produces only optimal equilibria.

**Results** We propose a framework for the design of (universal) network cost-sharing protocols with good pure Nash equilibria, in situations where the designer has incomplete information about the input. We consider two different models, the *adversarial* and the *stochastic*. In both models, the designer has prior knowledge of the underlying graph metric but the requested subset of players is *not known* and is activated either in an adversarial manner (adversarial model) or is drawn from a known probability distribution (stochastic model). The central question we address is: *to what extent does prior knowledge of the graph metric help in good network design under uncertainty?*

For the adversarial model, we first demonstrate that there exist classes of graph metrics where prior knowledge of the underlying graph metric can dramatically improve the performance of good network cost-sharing design. For *outerplanar* graph metrics, we provide a universal ordered cost-sharing protocol with constant PoA, against any choice of the adversary. This is in contrast to uniform protocols that ignore the graph and cannot achieve PoA better than  $\Omega(\log k)$  in outerplanar graph metrics.

**Open Question** Can the design of universal protocols for planar graph metrics or even the grid graph improve the PoA guarantees compared to uniform protocols that ignore the underlying graph metric?

Our main technical result shows that there exist graph metrics, for which knowing the underlying graph metric does not help the designer, and *any universal* protocol has PoA of  $\Omega(\log k)$ . This matches the upper bound of  $O(\log k)$  that can be achieved without prior knowledge of the graph metric [43, 23].

Then we switch to the stochastic model, where the players (terminal vertices) are activated according to some probability distribution that is known to the designer. We show that there exists a *randomized* ordered protocol that achieves constant PoA. If each player is activated independently

with some probability, by using standard derandomization techniques [58, 54], we produce a *deterministic* ordered protocol that achieves constant PoA. We remark, that the first result holds also for the *black-box* model, where the probability distribution is not known to the designer, but is allowed to draw independent (polynomially many) samples.

Our results for the adversarial model motivate the following question that is left open.

**Open Question** For which metric spaces can one design universal protocols with constant PoA? What sort of structural graph properties are needed to obtain good guarantees?

**Techniques** We prove our main lower bound for the adversarial model in two parts. In the first part (Section 4) we bound the PoA achieved by *any ordered* protocol. Our origin is a well-known “zig-zag” *ordered* structure that has been used to show a lower bound on the Greedy Algorithm of the Online Steiner Tree problem (see the labeled path  $(q_1, q_6, q_4, \dots, q_2)$  in Figure 1(b)). The challenge is to show that high dimensional hypercubes exhibit such a distance preserving structure *no matter how the vertices are ordered*. Section 4 is devoted to this task and we believe that this is of independent interest.

We show the existence proof by employing powerful tools from Extremal Combinatorics and in particular Ramsey Theory [39]. We are inspired by a Ramsey-type result due to Alon et al. [4], in which they show that for any given length  $\ell \geq 5$ , any  $r$ -edge coloring of a high dimensional hypercube contains a monochromatic cycle of length  $2\ell$ . Unfortunately, we cannot immediately use their results, but we show a similar Ramsey-type result for a different, carefully constructed structure; we assert that every 2-edge coloring of high dimensional hypercubes  $Q_n$  contains a monochromatic copy of that structure. Then, we prescribe a special 2-edge-coloring that depends on the ordering of  $Q_n$ , so that the special subgraph preserves some nice labeling properties. A suitable subset of the subgraph’s vertices can be 1-embedded into a hypercube of lower dimension. Recursively, we show existence of the desired *distance preserving* “zig-zag” structure.

In the second part (Section 5), we extend the lower bound to *all universal* cost-sharing protocols, by using the characterization of [23]. At a high level, we use as basis the construction for the ordered protocol and create “multiple copies”<sup>2</sup>. The adversary will choose different subsets of players, depending on whether the designer chose protocols “closer” to Shapley or to ordered. In the latter case, we use arguments from Matching Theory to guarantee existence of ordered-like players in one of the hypercubes.

For the stochastic model (Section 6), we construct an approximate minimum Steiner tree over a subset of vertices which is drawn from the known probability distribution. This tree is used as a base to construct a spanning tree, which determines a total order over the vertices. We finally produce a deterministic order by applying standard derandomization techniques [58, 54].

**Related Work** Following the work of [5, 6], a long line of research studies network cost-sharing games, mainly focusing on the PoS of the Shapley cost-sharing mechanism. [5] showed a tight  $\Theta(\log k)$  bound for directed networks, while for undirected networks several variants have been studied [15, 16, 17, 22, 24, 32, 33, 50] but the exact value of PoS still remains a big open problem. For multicast games, an improved upper bound of  $O(\log k / \log \log k)$  is known due to Li [50], while for broadcast games (where every vertex is a terminal of some player) a series of work [33, 49] lead

---

<sup>2</sup>Note that the standard complexity measure, to analyze the inefficiency of equilibria, is the number of participants,  $k$ , and not the total number of vertices in the graph (see for example [5, 23]).

finally to a constant due to Bilò et al. [17]. The PoA of some special equilibria has been also studied in [20, 21].

Chen, Roughgarden and Valiant [23] initiated the study of network cost-sharing design with respect to PoA and PoS. They characterized a class of protocols that satisfy certain desired properties (which was later extended by Gopalakrishnan, Marden and Wierman, in [37]), and they thoroughly studied PoA and PoS for several cases. Recently, Christodoulou, Leonardi and Sgouritsa [26] studied the Bayesian network design showing a lower bound of  $\Omega(\sqrt{k})$  for any cost-sharing protocol satisfying the same properties. von Falkenhausen and Harks [57] studied singleton and matroid games with weighted players, while Gkatzelis, Kollias and Roughgarden [35] focus on weighted congestion games with polynomial cost functions. The very recent work of Harks, Huber and Surek [42] is very similar to our work in the sense that the designer has some knowledge of the underlying graph; they thoroughly characterized the topological properties of the underlying graph so that the optimum solution is a pure Nash equilibrium.

Close in spirit to universal cost-sharing protocols is the notion of Coordination Mechanisms [25] that provides a way to improve the PoA in cases of incomplete information. The designer has to decide in advance local scheduling policies or increases in edge latencies, without knowing the exact input, and has been used for scheduling problems [1, 2, 8, 13, 19, 25, 30, 44, 47] as well as for simple routing games [14, 27].

As discussed in Example 1, the analysis of the equilibria induced by ordered protocols corresponds to the analysis of the Greedy Algorithm for the MST. In the uniform model, this corresponds to the analysis of the Greedy Algorithm [7, 43] for the (Generalized) Online Steiner Tree problem [3, 9, 56], which was shown to be  $\Theta(\log k)$ -competitive by Imase and Waxman [43] ( $O(\log^2 k)$ -competitive for the Generalized Online Steiner Tree problem by [7]). The universal model is closely related to universal network design problems [45], hence our choice for the term “universal”. In the universal TSP, given a metric space, the algorithm designer has to decide a *master* order so that tours that use this order have good approximation [10, 12, 29, 38, 41, 45, 52].

Much work has been done in stochastic models and we only mention the most related to our work. Karger and Minkoff [46] showed a constant approximation guarantee for the maybecast problem, where the designer needs to fix (before activation) some path for every vertex to the root. Garg et al. [34] gave bounds on the approximation of the stochastic online Steiner tree problem. A line of works [11, 38, 53, 54] has studied the *a priori TSP*. Shmoys and Talwar [54] assumed independent activations and demonstrated randomized and deterministic algorithms with constant approximations.

## 2 Model and definitions

**Universal Cost-Sharing Protocols** A *multicast network cost-sharing game*, is specified by a connected undirected graph  $G = (V, E)$ , with a designated root  $t$  and nonnegative weight  $c_e$  for every edge  $e$ , a set of players  $S = \{1, \dots, k\}$  and a cost-sharing protocol. Each player  $i$  is associated with a *terminal*<sup>3</sup>  $s_i$ , which she needs to connect with  $t$ . We say that a vertex is *activated* if there exists some requested player associated with it. In the *adversarial model* the designer *knows nothing* about the set  $S$  of activated vertices, while in the *stochastic model*, the vertices are activated according to some probability distribution  $\Pi$  which is known to the designer.

For any set,  $N$ , of players, a *cost-sharing method*  $\xi_e : 2^N \rightarrow \mathbb{R}_+^{|N|}$  is a function of the set of players,  $R \subseteq N$ , using edge  $e$  and decides the cost-share for each player  $i \in N$ .<sup>4</sup> A natural rule is

<sup>3</sup>We abuse notation and use  $S$  to refer both to the players and their associated vertices.

<sup>4</sup>We do not define the cost-sharing method for the set  $S$ , since in our setting the designer is not aware of  $S$ .

that the shares for players not included in  $R$  should always be 0. We use the notation  $\xi_e(i, R)$  to denote the cost-share of player  $i$  under input  $R$ ; note that if  $i \notin R$ , then  $\xi_e(i, R) = 0$ . For any graph  $G$  and any set of players  $N$ , a *cost-sharing protocol*  $\Xi$  assigns, for every  $e \in E$ , some cost-sharing method  $\xi_e$ .

Following previous work [23, 42, 57], we focus on cost-sharing protocols that satisfy the following natural properties:

- (1) *Budget-balance*: For every network game induced by the cost-sharing protocol  $\Xi$ , and every outcome of it,  $\sum_{i \in R} \xi_e(i, R) = c_e$ , for every edge  $e$  with cost  $c_e$ .
- (2) *Separability*: For every network game induced by the cost-sharing protocol  $\Xi$ , the cost-shares of each edge are completely determined by the set of players using it.
- (3) *Stability*: In every network game induced by the cost-sharing protocol  $\Xi$ , there exists at least one pure Nash equilibrium, regardless of the graph structure.

Under the assumption that each player is associated with a distinct vertex, which can be done w.l.o.g.<sup>5</sup>, we call a cost-sharing protocol  $\Xi$  *universal*, if it satisfies the above properties for any graph  $G$ , and it assigns the cost-sharing method<sup>6</sup>  $\xi_e : 2^V \rightarrow \mathbb{R}_+^{|V|}$  to any edge  $e$  based only on knowledge of  $G$  (without any knowledge of  $S$ ) for the adversarial model, while in the stochastic model the method can in addition depend on  $\Pi$ .

Note that in the way that we define the Stability property, we require that the protocol admits a pure Nash equilibrium *regardless of the graph structure*. This is because the network may evolve over time and it is essential that existence of pure Nash equilibrium is always guaranteed. As we discuss next, due to the work of [23] we may consider only the generalized weighted Shapley protocols. If one drops the requirement of pure Nash equilibrium existence regardless of the graph structure and requires only the existence of a pure Nash equilibrium on a specific network, then other protocols may exist; see for example the work of Marden and Wierman [51] for the case of parallel links.

**Generalized Weighted Shapley Protocol (GWSP)** The *generalized weighted Shapley protocol (GWSP)* is defined by the players' weights (parameters)  $\{w_1, \dots, w_n\}$  and an *ordered* partition of the players  $\Sigma = (U_1, \dots, U_h)$ . An interpretation of  $\Sigma$  is that for  $i < j$ , players from  $U_i$  appear before players from  $U_j$  and therefore  $U_j$  are charged zero cost for their common edges with  $U_i$ . More formally, for every edge  $e$  of cost  $c_e$ , every set of players  $R_e$  that uses  $e$  and for  $s = \arg \min_j \{U_j | U_j \cap R_e \neq \emptyset\}$ , the GWSP assigns the following method to  $e$ :

$$\xi_e(i, R_e) = \begin{cases} \frac{w_i}{\sum_{j \in U_s \cap R_e} w_j} c_e, & \text{if } i \in U_s \cap R_e \\ 0, & \text{otherwise} \end{cases}$$

In the special case that each  $U_i$  contains exactly one player, the protocol is called *ordered*. The order of the  $U_i$  sets indicates a permutation of the players, denoted by  $\pi$ .

In this paper we restrict ourselves to the family of generalized weighted Shapley protocols which is justified due to the characterization of Chen, Roughgarden and Valiant [23]. Next we restate their Theorem 3.8 by using their proof of Lemma 4.4.

Instead, we define it for the set of all potential players.

<sup>5</sup>To see this, if there are two players with  $s_1 = s_2 = v$ , for some  $v \in V$ , we modify the graph by connecting a new vertex  $v'$  with  $v$  via a zero-cost edge and then we set  $s_1 = v$  and  $s_2 = v'$ . Neither the optimum solution, nor any NE is affected by this modification.

<sup>6</sup>The methods should be defined on  $V$ , since every vertex is potentially associated with some player.

**Theorem 3.** (Theorem 3.8 of [23]). A cost-sharing protocol satisfies budget-balance, separability and stability in any multicast game with unit edge costs if and only if it is a generalized weighted Shapley protocol.

**(Pure) Nash Equilibrium (NE)** We denote by  $\mathcal{P}_i$  the strategy space of player  $i$ , i.e. the set of all the paths connecting  $s_i$  to  $t$ .  $\mathbf{P} = (P_1, \dots, P_k)$  denotes an *outcome* or a *strategy profile*, where  $P_i \in \mathcal{P}_i$  for all  $i \in S$ . As usual,  $\mathbf{P}_{-i}$  denotes the strategies of all players but  $i$ . Let  $R_e$  be the set of players using edge  $e \in E$  under  $\mathbf{P}$ . The cost-share of player  $i$  induced by  $\xi_e$ 's is equal to  $c_i(\mathbf{P}) = \sum_{e \in P_i} \xi_e(i, R_e)$ . The players' objective is to minimize their cost-share  $c_i(\mathbf{P})$ . A strategy profile  $\mathbf{P} = (P_1, \dots, P_k)$  is a *Nash equilibrium (NE)* if for every player  $i \in S$  and every strategy  $P'_i \in \mathcal{P}_i$ ,  $c_i(\mathbf{P}) \leq c_i(\mathbf{P}_{-i}, P'_i)$ .

**Price of Anarchy (PoA)** The cost of an outcome  $\mathbf{P} = (P_1, \dots, P_k)$  is defined as  $c(\mathbf{P}) = \sum_{e \in \cup_i P_i} c_e$ , while  $\mathbf{O} = (O_1, \dots, O_k) \in \arg \min_{\mathbf{P}} c(\mathbf{P})$  is an optimum solution. The *Price of Anarchy (PoA)* is defined as the worst-case ratio of the cost in a NE over the optimal cost in the game induced by  $S$ . In the adversarial model the *worst-case*  $S$  is chosen, while in the stochastic model  $S$  is drawn from a known distribution  $\Pi$ . Formally, in the adversarial model we define the PoA of a protocol  $\Xi$  on  $G$  as

$$PoA(G, \Xi) = \max_{S \subseteq V \setminus \{t\}} \frac{\max_{\mathbf{P} \in \mathcal{N}} c(\mathbf{P})}{c(\mathbf{O})},$$

where  $\mathcal{N}$  is the set of all NE of the game induced by  $\Xi$  and  $S$  on  $G$ .

In the stochastic model, the PoA of  $\Xi$ , given  $G$  and  $\Pi$  is

$$PoA(G, \Xi, \Pi) = \frac{\mathbb{E}_{S \sim \Pi} [\max_{\mathbf{P} \in \mathcal{N}} c(\mathbf{P})]}{\mathbb{E}_{S \sim \Pi} [c(\mathbf{O})]}.$$

In both models the objective of the designer is to come up with protocols that *minimize* the above ratios. Finally, the Price of Anarchy for a class of graph metrics  $\mathcal{G}$ , is defined in the two models, respectively, as

$$PoA(\mathcal{G}) = \max_{G \in \mathcal{G}} \min_{\Xi(G)} PoA(G, \Xi); \quad PoA(\mathcal{G}) = \max_{G \in \mathcal{G}} \min_{\Xi(G, \Pi)} \max_{\Pi} PoA(G, \Xi, \Pi).$$

**Graph Theory** For every graph  $G$ , we denote by  $V(G)$  and  $E(G)$  the set of vertices and edges of  $G$ , respectively. For any  $v, u \in V(G)$ ,  $(v, u)$  denotes an edge between  $v$  and  $u$  and  $d_G(v, u)$  denotes the shortest distance between  $v$  and  $u$  in  $G$ ; if  $G$  is clear from the context, we simply write  $d(v, u)$ . A graph  $G$  is an *induced* subgraph of  $H$ , if  $G$  is a subgraph of  $H$  and for every  $v, u \in V(G)$ ,  $(v, u) \in E(G)$  if and only if  $(v, u) \in E(H)$ .  $G$  is a *distance preserving (isometric)* subgraph of  $H$ , if  $G$  is a subgraph of  $H$  and for every  $v, u \in V(G)$ ,  $d_G(v, u) = d_H(v, u)$ .

### 3 Outerplanar Graphs

In this section we show that there exists a class of graph metrics, prior knowledge of which can dramatically improve the performance of good network cost-sharing design. For *outerplanar* graphs, we provide a universal cost-sharing protocol with constant PoA. In contrast, we stress that uniform protocols *cannot* achieve PoA better than  $\Omega(\log k)$ , because the lower bound for the greedy algorithm of the Online Steiner Tree problem can be embedded in an outerplanar graph (see Figure 2

for an illustration). An *outerplanar* graph is a planar graph where all the vertices belong to the outer face. For a *biconnected*<sup>7</sup> *outerplanar* graph the outer face forms a (unique) Hamiltonian cycle.

We next define an ordered universal cost-sharing protocol,  $\Xi_{tour}$ , and we show that it has constant PoA. We describe  $\Xi_{tour}$  only for metric spaces that are defined by *biconnected* outerplanar graphs. In order to define  $\Xi_{tour}$  for an outerplanar graph  $G$  that is not biconnected, we first turn it into an *equivalent*<sup>8</sup> biconnected graph  $G^*$ , by appropriately adding edges of sufficiently high cost  $h$ . In order to do this, we can set  $h$  to be a value strictly greater than  $\sum_{e \in E(G)} c_e$ . Then, equivalence is obvious since we only add edges that cannot be used in either any NE or the minimum Steiner tree outcome. Hence, it is w.l.o.g. to consider only biconnected outerplanar graphs. It is known that every biconnected outerplanar graph admits a *unique* Hamiltonian cycle [55] that can be found in linear time [31].

**Definition of  $\Xi_{tour}$ :**  $\Xi_{tour}$  orders the vertices according to the cyclic order in which they appear in the Hamiltonian tour, starting from  $t$  and proceeding in a clockwise order  $\pi$ . In Figure 2,  $\pi(t) < \pi(q_8) < \pi(q_4) < \pi(q_9) < \dots < \pi(q_{15})$ .

In the following theorem we show that, for outerplanar graphs, the PoA of  $\Xi_{tour}$  is constant and more precisely is upper bounded by 2.




Figure 2: The figure shows an example of an outerplanar graph where the order  $q_i < q_{i+1}$  gives PoA of  $\Omega(\log k)$ .

**Theorem 4.** The PoA of  $\Xi_{tour}$  in outerplanar graphs is at most 2.

*Proof.* Let  $G = (V, E, t)$  be any biconnected outerplanar graph and  $S$  be the set of activated vertices. Let  $T^*$  be the minimum Steiner tree that connects  $S \cup \{t\}$  and suppose that is rooted at  $t$ . We denote by  $P_{T^*}(i, i')$  the unique path from  $i$  to  $i'$  in  $T^*$  and by  $T_v^*$  the subtree of  $T^*$  rooted at vertex  $v$ .

We first show the following claims that will be useful to complete the proof.

*Claim 5.* For any  $i, i', j, j' \in V(T^*)$  such that  $\pi(i) < \pi(j) < \pi(i') < \pi(j')$ , the paths  $P_{T^*}(i, i')$  and  $P_{T^*}(j, j')$  are not vertex-disjoint, (i.e.  $P_{T^*}(i, i')$  and  $P_{T^*}(j, j')$  share a common vertex) .

*Proof.* Consider the representation of  $G$  as a planar graph where the *unique* Hamiltonian tour of  $G$  is the outer face, meaning that all the edges of  $P_{T^*}(i, i')$  and  $P_{T^*}(j, j')$  are either edges of the Hamiltonian tour or chords of it. The Hamiltonian tour defines two paths between  $i$  and  $i'$ , one

<sup>7</sup>A graph is biconnected if, after removing any vertex, the graph remains connected.

<sup>8</sup>We mean that any NE outcome and the minimum Steiner tree solution remain unchanged after the transformation and therefore the PoA of  $G$  is exactly the same with the one of  $G^*$ .

containing  $j$  and the other containing  $j'$ . Hence, the paths  $P_{T^*}(i, i')$ ,  $P_{T^*}(j, j')$  have either crossing edges or some common vertex. Due to the planarity of  $G$  the first case is excluded and the claim follows.  $\square$

*Claim 6.* For any two vertex-disjoint subtrees  $T_v^*$  and  $T_u^*$  of  $T^*$ , rooted at vertices  $v$  and  $u$ , respectively, either all the vertices of  $T_v^*$  precedes all the vertices of  $T_u^*$ , or the opposite.

*Proof.* Assume on the contrary that w.l.o.g. there exist  $i' \in V(T_v^*)$  and  $j, j' \in V(T_u^*)$  such that  $\pi(j) < \pi(i') < \pi(j')$ . Since  $T_v^*$  and  $T_u^*$  are vertex-disjoint subtrees, first  $t \notin V(T_v^*)$  and  $t \notin V(T_u^*)$  and further the paths  $P_{T^*}(t, i')$  and  $P_{T^*}(j, j')$  should also be vertex-disjoint. Notice, though, that  $\pi(t) < \pi(j) < \pi(i') < \pi(j')$  and so, by Claim 5 we end up with a contradiction.  $\square$

*Claim 7.* For any two vertices  $v, u \in V(T^*)$  where  $v$  is an ancestor of  $u$ , then  $v$  either precedes or follows all the vertices of the subtree  $T_u^*$ .

*Proof.* If  $v$  is the root then trivially  $v$  precedes all the vertices of  $T_u^*$ . Suppose now that  $v \neq t$  and consider the case that  $v$  precedes  $u$  (the other case is similar). For the sake of contradiction, assume that there exists a vertex  $u' \in V(T_u^*)$  such that  $u'$  precedes  $v$  and therefore,  $\pi(t) < \pi(u') < \pi(v) < \pi(u)$ . Note that  $P_{T^*}(t, v)$  and  $P_{T^*}(u, u')$  are vertex-disjoint and therefore, by Claim 5 we end up with a contradiction.  $\square$

For convenience, we next refer to the set of the activated vertices as  $S = \{1, 2, \dots, k\}$ , based on their order  $\pi$ , from smaller label to larger, i.e. vertex  $i$  has the  $i^{\text{th}}$  smallest label among  $S$ . We further adopt the convention that  $t = 0$ .

Consider any NE,  $\mathbf{P} = (P_i)_{i \in N}$ . We bound from above the cost-share of each player at vertex  $i \in [k]$  by the cost of the path in  $T^*$  that connects her with  $i - 1$ , i.e.,

$$c_i(\mathbf{P}) \leq c(P_{T^*}(i, i - 1)) = \sum_{e \in P_{T^*}(i, i - 1)} c_e.$$

Then, by summing over  $S$ ,

$$c(\mathbf{P}) = \sum_{i \in [k]} c_i(\mathbf{P}) \leq \sum_{i \in [k]} c(P_{T^*}(i, i - 1)) = \sum_{i \in [k]} \sum_{e \in P_{T^*}(i, i - 1)} c_e = \sum_{e \in E(T^*)} \sum_{i: e \in P_{T^*}(i, i - 1)} c_e.$$

We argue next that by Claims 6 and 7 we can infer that, for each edge  $e$  of  $E(T^*)$ , there exist at most two paths  $P_{T^*}(i, i - 1)$  containing  $e$ , leading to:

$$c(\mathbf{P}) \leq \sum_{e \in E(T^*)} 2c_e = 2c(T^*).$$

To explain the last argument, consider any edge  $e = (v', v) \in E(T^*)$  and let  $v$  be the child of  $v'$  in  $T^*$ . For any vertex  $u \notin V(T_v^*)$ , either  $T_v^*$  and  $T_u^*$  are vertex-disjoint, or  $u$  is an ancestor of  $v$  in  $T^*$ . In either case, by Claims 6 and 7,  $u$  either precedes or follows *all* vertices of  $T_v^*$ . Let  $\ell, h \in [k]$  be the vertices of  $S \cap V(T_v^*)$ <sup>9</sup> with the lowest and the highest labels, respectively (it is possible that  $\ell = h$ ). It is easy to see that only the paths  $P_{T^*}(\ell, \ell - 1)$  and  $P_{T^*}(h + 1, h)$  use edge  $e$  (the second path exists only if  $h < k$ ).  $\square$

We next demonstrate that our analysis is tight.

<sup>9</sup>There should exist at least one activated vertex in  $T_v^*$ , otherwise  $e \notin E(T^*)$ .

**Proposition 8.** The PoA of  $\Xi_{tour}$  in outerplanar graphs is at least 2.

*Proof.* Consider a cycle graph  $C = (V, E, t)$  with  $2k$  vertices and unit-cost edges. Let the vertices  $V = \{t = 0, 1, 2, \dots, 2k - 1\}$  be named based on their order  $\pi$ , from smaller to larger label. We consider the set of the  $k$  activated vertices to be  $S = \{k, k + 1, \dots, 2k - 1\}$ . It is a NE if the player on vertex  $k$  connects with  $t$  through the path  $(0, 1, \dots, k)$  and each other player on vertex  $k + i$ , for  $i \in [k - 1]$ , connects with the vertex  $k + i - 1$  and follow their path to the root. The cost of this NE is  $2k - 1$ .

The optimum solution would be to connect  $S$  with  $t$  through the path  $(k, k + 1, \dots, 2k - 1, t)$  with cost  $k$ . Therefore,  $\text{PoA} \geq \frac{2k-1}{k}$  which for large  $k$  it converges to 2.  $\square$

## 4 Lower Bound of Ordered Protocols

The main result of this section is that the PoA of *any* ordered protocol is  $\Omega(\log k)$  which is tight<sup>10</sup>. We formally define (Definition 10) the ‘zig-zag’ pattern of the lower bounds of the Greedy Algorithm of the Online Steiner Tree problem (see Example 1(b) and Figure 3). Then the main technical challenge is to show that *for any* ordering of the vertices of high dimensional hypercubes, there always exists a *distance preserving* path, such that the order of its vertices follows that zig-zag pattern. Finally, by connecting any two vertices of the hypercube with a direct edge of suitable cost, similar to the example in Figure 1(b), we get the final lower bound construction.

Before defining the zig-zag pattern, we give the definition of Classes (Definition 9), which is a partition of the path’s vertices. Informally, given a path graph  $P = (v_0, \dots, v_{2^r})$ , we define a partition of its vertices into  $r + 1$  classes,  $D_0, D_1, \dots, D_r$ , as described next. Suppose that we construct  $P$  in  $r + 1$  steps, where at each step  $j \in \{0, \dots, r\}$  we introduce the vertices of class  $D_j$  as follows. At step 0 we connect the two endpoints  $v_0, v_{2^r}$  via an edge and  $D_0 = \{v_0, v_{2^r}\}$ . At step 1 we replace the edge  $(v_0, v_{2^r})$  by a two-length path  $(v_0, v_{2^{r-1}}, v_{2^r})$ , i.e. we place the vertex of  $D_1 = \{v_{2^{r-1}}\}$  between the existing vertices. Repeatedly, at each step  $j$  we replace each of the current edges by a two-length path, in the middle of which we place a vertex of  $D_j$ . We next give a formal definition of the vertices’ partition.

**Definition 9** (Classes). For  $r > 0$  and for a path graph  $P = (v_0, \dots, v_{2^r})$  of  $2^r + 1$  vertices, we define a partition of the vertices into  $r + 1$  classes,  $D_0, D_1, \dots, D_r$ , as follows: Class 0 contains the endpoints of  $P$ ,  $D_0 = \{v_0, v_{2^r}\}$ . For every  $j \in [r]$ ,  $D_j = \{v_i \mid i \equiv 0 \pmod{2^{r-j}} \text{ and } i \not\equiv 0 \pmod{2^{r-j+1}}\}$ .

As an example, consider the path  $P = (v_0, v_1, \dots, v_8)$  of Figure 3, where  $r = 3$ . Then,  $D_0 = \{v_0, v_8\}$ ,  $D_1 = \{v_4\}$ ,  $D_2 = \{v_2, v_6\}$  and  $D_3 = \{v_1, v_3, v_5, v_7\}$ . Note that always  $|D_0| = 2$  and for  $j \neq 0$ ,  $|D_j| = 2^{j-1}$ .

For  $j > 0$  and  $v_i \in D_j$ , we define the *parents* of  $v_i$  as  $\{w \mid d_P(v_i, w) = 2^{r-j}\}$ , i.e. the closest vertices that belong to lower-labeled classes. For example, for the path of Figure 3, the parents of  $v_4, v_5, v_6$  are respectively the sets  $\{v_0, v_8\}$ ,  $\{v_4, v_6\}$ ,  $\{v_4, v_8\}$ . Remark that for all  $v \notin \{v_0, v_{2^r}\}$ ,  $v$  has two parents belonging to lower-labeled classes than  $v$  and all vertices between  $v$  and any of its parents belong to higher-labeled classes than  $v$ . We are now ready to define the ‘zig-zag’ pattern.

**Definition 10** (Zig-zag pattern). We call a path graph  $P = (v_0, v_1, \dots, v_{2^r})$ , with distinct integer labels  $\pi$ , *zig-zag*, and we denote it by  $P_r(\pi)$ , if for every  $v \notin \{v_0, v_{2^r}\}$ ,  $v$  has greater label than both its parents  $w_1, w_2$ , i.e.  $\pi(v) > \pi(w_1)$  and  $\pi(v) > \pi(w_2)$ .

<sup>10</sup>We clarify that the result is with respect to the number of players,  $k$ , and does not indicate any lower bound with respect to the number of vertices,  $n$ .

An example of such a path for  $r = 3$  is shown in Figure 3. Our main result of this section is that there exist graphs, high dimensional hypercubes, such that for any order  $\pi$ ,  $P_r(\pi)$  always appears as a *distance preserving* subgraph. Our proof is existential and uses Ramsey theory.




Figure 3: An example of a  $P_3(\pi)$  path. The numbers correspond to the labels.

**Example 11.** In order to give some intuition, we will first use a Ramsey-type result due to Alon et al. [4] to show that, for any  $\pi$ ,  $P_2(\pi)$  appears as a subgraph. Alon et al. [4] showed that for any given integer  $\ell \geq 5$ , any edge-colouring of a sufficiently high dimensional hypercube contains a monochromatic cycle of length  $2\ell$ . Let  $Q_n$  be the hypercube of [4] for  $\ell = 5$ , and notice that it is bipartite i.e.  $Q_n = (A, B, E)$ . For any ordering of vertices of  $Q_n$  we define a colouring as follows: for any edge  $(v, u)$ , with  $v \in A$  and  $u \in B$ , if  $\pi(v) < \pi(u)$ , we paint the edge blue, otherwise we paint it red.

Suppose w.l.o.g. that the monochromatic cycle  $C_{10}$  of length 10 is blue (see also Figure 4 for an illustration). Then, for any  $v \in A \cap V(C_{10})$  (continuous circles), its neighbours in  $C_{10}$  should have *higher* label (dashed circles). The vertices of  $A \cap V(C_{10})$  can be 1-embedded into a cycle  $C_5$  of length 5 (dotted cycle). We appropriately choose three consecutive vertices of  $C_5$ , such that the middle one has higher label than the others in  $\pi$  ((5, 8, 1) in Figure 4). It is not hard to see that such a triplet is guaranteed because  $C_5$  is a cycle. These three vertices with their intermediate ones in  $C_{10}$  form a path isomorphic to  $P_2(\pi)$ ; that path in Figure 4 is the (5, 10, 8, 9, 1).




Figure 4: An example of  $P_2(\pi)$  by using the result of [4].

There are two limitations in using the results of [4] in our proof. a) The induced monochromatic cycle of any length can only be used in order to prove the existence of a zig-zag pattern of length 4 and it doesn't help for paths of higher lengths as required for our lower bound. b) The induced zig-zag pattern is *not* necessarily distance-preserving, because the monochromatic cycle derived by [4] is *not distance preserving*, which is a crucial property for our lower bound to hold. Therefore, in order to overcome those limits, we prove a similar Ramsey-type result, but for a different monochromatic subgraph which is distance preserving and has some special properties (to be described in Section 4.1).

**Proof Overview** The proof is by induction and in the inductive step our starting point is the  $n$ -th dimensional hypercube  $Q_n$ . Given an ordering/labeling  $\pi$  of the vertices of  $Q_n$  we first show

that  $Q_n$  contains a subgraph  $W$  which is isomorphic to a ‘pseudo-hypercube’  $Q_m^2$  ( $m < n$ ) where the labeling of its vertices satisfies a special property (to be described shortly).  $Q_m^2$  is defined by replacing each edge of  $Q_m$  by a 2-edge path (of length two)<sup>11</sup>.

**Definition 12. (Labeling property):** For the subgraph  $W$  we require that all such newly formed 2-edge paths, are  $P_1(\pi)$  paths, i.e. the label of the middle vertex is greater than the labels of the endpoints (Figure 5(a) shows such a labeling).

Next, we contract all such 2-edge paths of  $W$  into single edges, resulting in a graph isomorphic to  $Q_m$ ; this is the hypercube used for the next step. Note that each contracted edge still corresponds to a path in  $Q_n$ . Therefore, after  $r$  recursive steps, each edge corresponds to a  $2^r$  path of  $Q_n$ . Further, note that such a path is a  $P_r(\pi)$  path, due to the labeling property that we preserve at each step. We require that, at the end of the last inductive step,  $Q_m = Q_1$  (a single edge), and (by unfolding it) we show that this edge corresponds to a distance preserving subgraph of the original graph/hypercube. At each step, we have  $m < n$  and the relation between  $n$  and  $m$  is determined by a Ramsey-type argument.

We next describe the basic ingredients that we use to show existence of  $W$ . We apply a coloring scheme to the edges of  $Q_n$  that depends on the vertices’ order.

**Definition 13. (Coloring Scheme):** Consider  $Q_n$  as a bipartite graph  $Q_n = (A, B, E)$ . For any edge  $(v, u)$ , with  $v \in A$  and  $u \in B$ , if the  $v$ ’s label is smaller than  $u$ ’s, we paint the edge blue, otherwise we paint it red.

By a Ramsey-type argument we show that  $Q_n$  has a *monochromatic* subgraph isomorphic to a specially defined graph  $G_m$ ;  $G_m$  is carefully specified in such a way that it contains at least two subgraphs isomorphic to pseudo-hypercubes  $Q_m^2$ . The special property of those two subgraphs is described next.

Let  $H_1$  and  $H_2$  be the two half cubes<sup>12</sup> of  $Q_n$  and let  $V(H_1) = A$  and  $V(H_2) = B$ . Observe that if  $Q_m^2$  is a subgraph of  $Q_n$  then the corresponding  $Q_m$  is an induced subgraph of either  $H_1$  or  $H_2$ . We carefully construct  $G_m$  such that it contains subgraphs  $W_1$  and  $W_2$  isomorphic to  $Q_m^2$ , whose corresponding  $Q_m$ ’s are induced subgraphs of  $H_1$  and  $H_2$ , respectively. The color of  $G_m$  determines which of the  $W_1$  and  $W_2$  will serve as the desired  $W$ . In particular, if the color is blue, then for every edge  $(v, u)$ , with  $v \in V(H_1)$  and  $u \in V(H_2)$ , it should hold that  $v$ ’s label is smaller than  $u$ ’s and therefore the labeling property is satisfied for  $W_1$ ; similarly, if the color is red,  $W_2$  serves as  $W$ .

**Proof Roadmap** The whole proof of the lower bound proceeds in several steps in the following sections. In Section 4.1 we give the formal definition of the subgraph  $G_m$ . Section 4.2 is devoted to show that every 2-edge coloring of a (suitably) high dimensional hypercube contains a monochromatic copy of  $G_m$  (Lemma 15), by using Ramsey theory. Then, in Section 4.3 we show that, for any ordering of the vertices of  $Q_n$ , we can define a special 2-edge-coloring, so that there exists a  $Q_m^2$  subgraph of  $G_m$  that preserves the Labeling property (Lemma 17). At last, in Section 4.4, by a recursive application of the combination of the Ramsey-type result and the coloring, we prove the existence of the *zig-zag* path in high dimensional hypercubes (Theorem 18). We then show how to construct a graph that serves as lower bound for all ordered protocols (Theorem 20). This is done

<sup>11</sup>See  $Q_m^2$  of Definition 16 and Figure 5(a) for an illustration

<sup>12</sup>The two half-cubes of order  $n$  are formed from  $Q_n$  by connecting all pairs of vertices with distance exactly two and dropping all other edges.

by connecting any two edges of the hypercube with a direct edge of appropriate cost, similar to the example in Figure 1(b).

**Definitions and notation on Hypercubes** We denote by  $[r, s]$  (for  $r \leq s$ ) the set of integers  $\{r, r+1, \dots, s-1, s\}$ , but when  $r = 1$ , we simply write  $[s]$ . We follow definitions and notation of [4]. Let  $Q_n$  be the graph of the  $n$ -dimensional hypercube whose vertex set is  $\{0, 1\}^n$ . We represent a vertex  $v$  of  $V(Q_n)$  by an  $n$ -bit string  $x = \langle x_1 \dots x_n \rangle = \langle x_j \rangle_{j=1}^n$ , where  $x_i \in \{0, 1\}$ . By  $\langle xy \rangle$  or  $xy$  we denote the concatenation of an  $r$ -bit string  $x$  with an  $s$ -bit string  $y$ , i.e.  $xy = \langle x_1 \dots x_r y_1 \dots y_s \rangle$ . An edge is defined between any two vertices that differ only in a single bit. We call this bit, *flip-bit*, and we denote it by  $*$ . For example,  $x = \langle 11100 \rangle, y = \langle 11000 \rangle$  are two vertices of  $Q_5$  and  $(x, y) = \langle 11 * 00 \rangle$  is the edge that connects them. The distance between two vertices  $x, y$  is defined by their Hamming distance,  $d(v, u) = |\{j : x_j \neq y_j\}|$ . For a fixed subset of coordinates  $R \subseteq [n]$ , we extend the definition of the distance as follows,

$$d(x, y, R) = \begin{cases} d(x, y), & \text{if } \forall j \in R, x_j = y_j \\ \infty, & \text{otherwise.} \end{cases}$$

We define the *level* of a vertex  $x$  by the number of ‘ones’ it contains,  $w(x) = \sum_{i=1}^n x_i$ . We denote by  $L_i$  the set of vertices of level  $i \in [0, n]$ . We define the *prefix sum* of an edge  $e = (x, y)$ , where the flip-bit is in the  $j$ -th coordinate, by  $p(e) = \sum_{i=1}^{j-1} x_i$ . We represent any ordering  $\pi$  of  $V(Q_n)$ , by labeling the vertices with labels  $1, \dots, 2^n$ , where label  $i$  corresponds to ranking  $i$  in  $\pi$ .

#### 4.1 Description of $G_m$

For a positive integer  $m$ , we define a graph  $G_m = (V_m, E_m)$  that is an induced subgraph of  $Q_{4m}$  on  $V_m = V_1 \cup V_2 \cup V_3 \subseteq V(Q_{4m})$ . A vertex of  $V_1$  is defined by  $2m - 1$  concatenations of pairs  $\langle 01 \rangle$  and  $\langle 10 \rangle$  and a single pair  $\langle 00 \rangle$  that appears in the *second* half of the string. A vertex of  $V_2$  is defined by  $2m$  concatenations of  $\langle 01 \rangle$  and  $\langle 10 \rangle$ . A vertex of  $V_3$  is defined by  $2m - 2$  concatenations of  $\langle 01 \rangle$  and  $\langle 10 \rangle$ , one pair  $\langle 11 \rangle$  that appears on the first half of the string, and one pair  $\langle 00 \rangle$  that appears on the second half. For example, for  $m = 2$ ,  $\langle 01 \ 10 \ 00 \ 10 \rangle \in V_1$ ,  $\langle 01 \ 10 \ 10 \ 10 \rangle \in V_2$ ,  $\langle 01 \ 11 \ 10 \ 00 \rangle \in V_3$ . More formally, let  $A = \{\langle 01 \rangle, \langle 10 \rangle\}$ , then the subsets  $V_1, V_2, V_3$  are defined as follows:

$$\begin{aligned} V_1 &:= V_1(m) = \{\langle a_j b_j \rangle_{j=1}^{2m} \mid \exists i \in [m+1, 2m] \text{ s.t. } \langle a_i b_i \rangle = \langle 00 \rangle \text{ and } \forall j \neq i, \langle a_j b_j \rangle \in A\}, \\ V_2 &:= V_2(m) = \{\langle a_j b_j \rangle_{j=1}^{2m} \mid \forall j, \langle a_j b_j \rangle \in A\}, \\ V_3 &:= V_3(m) = \{\langle a_j b_j \rangle_{j=1}^{2m} \mid \exists i_1 \in [m], \exists i_2 \in [m+1, 2m] \text{ s.t.} \\ &\quad \langle a_{i_1} b_{i_1} \rangle = \langle 11 \rangle, \langle a_{i_2} b_{i_2} \rangle = \langle 00 \rangle \text{ and } \forall j \neq i_1, i_2, \langle a_j b_j \rangle \in A\}. \end{aligned}$$

Observe that  $G_m$  is bipartite with vertex partitions  $V_1$  and  $V_2 \cup V_3$ , as vertices of  $V_1$  belong to level  $2m - 1$ , while vertices of  $V_2 \cup V_3$  to level  $2m$ .

**Lemma 14.** Every pair of vertices  $x, x' \in V_1(m)$  with  $d(x, x', [2m+1, 4m]) = 2$ , have a unique common neighbor  $y \in V_3(m)$ . Also, every pair of vertices  $x, x' \in V_2(m)$ , with  $d(x, x', [2m]) = 2$ , have a unique common neighbor  $y \in V_1(m)$ .

*Proof.* Recall that (by definition) if  $d(x, x', R) \neq \infty$  then  $x, x'$  should coincide in *all*  $R$  coordinates. For the first statement, observe that the premises of the Lemma hold only if there exists  $s \in [m]$  such that  $x_{2s-1}x_{2s} = \langle 10 \rangle$  and  $x'_{2s-1}x'_{2s} = \langle 01 \rangle$  (or the other way around), in which case the required vertex  $y$  from  $V_3(m)$  has  $y_{2s-1}y_{2s} = \langle 11 \rangle$ ; the rest of the bits are the same among  $x, x', y$ . For the

second statement, the premises of the Lemma hold only if there exists an  $s \in [m+1, 2m]$  such that  $x_{2s-1}x_{2s} = \langle 10 \rangle$  and  $x'_{2s-1}x'_{2s} = \langle 01 \rangle$  (or the other way around), in which case the required vertex  $y$  from  $V_1(m)$  has  $y_{2s-1}y_{2s} = \langle 00 \rangle$  and the rest of the bits are the same among  $x, x', y$ .  $\square$

## 4.2 Ramsey-type Theorem

In the following lemma we showed that there exists a sufficiently large hypercube  $Q_n$  such that, no matter how we paint its edges with two colors, it contains a monochromatic copy of  $G_m$ . We note here that the lemma holds for any  $n > g(m)$ , where the value of  $g(m)$  is determined by a Ramsey-type argument. Therefore, the proof is existential and it doesn't provide any bounds on  $g(m)$ . Lemma 15 is only used in Lemma 17, but we believe that it is of independent interest.

**Lemma 15.** For any positive integer  $m$ , and for sufficiently large  $n \geq n_0 = g(m)$ , any 2-edge coloring  $\chi$  of  $Q_n$ , contains a monochromatic copy of  $G_m$ <sup>13</sup>.

*Proof.* The proof follows ideas of Alon et al. [4]. Consider a hypercube  $Q_n$ , with sufficiently large  $n > 6m$  to be determined later, and some arbitrary 2-edge-coloring  $\chi : E(Q_n) \rightarrow \{1, 2\}$ . Let  $E^*$  be the set of edges between vertices of  $L_{4m-1}$  and  $L_{4m}$  (recall that  $L_i = \{v | w(v) = i\}$ ).

Each edge  $e \in E^*$  contains  $4m-1$  1's, a flip-bit represented by  $*$  and the rest of the coordinates are 0. Moreover,  $e$  is uniquely determined by its  $4m$  non-zero coordinates  $R_e \subseteq [n]$  and its prefix sum  $p(e) \in [0, 4m-1]$  (number of 1's before the flip-bit). Therefore, the color  $\chi(e)$  defines a coloring of the pair  $(R_e, p(e))$ , i.e.  $\chi(e) = \chi(R_e, p(e))$ . For each subset  $R \subset [n]$  of  $4m$  coordinates, we denote by  $c(R) = (\chi(R, 0), \dots, \chi(R, 4m-1))$  the color induced by the edge coloring. The coloring of all subsets  $R$  defines a coloring of the complete  $4m$ -uniform hypergraph of  $[n]$ <sup>14</sup> using  $2^{4m}$  colors.

By Ramsey's Theorem for hypergraphs [39], there exists  $n_0 = g(m)$  such that for any  $n \geq n_0$  there exists some subset  $U \subset [n]$  of size  $6m$  such that all  $4m$ -subsets  $R \subset U$  have the same color  $c(R) = c^*$ . Therefore, for every  $4m$ -subsets  $R_1, R_2 \subset U$  and  $p \in [0, 4m-1]$ , it is  $\chi(R_1, p) = \chi(R_2, p) = \chi_p^*$ . Since  $p$  takes  $4m$  values and there are only two different colors, there must exist  $2m$  indices  $p_0, \dots, p_{2m-1} \in [0, 4m-1]$  with the same color  $\chi(R, p_i) = \chi^*$ , for all  $R \subset U$ ,  $|R| = 4m$  and  $i \in [0, 2m-1]$ .

It remains to show that the graph formed by those monochromatic edges contain a copy of  $G_m$ . We will show this by placing the bits of each edge from  $E_m$  (the set of edges of  $G_m$ ) to suitable coordinates of  $[n]$  and filling the rest of the coordinates suitably by zeros and ones. More precisely, we insert blocks of 1's of suitable length among the bits of the edges of  $E_m$ , and all those bits are placed at the coordinates of  $U$ . The rest of the bits ( $n - |U|$ ) are set to zero.

Let  $1^r$  be a string of  $r$  1's and define  $\beta_i = 1^{p_i - p_{i-1} - 1}$  for  $i \in [2m-1]$ ,  $\beta_0 = 1^{p_0}$  and  $\beta_{2m} = 1^{4m-1-p_{2m-1}}$ . For any edge  $e = \langle a_j b_j \rangle_j \in E_m$ , we insert  $\beta_0$  at the beginning of the string, for  $j \in [m]$  we insert  $\beta_j$  between  $a_j$  and  $b_j$  and for  $j \in [m+1, 2m]$  we insert the string  $\beta_j$  after  $b_j$ . The following illustrates these insertions:

$$\underbrace{1 \dots 1}_{p_0} a_1 \underbrace{1 \dots 1}_{p_1 - p_0 - 1} \quad b_1 a_2 \dots a_m \underbrace{1 \dots 1}_{p_m - p_{m-1} - 1} b_m a_{m+1} b_{m+1} \underbrace{1 \dots 1}_{p_{m+1} - p_m - 1}$$

$$a_{m+2} b_{m+2} \dots a_{2m} b_{2m} \underbrace{1 \dots 1}_{4m-1-p_{2m-1}}$$

Recall that each edge of  $E_m$  contains exactly  $2m$  zero bits and  $2m$  non-zero bits (one of which is the flip bit). Also notice that  $\sum_j |\beta_j| = p_0 + \sum_{i=1}^{2m-1} (p_i - p_{i-1} - 1) + 4m - 1 - p_{2m-1} =$

<sup>13</sup>The result could be extended to any (fixed) number of colors, but we need only two for our application.

<sup>14</sup>A  $k$ -uniform hypergraph is a hypergraph such that all its hyperedges have size  $k$ .

$-(2m - 1) + 4m - 1 = 2m$ . Therefore, in total we have  $6m$  bits (same as the size of  $U$ ) and  $4m$  non-zero bits (same as the size of  $R$ ). We place these  $6m$  bits precisely at the coordinates of  $U$ . The rest  $n - 6m$  of the coordinates are filled with zeros.

It remains to show that for such edges the prefix of the flip-bit is *always* one of the  $p_0, \dots, p_{2m-1}$ . This would imply that all these edges are monochromatic. Furthermore, all but  $4m$  coordinates are fixed and the  $4m$  coordinates form exactly the sets  $V_1(m), V_2(m), V_3(m)$ ; therefore, the monochromatic subgraph is isomorphic to  $G_m$ .

For any edge  $e = \langle a_j b_j \rangle_j \in E_m$ , let the flip-bit be at position:

- $a_j$  for  $j \in [m]$ . Its prefix is  $\sum_{i=0}^{j-1} |\beta_i| + (j - 1) = p_{j-1}$ , where the term  $j - 1$  corresponds to the number of pairs  $\langle a_s b_s \rangle$  with  $s < j$ , each of which contributes to the prefix with a single 1.
- $b_j$  for  $j \in [m]$ . Since  $j \leq m$ ,  $a_j = 1$ . Then the prefix equals to  $\sum_{i=0}^j |\beta_i| + (j - 1) + 1 = p_j$ .
- $a_j$  or  $b_j$  for  $j \in [m + 1, 2m]$ . For such  $j$ ,  $\langle a_j b_j \rangle \in \{\langle 0* \rangle, \langle *0 \rangle\}$  and all other pairs belong to  $A$ . Therefore, the prefix is equal to  $\sum_{i=0}^{j-1} |\beta_i| + (j - 1) = p_{j-1}$ .

□

### 4.3 Coloring based on the labels

This part of the proof shows that for any ordering of the vertices of a hypercube  $Q_n$ , there is a 2-edge coloring with the following property: in the monochromatic  $G_m$ , either all the vertices of  $V_1$  or all the vertices of  $V_2$  have neighbors in  $G_m$  with only higher label. This implies a desired labeling property for a  $Q_m^2$  subgraph of  $Q_n$ , the structure of which is defined next.

**Definition 16.** We define  $Q_n^s$  to be a subdivision of  $Q_n$ , by replacing each edge by a path of length  $s$ .  $Q_n^1$  is simply  $Q_n$ . We denote by  $Z(Q_n^s)$  the set of all pairs of vertices  $(x, x')$ , which correspond to edges of  $Q_n$ ;  $P(x, x')$  is the corresponding path in  $Q_n^s$ . For every  $(x, x') \in Z(Q_m^2)$ , we denote by  $\theta(x, x')$  the middle vertex of  $P(x, x')$ .




Figure 5: Examples of (a)  $Q_3^2$  and (b)  $Q_2^4$ . The labels on the nodes are examples of the labeling property, (a) after one inductive step, (b) after two inductive steps.

In the next lemma we show that for any ordering of the vertices of  $Q_n$ , there exists a subgraph isomorphic to  $Q_m^2$ , such that the ‘middle’ vertices have higher label than their neighbors (Labeling Property). This lemma is only used in Theorem 18.

**Lemma 17.** For any positive integer  $m$ , for all  $n \geq n_0 = g(m)$  and for any ordering  $\pi$  of  $V(Q_n)$ , there exists a subgraph  $W$  of  $Q_n$  that is isomorphic to  $Q_m^2$ , such that, for every  $(x, x') \in Z(W)$ , it is  $\pi(\theta(x, x')) > \max\{\pi(x), \pi(x')\}$ .

*Proof.* Choose a sufficiently large  $n \geq n_0 = g(m)$  as in Lemma 15. Partition the vertices of  $Q_n$  into sets  $\mathcal{O}, \mathcal{E}$  of vertices of odd and even level, respectively. We color the edges of  $Q_n$  as follows. For every edge  $e = (z, z')$  with  $z \in \mathcal{O}$  and  $z' \in \mathcal{E}$ , if  $\pi(z) < \pi(z')$ , then paint  $e$  blue. Otherwise paint it red. Therefore, for every blue edge, the endpoint in  $\mathcal{O}$  has smaller label than the endpoint in  $\mathcal{E}$ . The opposite holds for any red edge.

Lemma 15 implies that  $Q_n$  contains a monochromatic copy (blue or red) of  $G_m$ . Recall that this monochromatic subgraph of  $Q_n$  is bipartite between vertices of levels  $L_{4m-1}$  and  $L_{4m}$  and that  $V_1 \subset L_{4m-1} \subset \mathcal{O}$  and  $V_2 \cup V_3 \subset L_{4m} \subset \mathcal{E}$ . Let  $R \subset [n]$  be the subset of the  $4m$  coordinates that correspond to vertices of  $G_m$ . Also let  $R_1$  and  $R_2$  be the subsets of the first  $2m$  and the last  $2m$  coordinates of  $R$ , respectively.

First suppose that the subgraph isomorphic to  $G_m$  is blue. An immediate implication of our coloring is that for every edge  $(z, z') \in E_m$  with  $z \in V_1$ ,  $z' \in V_2 \cup V_3$  it must be  $\pi(z) < \pi(z')$ . Fix a  $2m$ -bit string  $s$  that corresponds to a permissible bit assignment to the  $R_2$  coordinates of some vertex in  $V_1$  (see Section 4.1). Define  $W_s$  as the subset of vertices of  $V_1$  where the  $R_2$  coordinates are set to  $s$ . Recall that each of the first  $m$  pairs  $\langle a_j b_j \rangle$ ,  $j \in [m]$ , of a vertex  $z \in W_s$ , may take any of the two bit assignments  $\langle 01 \rangle$  and  $\langle 10 \rangle$ . Hence,  $|W_s| = 2^m$ .

Observe that we can embed  $W_s$  into  $Q_m$  with distortion<sup>15</sup> 1 and scaling factor  $1/2$ , by mapping the first  $m$  pairs of bits into single bits; map  $\langle 01 \rangle$  to 1 and  $\langle 10 \rangle$  to 0. Every two vertices with distance  $d$  in  $Q_m$ , have distance  $2d$  in  $Q_n$ . For every  $x, x' \in W_s \subset V_1$  with  $d(x, x') = 2$ , it holds that  $d(x, x', R_2) = 2$ , since  $x, x'$  have the same  $R_2$  coordinates. Lemma 14 implies that there exists  $y = \theta(x, x') \in V_3$ , such that  $d(x, y) = d(x', y) = 1$ , and therefore,  $\pi(y) > \max\{\pi(x), \pi(x')\}$ . Take the union  $Y = \cup_y$  of all such vertices  $y$ , then  $W_s \cup Y$  induces a subgraph  $W$  isomorphic to  $Q_m^2$ , that fulfills the labeling requirements.

The case of  $G_m$  being red is similar. We focus only on the vertices of  $V_2$ . Fix now a  $2m$ -bit string  $s$  that corresponds to a permissible bit assignment of the  $R_1$  coordinates of a vertex in  $V_2$ . Define  $W_s$  as the subset of vertices of  $V_2$  where the  $R_1$  coordinates are set to  $s$ . Similarly, we can embed  $W_s$  into  $Q_m$  with distortion 1 and scaling factor  $1/2$ .

For every  $x, x' \in W_s \subset V_2$  with  $d(x, x') = 2$ , where the  $R_1$  coordinates are fixed to  $s$ , Lemma 14 implies that there exists  $y = \theta(x, x') \in V_1$ , such that  $d(x, y) = d(x', y) = 1$ , and therefore,  $\pi(y) > \max\{\pi(x), \pi(x')\}$ . Take the union  $Y = \cup_y$  of all such vertices  $y$ , then  $W_s \cup Y$  induces a subgraph  $W$  isomorphic to  $Q_m^2$ , that fulfills the labeling requirements.  $\square$

#### 4.4 Lower Bound Construction

Now we are ready to prove the main theorem of this section.

**Theorem 18.** For every positive integer  $r$ , and for sufficiently large  $n = n(r)$ , there exists a graph  $Q_n$  such that, for every ordering  $\pi$  of its vertices, it contains a zig-zag distance preserving path  $P_r(\pi)$ .

*Proof.* Let  $g$  be a function as in Lemma 15. We recursively define the sequence  $n_0, n_1, \dots, n_r$ , such that  $n_r = 1$  and  $n_{i-1} = g(n_i)$ , for  $i \in [r]$ . We will show that  $Q_{n_0}$  ( $n_0 = n(r)$ ) is the graph we are looking for.

*Claim 19.* For every  $i \in [0, r]$ , and for any vertex ordering  $\pi$  of  $Q_{n_0}$ ,  $Q_{n_0}$  contains a subgraph isomorphic to  $Q_{n_i}^2$ , such that for every  $(x, x') \in Z(Q_{n_i}^2)$ ,  $P(x, x')$  is a zig-zag path  $P_i(\pi)$ .

<sup>15</sup>We give the definition of distortion: Let  $(X, \rho)$  and  $(Y, \sigma)$  be metric spaces, and let  $D \geq 1$  be a real number. A map  $f : X \rightarrow Y$  is said to have *distortion* at most  $D$  if there exists a real number  $r > 0$  (which is called *scaling factor*) such that for all  $x, y \in X$  it holds that  $r \cdot \rho(x, y) \leq \sigma(f(x), f(y)) \leq D \cdot r \cdot \rho(x, y)$ .

*Proof.* The proof is by induction on  $i$ . As a base case,  $Q_{n_0}^{2^0} = Q_{n_0}$  is the graph itself. An edge is trivially a path  $P_0(\pi)$ , for any  $\pi$ . Suppose now that  $Q_{n_0}$  contains a subgraph isomorphic to  $Q_{n_i}^{2^i}$ , for some  $i < r$ , such that for every  $q \in Z(Q_{n_i}^{2^i})$ ,  $P(q)$  is a zig-zag path  $P_i(\pi)$ . It is sufficient to show that  $Q_{n_i}^{2^i}$  contains a subgraph isomorphic to  $Q_{n_{i+1}}^{2^{i+1}}$ , such that for every  $q \in Z(Q_{n_{i+1}}^{2^{i+1}})$ ,  $P(q)$  is a zig-zag path  $P_{i+1}(\pi)$ .

For every  $(x, x') \in Z(Q_{n_i}^{2^i})$ , if we replace  $P(x, x')$  with a direct edge  $e = (x, x')$ , the resulting graph is a copy of  $Q_{n_i}$ . Applying Lemma 17 on  $Q_{n_i}$ , guarantees the existence of a subgraph  $W$  isomorphic to  $Q_{n_{i+1}}^2$  ( $n_i = g(n_{i+1})$ ), where for every  $(y, y') \in Z(W)$ ,  $\pi(\theta(y, y')) > \max\{\pi(y), \pi(y')\}$ . Each of the edges  $(y, \theta(y, y'))$  and  $(y', \theta(y, y'))$  of  $Q_{n_{i+1}}^2$  are replaced by a path  $P_i(\pi)$  in  $Q_{n_i}^{2^i}$ . Therefore,  $W$  is a copy of  $Q_{n_{i+1}}^{2^{i+1}}$ , with  $P(y, y')$  being a zig-zag path  $P_{i+1}(\pi)$ .  $\square$

We now argue that the resulting  $P_r(\pi)$  is a distance preserving path. Our analysis indicates a sequence of hypercubes  $Q_{n_0}, Q_{n_1}, \dots, Q_{n_r}$ . Recall that in Lemma 17, in order to get  $Q_{n_{i+1}}$  from  $Q_{n_i}$  we mapped  $\langle 01 \rangle$  to 1 and  $\langle 10 \rangle$  to 0 and the vertices of  $Q_{n_{i+1}}$  did not differ in any other bit but the ones we mapped. Consider now the two vertices  $x, x'$  of  $Q_{n_r} = Q_1$  with bit-strings  $\langle 0 \rangle$  and  $\langle 1 \rangle$ , respectively. Their Hamming distance in their original bit representation (in  $Q_{n_0}$ ) should be  $2^r$ , the same with their distance in  $P_r(\pi)$ .

For instance, for  $r = 4$ , Table 1 shows the bit sequences in  $Q_{n_3}, Q_{n_2}, Q_{n_1}$  and  $Q_{n_0}$  that correspond to the bits  $\langle 0 \rangle$  and  $\langle 1 \rangle$  of the vertices  $x, x'$  of  $Q_{n_4} = Q_1$ . In any  $Q_{n_i}$ , both bit sequences

Table 1: Example of unfolding the bit mapping.

	$Q_{n_4}$	$Q_{n_3}$	$Q_{n_2}$	$Q_{n_1}$	$Q_{n_0}$
$x$	$\langle 0 \rangle$	$\langle 10 \rangle$	$\langle 0110 \rangle$	$\langle 10010110 \rangle$	$\langle 0110100110010110 \rangle$
$x'$	$\langle 1 \rangle$	$\langle 01 \rangle$	$\langle 1001 \rangle$	$\langle 01101001 \rangle$	$\langle 1001011001101001 \rangle$

occupy exactly the same coordinates. The rest of the coordinates of  $x, x'$  are occupied by identical bits in all bit representations. Therefore,  $d_{Q_{n_0}}(x, x') = 16 = 2^r$ .

Moreover, if any two vertices of  $P_r(\pi)$  are closer in  $Q_{n_0}$  than in  $P_r(\pi)$ , then this would contradict the fact that  $d_{Q_{n_0}}(x, x') = 2^r$ .  $\square$

Finally we extend  $Q_n$  so that for any order  $\pi$  of its vertices, a path  $P_r(\pi)$  exists along with the shortcuts similar to the example of Figure 1(b).

**Theorem 20.** Any ordered universal protocol on undirected graphs admits a PoA of  $\Omega(\log k)$ , where  $k$  is the number of activated vertices.

*Proof.* Let  $k = 2^r + 1$  for some positive integer  $r$ . From Theorem 18, we know that for any vertex ordering  $\pi$  of  $Q_{n(r)}$  there is a distance preserving path  $P_r(\pi)$ .

We use  $Q_{n(r)}$  as a basis to construct the weighted graph  $\tilde{Q}_{n(r)}$  with vertex set  $V(\tilde{Q}_{n(r)}) = Q_{n(r)} \cup \{t\}$ , where  $t$  is the designated root. We connect every pair of vertices  $x, y$  with a direct edge of cost  $c_e = 2^r$ , if  $t$  is one of its endpoints, otherwise its cost is  $c_e = d_{Q_{n(r)}}(x, y)$  (similar to Figure 1(b)).

The adversary selects to activate the vertices of  $P_r(\pi)$ , and the lower bound follows; in the NE the players choose their direct edges to connect with one of their parents (see at the beginning of Section 4 for the term “parent”).  $\square$

## 5 Lower Bound for all universal protocols

In this section, we exhibit a graph metric for which no universal cost-sharing protocol admits a PoA better than  $\Omega(\log k)$ . Due to the characterization of [23], we can restrict ourselves to generalized weighted Shapley protocols (GWSPs). We refer the reader to Section 2 for the definition. We remind the reader that ordered and Shapley protocols are the two extreme cases of the GWSPs. So, we will distinguish between protocols that are closer to ordered protocols and protocols that are closer to Shapley protocols.

We follow the notation of [23], and for the sake of self-containment we include here the most related definitions and lemmas.

### 5.1 Cost-Sharing Preliminaries

A strictly positive function  $f : 2^N \rightarrow \mathbb{R}^+$  is an *edge potential* on  $N$ , if it is strictly increasing, i.e. for every  $R \subset S \subseteq N$ ,  $f(R) < f(S)$ , and for every  $S \subseteq N$ ,

$$\sum_{i \in S} \frac{f(S) - f(S \setminus \{i\})}{f(\{i\})} = 1.$$

For simplicity, instead of  $f(\{i\})$ , we write  $f(i)$ . A cost-sharing protocol is called *potential-based*, if it is defined by assigning to each edge of cost  $c$ , the cost-sharing method  $\xi$ , where for every  $S \subseteq N$  and  $i \in S$ ,

$$\xi(i, S) = c \cdot \frac{f(S) - f(S \setminus \{i\})}{f(i)}.$$

Let  $\Xi_1$  and  $\Xi_2$  be two cost-sharing protocols for *disjoint* sets of vertices  $U_1$  and  $U_2$ , with methods  $\xi_1$  and  $\xi_2$ , respectively. The *concatenation* of  $\Xi_1$  and  $\Xi_2$  is the cost sharing protocol  $\Xi$  of the set  $U_1 \cup U_2$ , with method  $\xi$  defined as

$$\xi(i, S) = \begin{cases} \xi_1(i, S \cap U_1) & \text{if } i \in U_1, \\ \xi_2(i, S) & \text{if } S \subseteq U_2, \\ 0 & \text{otherwise.} \end{cases}$$

Note that the concatenation of two protocols for disjoint sets of vertices defines an order among these two sets. The GWSPs are concatenations of potential-based protocols.

The following two lemmas from [23] give some bounds on the players' cost-shares given a relation of their edge potential values. Both lemmas will be used in Theorem 26 and more specifically Lemma 21 will be used for the *Shapley-like* protocols and Lemma 22 for the *ordered-like* ones.

**Lemma 21.** (Lemma 4.10 of [23]). Let  $f$  be an edge potential on  $N$  and  $\xi$  the induced (by  $f$ ) cost-sharing method, for unit costs. For  $k \geq 1$  and a constant  $\alpha$ , with  $1 \leq \alpha^{2k} \leq 1 + k^{-3}$ , let  $S \subseteq N$  be a subset of vertices with  $f(i) \leq \alpha f(j)$ , for every  $i, j \in S$ . If  $|S| \leq k$ , then for any  $i, j \in S$ ,

$$\xi(i, S) \leq \alpha(\xi(j, S) + 2k^{-2}).$$

**Lemma 22.** (Lemma 4.11 of [23]). Let  $f$  be an edge potential on  $N$ , and  $\xi$  be the cost-sharing method induced by  $f$ , for unit cost. For any two vertices  $i, j \in N$ , such that  $f(i) \geq \beta f(j)$ ,  $\xi(i, \{i, j\}) \geq \beta/(\beta + 1)$ , and for every  $S \supseteq \{i, j\}$ ,  $\xi(j, S) \leq 1/(\beta + 1)$ .

*Remark 23.* Consider a GWSP with  $\Sigma = (U_1, \dots, U_h)$  as the ordered partition of the players. We remark that Lemma 22 refers to players that belong to the same set  $U_a$ . However, the same bounds hold for any two players  $i, j \in N$  that belong to different sets  $U_a, U_b$  respectively, i.e. for  $a < b$ ,  $\xi(i, \{i, j\}) = 1 > \beta/(\beta + 1)$  and for every  $S \supseteq \{i, j\}$ ,  $\xi(j, S) = 0 < 1/(\beta + 1)$ .

## 5.2 Lower Bound

The following two technical lemmas will be used in our main theorem.

**Lemma 24.** Let  $X$  be a finite set of size  $msr^2$ , and  $X_1, \dots, X_m$  be a partition of  $X$ , with  $|X_i| = sr^2$ , for all  $i \in [m]$ . Then, for any coloring  $\chi$  of  $X$  such that *no more* than  $r$  elements have the same color, there exists a *rainbow* subset  $S \subset X$  (i.e.  $\chi(v) \neq \chi(u)$  for all  $v, u \in S$ ), with  $|S \cap X_i| = s$  for every  $i \in [m]$ .

*Proof.* Given the partition  $X_1, \dots, X_m$  of  $X$  and the coloring  $\chi$ , we construct a bipartite graph  $G = (A, B, E)$ , where  $A$  is the set of colors used in  $\chi$ . For every  $X_i$  we create a set  $B_i$  of size  $s$ ; then  $B = \cup B_i$ . If color  $j$  is used in  $X_i$ , we add an edge  $(j, l)$  for all  $l \in B_i$ .

Each color  $j \in A$  appears in at most  $r$  distinct  $X_i$  sets, and since for each  $X_i$  there are  $s$  vertices ( $B_i$ ), the degree of  $j$  is at most  $rs$ . On the other hand, each  $X_i$  has size  $r^2s$  and hence, it has at least  $rs$  different colors. Therefore, the degree of each vertex of  $B$  is at least  $rs$ .

Consider any set  $R \subseteq B$ , and let  $E(R)$  be the set of edges with at least one endpoint in  $R$ . If  $N(R)$  denotes the set of neighbors of  $R$ , observe that  $E(R) \subseteq E(N(R))$ . By using the degree bound on vertices of  $B$ ,  $|E(R)| \geq rs|R|$  and by using the degree bound on vertices of  $A$ ,  $|E(N(R))| \leq rs|N(R)|$ . Therefore,  $|R| \leq |N(R)|$ . By Hall's Theorem there exists a matching which covers every vertex in  $B$ . Each vertex in  $B_i$  is matched with a distinct color and therefore in each  $X_i$  there exists a subset with at least  $s$  elements with distinct colors; let  $W_i$  be such a subset with exactly  $s$  elements. In addition the colors in different  $W_i$  subsets should be distinct by the matching. Then,  $S = \cup W_i$ .  $\square$

**Lemma 25.** Let  $X = (X_1, \dots, X_m)$  be a partition of  $[m^2]$ , with  $|X_i| = m$ , for all  $i \in [m]$ . Then, there exists a subset  $S \subset [m^2]$  with exactly one element from each subset  $X_i$ , such that no two distinct  $x, y \in S$  are consecutive, i.e. for every  $x, y \in S$ ,  $|x - y| \geq 2$ .

*Proof.* For every  $i$ , let  $X_i = \{x_{i1}, \dots, x_{im}\}$ . W.l.o.g we can assume that the  $x_{ij}$ 's are in increasing order with respect to  $j$  and in addition that  $X_i$ 's are sorted such that  $x_{ii} < x_{ji}$ , for all  $j > i$  (otherwise rename the elements recursively to fulfill the requirement). Then, it is not hard to see that  $S = \{x_{kk} | k \in [m]\}$  can serve as the required set.  $\square$

Now we proceed with the main theorem of this section. We create a graph where every GWSP has high PoA. At a high level, we construct a high dimensional hypercube with sufficiently large number of potential players at each vertex (by adding many copies of each vertex connected via zero-cost edges). Moreover, we add shortcuts among the vertices of suitable costs and we connect each vertex with  $t$  via two parallel links with costs that differ by a large factor (see Figure 6). If the protocol induces a large enough set of potential players with *Shapley-like* values in some vertex, then it is a NE that all these players follow the most costly link to  $t$ . Otherwise, by using Lemmas 24 and 25 we show that there exists a set of potential players with *ordered-like* values, one at each vertex of the hypercube. Then, by using the results of Section 4, there exists a path where the vertices are *zig-zag*-ordered.

The separation into these two extreme cases was first used in [23]. The crucial difference, is that for their problem the protocol is specified independently of the underlying graph, and therefore the adversary knows the case distinction (ordered or Shapley) and bases the lower bound construction on that. However, our problem requires more work as the graph should be constructed *in advance*, and should work for both cases.

**Theorem 26.** There exist graph metrics, such that the PoA of any universal cost-sharing protocol is at least  $\Omega(\log k)$ , where  $k$  is the number of activated vertices.

*Proof.* Let  $k = 2^{r-1} + 1$  be the number of activated vertices with  $r \geq 4$ , (so  $k \geq 9$ ).

**Graph Construction.** We use as a base of our lower bound construction, a hypercube  $Q := Q_n$ , with edge costs equal to 1 and  $n = n(r)$  as in Theorem 18. Based on  $Q$ , for  $M = 16k^{12}2^{3n}$  we construct the following network with  $N = 2^n M$  vertices, plus the designated root  $t$ . We add to  $Q$  direct edges/shortcuts as follows: for every two vertices  $v, u$  of distance  $2^j$ , for  $j \in [r]$ , we add an edge/shortcut,  $(v, u)$ , with cost equal to  $\hat{c}_j = 2^j \left(\frac{k-1}{k}\right)^j = \Omega(2^j)$ . Moreover, for every vertex  $v_q$  of  $Q$ , we create  $M - 1$  new vertices, each of which we connect with  $v_q$  via a zero-cost edge. Let  $V_q$  be the set of these vertices (including  $v_q$ ). Finally, we add a root  $t$ , which we connect with every vertex  $v_q$  of  $Q$ , via two edges  $e_{q1}$  and  $e_{q2}$ , with costs  $2k$  and  $2k \cdot k/6$ , respectively. We denote this new network by  $Q^*$  (see Figure 6).




Figure 6: An example of  $Q^*$  for  $Q_2$  as the base hypercube.

We will show that any GWSP for  $Q^*$  has PoA of  $\Omega(\log k)$ . Any GWSP can be described by concatenations of potential-based cost-sharing protocols  $\Xi_1, \dots, \Xi_h$  for a partition of the  $V(Q^*)$  into  $h$  subsets  $U_1, \dots, U_h$ , where  $\Xi_j$  is induced by some edge potential  $f_j$ . Following the analysis of Chen, Roughgarden and Valiant [23], we scale the  $f_j$ 's such that for every  $i, j$ ,  $f_j(i) \geq 1$ . For nonnegative integers  $s$  and for  $\alpha = (1 + k^{-3})^{\frac{1}{2k}}$ , we form subgroups of vertices  $A_{js}$ , for each  $U_j$ , as  $A_{js} = \{i \in U_j : f_j(i) \in [\alpha^s, \alpha^{s+1}]\}$  (note that some of  $A_{js}$ 's may be empty).

The adversary proceeds in two cases, depending on the intersection of the  $A_{js}$ 's with the  $V_q$ 's. **Shapley-like cost-sharing.** Suppose first that there exist  $A_{js}$  and  $V_q$  such that  $|A_{js} \cap V_q| \geq k$ , and take a subset  $R \subseteq A_{js} \cap V_q$  with exactly  $k$  vertices. The adversary will request precisely the set  $R$ . We argue that there is a NE where all players follow the edge  $e_{q2}$ , with cost  $2k \cdot k/6$ .

Budget-balance implies that there exists some player  $i^* \in R$  who is charged at most  $1/k$  proportion of the cost. Moreover, Lemma 21 implies that, all  $i \in R$  are charged at most  $\alpha(1/k + 2k^{-2}) \leq 2 \cdot (3/k) = 6/k$  proportion of the cost. Therefore, no player's share is more than  $2k$  and any alternative path would cost at least  $2k$ . However, the optimum solution is to use the parallel link  $e_{q1}$  of cost  $2k$ . Hence, the PoA is  $\Omega(k)$  for this case.

**Ordered-like cost-sharing.** If there is no such  $R$  with at least  $k$  vertices, then  $|A_{js} \cap V_q| < k$  for all  $j, s$  and  $q$ , which means that each  $A_{js}$  has size of at most  $k2^n$ . For every  $j \in [h]$ , we group consecutive sets  $A_{js}$  (starting from  $A_{j0}$ ) into sets  $B_{jl}$ , such that each  $B_{jl}$ , (except perhaps from the last one), contains *exactly*  $4k^5$  *nonempty*  $A_{js}$ 's. The last  $B_{jl}$  contains at most  $4k^5$  *nonempty*  $A_{js}$  sets. Consider the lexicographic order among  $B_{jl}$ 's, i.e.  $B_{jl} < B_{j'l'}$  if either  $j < j'$  or  $j = j'$  and  $l < l'$ . Rename these sets based on their total order as  $B_i$ 's. The size of each  $B_i$  is at most  $4k^6 2^n$ .

Now we apply Lemma 24 on the set  $V(Q^*) \setminus \{t\} = \cup_q V_q$ , for  $r = 4k^6 2^n$  and  $s = m = 2^n$ , by considering the subsets  $V_q$  as the partition of  $V(Q^*) \setminus \{t\}$  (recall that  $|V_q| = M = r^2 s$ ). As a

coloring scheme, we color all the vertices of each  $B_i$  with the same color and use different colors among the sets  $B_i$ . Lemma 24 guarantees that for each  $V_q$  there exists  $V'_q \subset V_q$  of size  $2^n$ , such that every  $v \in \cup_q V'_q = V'$  belongs to a distinct  $B_i$ .

The order of  $B_i$ 's suggests an order of the vertices of  $V'$ . Since the  $V'_q$ 's form a partition of  $V'$ , Lemma 25 guarantees the existence of a subset  $C \subset V'$ , such that  $C$  contains exactly one vertex from each  $V'_q$  and there are no consecutive vertices in  $C$ . This means that  $C$  contains exactly one vertex from each set  $V_q$  and all these vertices belong to *different* and *non-consecutive* sets  $B_i$ .

To summarize, so far we know that:

- (i) for any pair of vertices  $v, u \in C$ , either  $v$  and  $u$  come from different  $U_j$ 's or their  $f_j(v)$  and  $f_j(u)$  values differ by a factor of at least  $\alpha^{4k^5} \geq 8k+1$  (since there exist at least  $4k^5$  nonempty sets  $A_{j_s}$  between the ones that  $v$  and  $u$  belong to).
- (ii)  $C$  is a copy of  $Q_n$  (by ignoring zero-cost edges).

Let  $\pi$  be the order of vertices of  $C$  (recall that they are ordered according to the  $B_i$ 's they belong to). Theorem 18 guarantees that there always exists at least one distance preserving path  $P_r(\pi)$  (see Definition 10). Let  $S$  be the vertices of  $P_r(\pi)$  excluding the last class  $D_r$  (see Definition 9). The adversary will activate precisely the set  $S$  ( $|S| = k$ ). It remains to show that there exists a NE, the cost of which is a factor of  $\Omega(\log k)$  away from optimum. We will refer to these vertices as  $S = \{s_1, s_2, \dots, s_k\}$ , based on their order  $\pi$ , from smaller label to larger, and let player  $i$  be associated with  $s_i$ .

Let  $\mathcal{P}^*$  be the class of strategy profiles  $\mathbf{P} = (P_1, \dots, P_k)$  which are defined as follows:

- $P_1 = e_{11}$  and  $P_2 = (s_1, s_2) \cup P_1$ , where  $(s_1, s_2)$  is the shortcut edge between  $s_1$  and  $s_2$ .
- From  $i = 3$  to  $k$ , let  $s_\ell$  be any of  $s_i$ 's parents in the class hierarchy (we refer the reader to the beginning of Section 4); then  $P_i = (s_i, s_\ell) \cup P_\ell$ , where  $(s_i, s_\ell)$  is the shortcut edge between  $s_i$  and  $s_\ell$ .

We show in Claim 27 that there exists a strategy profile  $\mathbf{P}^* \in \mathcal{P}^*$  which is a NE.  $\mathbf{P}^*$  has cost:

$$c(\mathbf{P}^*) = c(e_{11}) + \hat{c}_r + \sum_{j=1}^{r-1} |D_j| \cdot \hat{c}_{r-j} = \Omega(2^r) + \Omega(2^r) + \sum_{j=1}^{r-1} 2^{j-1} \cdot \Omega(2^{r-j}) = \Omega(r2^r).$$

However, there exists the solution  $P_r(\pi) \cup e_{11}$ , which has cost of  $O(2^r)$ . Therefore, the PoA is  $\Omega(r) = \Omega(\log k)$ .

*Claim 27.* There exists  $\mathbf{P}^* \in \mathcal{P}^*$  which is a NE.

*Proof.* We prove the claim by using better-response dynamics. Note that any GWSP induces a potential game for which better-response dynamics always converge to a NE (see [23, 37]). We start with some  $\mathbf{P}_1 \in \mathcal{P}^*$  and we prove that, after a sequence of players' *best*-responses, we end up in  $\mathbf{P}_2 \in \mathcal{P}^*$ . Proceeding in a similar way we eventually converge to  $\mathbf{P}^*$ , which is the required NE.

We next argue that for any  $\mathbf{P} \in \mathcal{P}^*$ , players 1 and 2, have no incentive to deviate from  $P_1$  (argument (a)) and  $P_2$  (argument (b)), respectively. We further show that, given any strategy profile  $\hat{\mathbf{P}}$ , there exists some  $\mathbf{P} \in \mathcal{P}^*$  such that: for every player  $i \notin \{1, 2\}$ , if  $\mathbf{P}^i = (P_1, \dots, P_{i-1}, \hat{P}_{i+1}, \dots, \hat{P}_k)$  are the strategies of the other players,  $i$  prefers  $P_i$  to  $\hat{P}_i$  (arguments (c)-(e)). We define the desired  $\mathbf{P}$  recursively starting from  $\hat{\mathbf{P}}$  as follows:  $P_1 = e_{11}$ ,  $P_2 = (s_1, s_2) \cup P_1$  and from  $i = 3$  to  $k$ ,  $P_i \in A_i = \arg \min_{P'_i} \{c_i(\mathbf{P}^i, P'_i) \mid \exists (P'_{i+1}, \dots, P'_k) \text{ s.t. } (P_1, \dots, P_{i-1}, P'_i, \dots, P'_k) \in \mathcal{P}^*\}$ . If  $\hat{P}_i \in A_i$  then we set  $P_i = \hat{P}_i$ , otherwise we choose a path from  $A_i$  arbitrarily.

We first give some bounds on players' shares.

1. Let  $R \subseteq S$  be any set of players that use some edge  $e$  of cost  $c_e$  and let  $i$  be the one with the smallest label. The total share of players  $R \setminus \{i\}$  is upper bounded by  $\sum_{i=1}^{|R|-1} \frac{1}{(8k+1)^{i+1}} \cdot c_e < \frac{c_e}{8k}$  (Lemma 22 and Remark 23). Moreover,  $i$ 's share is at least  $\frac{8k+1}{8k+2} > \frac{8k-1}{8k} c_e$ .
2. The total cost of any  $P_i$  under  $\mathbf{P}^i$ , is at most  $8k$ . This is true because, for every player  $i'$  with  $i' \leq i$ , the first edge of  $P_{i'}$  is a shortcut to reach one of  $s_{i'}$ 's parents, with cost at most  $2^{r-j}$ , where  $D_j$  is the hierarchical class that  $s_{i'}$  belongs to (we refer the reader to the beginning of Section 4 for the definition of classes). Therefore, the cost of  $P_i$  is at most  $2k + \sum_{l=0}^{r-1} 2^{r-l} < 8k$ .
3. By combining the above two arguments, under  $\mathbf{P}^i$ , the total share of player  $i$  for the edges of  $P_i$  at which she is not the first according to  $\pi$ , is at most  $\frac{1}{8k} \cdot 8k \leq 1$ .

Here, we give the arguments for players 1 and 2.

- (a) The share of player 1 under  $\mathbf{P} \in \mathcal{P}^*$  is at most  $2k$  and any other path would incur a cost strictly greater than  $2k$ .
- (b) The share of player 2 under  $\mathbf{P} \in \mathcal{P}^*$  is at most  $2^r + 1 = 2k - 1$  (argument 3), whereas if she doesn't connect through  $s_1$ , her share would be at least  $2k$ . Moreover, if she connects to  $t$  through  $s_1$  but by using any other path  $p$  rather than the shortcut  $(s_1, s_2)$ , the cost of each edge of  $p$  is  $2^w \left(\frac{k-1}{k}\right)^w$  for some  $w \leq r-1$ , so it holds that it is at least  $2^w \left(\frac{k-1}{k}\right)^{r-1}$ . Given that the distance between  $s_1$  and  $s_2$  in  $P_r(\pi)$  is  $2^r$ , the total cost of  $p$  is at least  $2^r \left(\frac{k-1}{k}\right)^{r-1}$ . Player 2 is first according to  $\pi$  at  $p$  and by argument 1, her share is at least  $2^r \frac{8k-1}{8k} \left(\frac{k-1}{k}\right)^{r-1} > \hat{c}_r$ .

We next give the required arguments in order to show that  $P_i$  is a best response for player  $i \neq \{1, 2\}$  under  $\mathbf{P}^i$ . In the following, let  $s_i \in D_j$  and let  $s_\ell$  be the parent of  $s_i$  such that  $P_i = (s_i, s_\ell) \cup P_\ell$ . Also let  $s_{i'}$  be the predecessor of  $s_i$ , according to  $\pi$ , that is first met by following  $\hat{P}_i$  from  $s_i$  to  $t$ .

- (c) Suppose that  $s_{i'} = s_\ell$ .
  - Assume that  $\hat{P}_i$  doesn't use the shortcut  $(s_i, s_\ell)$ . The subpath of  $\hat{P}_i$  from  $s_i$  to  $s_\ell$  contains edges at which  $i$  is first according to  $\pi$  of total cost at least  $2^{r-j} \left(\frac{k-1}{k}\right)^{r-j-1}$ . By argument 1, her share is at least  $2^{r-j} \frac{8k-1}{8k} \left(\frac{k-1}{k}\right)^{r-j-1} > \hat{c}_{r-j}$ .
  - Assume that  $\hat{P}_i$  doesn't use  $P_\ell$ . The subpath of  $\hat{P}_i$  from  $s_\ell$  to  $t$  contains edges at which  $i$  is first according to  $\pi$  of total cost at least  $\hat{c}_1$  (the minimum distance between two activated vertices). By argument 1, her share is at least  $2 \frac{k-1}{k} \frac{8k-1}{8k} > 1$ , for  $k \geq 3$ , where 1 is at most her share for  $P_\ell$  (argument 3).

In both cases,  $c_i(\mathbf{P}^i, P_i) < c_i(\mathbf{P}^i, \hat{P}_i)$ .

- (d) Suppose that  $s_{i'}$  is  $s_i$ 's other parent. If  $\hat{P}_i \neq (s_i, s_{i'}) \cup P_{i'}$ , the above arguments still hold and so  $c_i(\mathbf{P}^i, P_i) < c_i(\mathbf{P}^i, \hat{P}_i)$ . Otherwise, by the definition of  $P_i$ , either  $P_i = \hat{P}_i$ , or  $c_i(\mathbf{P}^i, P_i) < c_i(\mathbf{P}^i, \hat{P}_i)$ .
- (e) Suppose that  $s_{i'}$  is not a parent of  $s_i$ . Player  $i$ 's share in  $P_i$  is at most  $\hat{c}_{r-j}$  for her first edge/shortcut and at most 1 for the rest of her path (argument 3). Let  $p$  be the path from  $s_i$  to  $s_{i'}$  in  $\hat{P}_i$ . We consider three cases for  $p$  and we will show that the cost-share of player  $i$  for  $p$  is at least  $\hat{c}_{r-j} + 1$ . Recall that between  $s_i$  and its parents there are only players that follow  $i$  in  $\pi$ , therefore the distance between  $s_i$  and  $s_{i'}$  in  $P_r(\pi)$  is at least  $2^{r-j} + 2$ .

- If  $c(p) \geq \hat{c}_{r-j+1}$  then player  $i$  has a cost-share of at least  $\frac{8k-1}{8k}\hat{c}_{r-j+1}$  (argument 1). But for  $k \geq 6$  and  $j < r$ ,  $\frac{8k-1}{8k}\hat{c}_{r-j+1} > \hat{c}_{r-j} + 1$ .
- If  $c(p) < \hat{c}_{r-j+1}$  and  $p$  contains a shortcut edge of cost  $\hat{c}_{r-j}$ , then this edge will only be used by player  $i$ . This is true because due to the bound on  $c(p)$  the shortcut edge of cost  $\hat{c}_{r-j}$  should have an endpoint between  $s_i$  and any of its parents (excluding them) and the only such edges that may be used by other players in  $\mathbf{P}^i$  is the ones with endpoints at  $s_i$  and any of its parents. Obviously, such an edges cannot belong to  $p$  because then  $s_{i'}$  would be one of  $s_i$ 's parents. This further means that  $p$  should contain another edge of cost at least  $\hat{c}_1$ . Therefore, (by argument 1)  $i$ 's has a cost-share for  $p$  of at least  $\hat{c}_{r-j} + 2\frac{k-1}{k}\frac{8k-1}{8k} > \hat{c}_{r-j} + 1$ , for  $k \geq 3$ .
- If  $c(p) < \hat{c}_{r-j+1}$  and  $p$  does not contain a shortcut edge of cost  $\hat{c}_{r-j}$ , then each edge of  $p$  has cost  $2^w \left(\frac{k-1}{k}\right)^w$  for some  $w \leq r-j-1$ . Suppose that the distance between  $s_i$  and  $s_{i'}$  in  $P_r(\pi)$  is  $2^{r-j} + a$ , where  $2 \leq a < 2^{r-j}$ , then it holds that  $p$  contains an edge whose endpoints have distance in  $P_r(\pi)$  at most  $a$ .<sup>16</sup> Let  $\hat{c}_y$  be the cost of that edge, where  $1 \leq y \leq r-j-1$ . Then the total cost of  $p$  is at least  $2^{r-j} \left(\frac{k-1}{k}\right)^{r-j-1} + 2y \left(\frac{k-1}{k}\right)^y$ . Player 2 is first according to  $\pi$  at  $p$  and by argument 1, her share is at least  $2^{r-j} \left(\frac{k-1}{k}\right)^{r-j-1} \frac{8k-1}{8k} + 2y \left(\frac{k-1}{k}\right)^y \frac{8k-1}{8k} > \hat{c}_{r-j} + 1$ , for  $k \geq 3$ .

We now describe a sequence of best-responses from some  $\hat{\mathbf{P}} \in \mathcal{P}^*$  to  $\mathbf{P}$  ( $\mathbf{P}$  is constructed based on  $\hat{\mathbf{P}}$  as described above). We follow the  $\pi$  order of the players and for each player we apply her best response. First note that players 1 and 2 have no better response, so  $P_1 = \hat{P}_1$  and  $P_2 = \hat{P}_2$ . When we process any other player  $i$ , we have already processed all her predecessors in  $\pi$  and so, the strategies of the other players are  $\mathbf{P}^i$ . Therefore,  $P_i$  is the best response for  $i$  (it may be that  $P_i = \hat{P}_i$ , where no better response exists for  $i$ ). The order that we process the vertices guarantees that  $\mathbf{P} \in \mathcal{P}^*$ . Best-response dynamics guarantee that eventually, no player could perform any best-response, resulting in the desired NE.  $\square$

$\square$

## 6 Stochastic Network Design

In this section we study the *stochastic* model, where the set of the activated vertices is no longer picked adversarially, but it is drawn from some probability distribution  $\Pi$ ; however,  $\Pi$  itself is chosen adversarially. The cost-sharing protocol is decided by the designer without the knowledge of the activated set and the designer may have knowledge of  $\Pi$  or access to some oracle of  $\Pi$ . We next design randomized and deterministic protocols with constant PoA. We note that both protocols can be determined in polynomial time.

### 6.1 Randomized Protocol

We show that there exists a *randomized* ordered protocol that achieves constant PoA. This result holds even for the *black-box* model [54], meaning that the probability distribution is not known to the designer, however she is allowed to draw independent (polynomially many) samples.

<sup>16</sup>Since there are only shortcut edges between vertices of distance in  $P_r(\pi)$  that is a power of 2, we can express the distance between  $s_i$  and  $s_{i'}$  in  $P_r(\pi)$  as  $\sum_i 2^{x_i}$ , with  $1 \leq x_i < 2^{r-j}$ . Then  $\sum_i 2^{x_i} = 2^{r-j} + a$ . Let  $x^*$  be the smallest value among  $x_i$ 's, then  $2^{x^*}$  divides all the terms on the left-hand side and  $2^{r-j}$ , therefore it should divide  $a$ , meaning that  $2^{x^*} \leq a$ .

This problem is closely related to the a priori TSP problem. The result of Shmoys and Talwar [54] on the a priori TSP immediately implies a protocol with a PoA of at most 8. However, by following their analysis for multicast games we can improve this upper bound to 6.78. For completeness we give here the whole proof.

The protocol's design highly relies on approximation algorithms for the minimum Steiner tree problem and therefore, the resulting PoA upper bound (Corollary 29) depends on known approximation ratios for this problem. More precisely, given an  $\alpha$ -approximate minimum Steiner tree, we show an upper bound of  $2(\alpha + 2)$  (Theorem 28). The approximate tree is used in our algorithm as a base in order to construct a spanning tree, which finally determines an order of all vertices; the detailed algorithm is given in Algorithm 1. This algorithm and its slight variants have been used in different contexts: rent-or-buy problem [40], a priori TSP [54] and, stochastic Steiner tree problem [34].

---

**Algorithm 1** Randomized order protocol  $\Xi_{rand}$

---

**Input:** A rooted graph  $G = (V, E, t)$  and an oracle for the probability distribution  $\Pi$ .

**Output:**  $\Xi_{rand}$ .

- Choose a random set of vertices  $R$  by drawing from distribution  $\Pi$  and construct an  $\alpha$ -approximate minimum Steiner tree,  $T_\alpha(R)$ , over  $R \cup \{t\}$ .
  - Connect all other vertices  $V \setminus V(T_\alpha(R))$  with their nearest neighbor in  $V(T_\alpha(R))$  (by breaking ties arbitrarily).
  - Double the edges of that tree and traverse some Eulerian tour starting from  $t$ . Order the vertices based on their first appearance in the tour.
- 

**Theorem 28.** Given an  $\alpha$ -approximate solution of the minimum Steiner tree problem,  $\Xi_{rand}$  has PoA at most  $2(\alpha + 2)$ .

*Proof.* Let  $\pi$  be the order of  $V$ , defined by  $\Xi_{rand}$ , and  $S$  be the random set of activated vertices that require connectivity with  $t$ . For the rest of the proof we denote by  $MST(S)$  a minimum spanning tree over the vertices  $S \cup \{t\}$  on the metric closure<sup>17</sup> of  $G$ .

Let  $s_1, \dots, s_r$  be the vertices of  $S$  as appeared in  $\pi$  and the strategy profile  $\mathbf{P}_R(S) = (P_1, \dots, P_r)$  be a NE of set  $S$ . Under the convention that  $s_0 = t$ ,  $c_{s_i}(\mathbf{P}_R(S)) \leq d_G(s_i, s_{i-1})$  for all  $s_i \in S$ . We construct a tree  $T_{R,S}$  from the  $T_\alpha(R)$  of Algorithm 1, by connecting only all vertices of  $S \setminus V(T_\alpha(R))$  with their nearest neighbor in  $V(T_\alpha(R))$  (by breaking ties in accordance to Algorithm 1). Note that, by doubling the edges of  $T_{R,S}$ , there exists an Eulerian tour starting from  $t$ , where the order of the vertices  $S$  (based on their first appearance in the tour) is  $\pi$  restricted to the set  $S$ <sup>18</sup>. Therefore,  $\sum_{s_i \in S} d_{T_{R,S}}(s_i, s_{i-1}) + d_{T_{R,S}}(s_0, s_r) = 2c(T_{R,S})$ . By combining the above,

$$\begin{aligned} c(\mathbf{P}_R(S)) &= \sum_{s_i \in S} c_{s_i}(\mathbf{P}_R(S)) \leq \sum_{s_i \in S} d_G(s_i, s_{i-1}) \\ &\leq \sum_{s_i \in S} d_{T_{R,S}}(s_i, s_{i-1}) \leq 2c(T_{R,S}). \end{aligned} \tag{1}$$

---

<sup>17</sup>The metric closure of an undirected graph  $G$  is the complete undirected graph on the vertex set  $V(G)$ , where the edge costs equal the shortest path distances in  $G$ .

<sup>18</sup>This Eulerian tour matches the tour constructed by shortcutting the Eulerian tour of Algorithm 1 to contain only the vertices  $R \cup S \cup \{t\}$ .

Let  $D_v(R)$  be the distance of  $v$  from its nearest neighbor in  $(R \cup \{t\}) \setminus \{v\}$ . In the special case that  $v = t$ , we define  $D_v(R) = 0$ . Then,

$$c(T_{R,S}) = c(T_\alpha(R)) + \sum_{v \in S \setminus R} D_v(R) \leq c(T_\alpha(R)) + \sum_{v \in S} D_v(R). \quad (2)$$

We use an indicator  $I(v \in S)$  which is 1 when  $v \in S$  and 0 otherwise; then

$$\sum_{v \in S} D_v(R) = \sum_v I(v \in S) D_v(R).$$

By taking the expectation over  $R$  and  $S$ ,

$$\mathbb{E}_R[\mathbb{E}_S[c(T_{R,S})]] \leq \mathbb{E}_R[c(T_\alpha(R))] + \mathbb{E}_R[\mathbb{E}_S[\sum_{v \in V} I(v \in S) D_v(R)]].$$

Since  $S$  and  $R$  are independent samples we can bound the second term as:

$$\begin{aligned} \mathbb{E}_R[\mathbb{E}_S[\sum_{v \in V} I(v \in S) D_v(R)]] &= \sum_{v \in V} \mathbb{E}_S[I(v \in S)] \mathbb{E}_R[D_v(R)] = \sum_{v \in V} \mathbb{E}_S[I(v \in S)] \mathbb{E}_S[D_v(S)] \\ &= \mathbb{E}_S[\sum_{v \in V} I(v \in S) D_v(S)] = \mathbb{E}_S[\sum_{v \in S} D_v(S)] \leq \mathbb{E}_S[c(MST(S))]. \end{aligned} \quad (3)$$

The third equality holds since  $D_v(S)$  is the distance of  $v$  from its nearest neighbor in  $(S \cup \{t\}) \setminus \{v\}$  and it is independent of the event  $I(v \in S)$ . For the inequality, note that  $D_v(S)$  is upper bounded by the distance of  $v$  from its parent in the  $MST(S)$ .

Let  $T_S^*$  be the minimum Steiner tree over  $S \cup \{t\}$ , then it is well known that  $c(MST(S)) \leq 2c(T_S^*)$ . Overall,

$$\mathbb{E}_R[\mathbb{E}_S[c(\mathbf{P}_R(S))]] \leq 2 \mathbb{E}_R[\mathbb{E}_S[c(T_{R,S})]] \leq 2(\mathbb{E}_R[c(T_\alpha(S))] + \mathbb{E}_S[c(MST(S))]) \leq 2(\alpha + 2) \mathbb{E}_S[c(T_S^*)].$$

□

By applying the 1.39-approximation algorithm of [18] we get the following corollary.

**Corollary 29.**  $\Xi_{rand}$  has PoA at most 6.78.

## 6.2 Deterministic Protocol

We now consider that each vertex  $v$  is activated independently with probability  $p_v$ ; the set of the activated vertices is sampled based on the probabilities  $p_v$ 's, i.e., the probability that set  $S$  is activated is  $\Pi(S) = \prod_{v \in S} p_v \cdot \prod_{v \notin S} (1 - p_v)$ . The probabilities  $p_v$ 's (and therefore  $\Pi$ ), are chosen adversarially. We additionally assume that the probabilities  $p_v$ 's are known to the designer. We show that there exists a *deterministic* ordered protocol that achieves constant PoA.

**Theorem 30.** There exists a deterministic ordered protocol with PoA at most 16.

*Proof.* We use derandomization techniques similar to [54, 58] and for completeness we give the full proof here. First we discuss how we can get a PoA of 6.78, if we drop the requirement of determining the protocol in polynomial time. Similar to the proof of Theorem 28 we define the tree  $T_{R,S}$  for the random activated set  $S$  as follows: we construct  $T_{R,S}$  from the  $T_\alpha(R)$  of Algorithm 1, by connecting only all vertices of  $S \setminus V(T_\alpha(R))$  with their nearest neighbor in  $V(T_\alpha(R))$  (by

breaking ties in accordance to Algorithm 1). We apply the standard derandomization approach of *conditional expectation method* on  $T_{R,S}$ . More precisely, we construct a deterministic set  $\hat{R}$  to replace the random set  $R$  in Algorithm 1, by deciding for each vertex of  $V \setminus \{t\}$ , one by one, whether it belongs to  $\hat{R}$  or not. Assume that we have already processed the set  $Q \subset V$  and we have decided that for its partition  $(Q_1, Q_2)$ ,  $Q_1 \subseteq \hat{R}$  and  $Q_2 \cap \hat{R} = \emptyset$  (starting from  $Q_1 = \{t\}$  and  $Q_2 = \emptyset$ ). Let  $v$  be the next vertex to be processed. From the conditional expectations and the independent activations we know that

$$\begin{aligned} \mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset] &= \mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset, v \in R] p_v \\ &+ \mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset, v \notin R] (1 - p_v), \end{aligned}$$

meaning that

$$\begin{aligned} \text{either } \mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset, v \in R] &\leq \mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset], \\ \text{or } \mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset, v \notin R] &\leq \mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset]. \end{aligned}$$

In the first case we add  $v$  in  $Q_1$  and in the second case we add  $v$  in  $Q_2$ . Therefore, after processing all vertices,  $Q_1 = \hat{R}$  and  $\mathbb{E}_S [c(T_{\hat{R},S})] \leq \mathbb{E}_{S,R} [c(T_{R,S})]$ . If we replace the sampled  $R$  of Algorithm 1 with the deterministic set  $\hat{R}$ , we can get the same bound on the PoA with the randomized protocol of Theorem 28.

However, the value of  $\mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset]$  seems difficult to be computed in polynomial time; the reason is that it involves the computation of  $\mathbb{E}_R [c(T_\alpha(R)) | Q_1 \subseteq R, Q_2 \cap R = \emptyset]$  which seems hard to be handled. To overcome this problem we use an estimator  $EST(Q_1, Q_2)$  of  $\mathbb{E}_{S,R} [c(T_{R,S}) | Q_1 \subseteq R, Q_2 \cap R = \emptyset]$ , which is constant away from the optimum  $\mathbb{E}_S [c(T_S^*)]$ , where  $T_S^*$  is the minimum Steiner tree over  $S \cup \{t\}$ . Following [58, 54], we use the optimum solution of the relaxed Connected Facility Location Problem (CFLP) on  $G$  in order to construct a feasible solution  $\bar{y}$  of the relaxed Steiner Tree Problem (STP) for a given set  $R$ . We show that the objective's value of the fractional STP for  $\bar{y}$  is constant away from  $\mathbb{E}_S [c(T_S^*)]$  and that its (conditional) expectation over  $R$  can be efficiently computed. This quantity is used in order to construct the estimator  $EST(Q_1, Q_2)$ . We apply the method of conditional expectations on  $EST(Q_1, Q_2)$  and after processing all vertices, by using the primal-dual algorithm [36], we compute a Steiner tree on  $Q_1$  with cost no more than twice the cost of the fractional solution.

In the rooted CFLP, a rooted graph  $G = (V, E, t)$  is given and the designer should select some facilities to open, including  $t$ , and connects them via some Steiner tree  $T$ . Every other vertex is assigned to some facility. The cost of the solution is  $M$  ( $M > 1$ ) times the cost of  $T$ , plus the distance of every other vertex from its assigned facility. Our analysis requires to consider a slightly different cost of the solution, which is the cost of  $T$ , plus the distance of every other vertex  $v$  from its assigned facility multiplied by the probability  $p_v$  of activating  $v$ . In the following LP relaxation of the CFLP,  $z_e$  and  $x_{ij}$  are 0-1 variables indicate, respectively, if  $e \in E(T)$  and whether the vertex  $j$  is assigned to facility  $i$ .  $\delta(U)$  denotes the set of edges with one endpoint in  $U$  and the other in  $V \setminus U$ ,  $d(i, j)$  denotes the minimum distance between vertices  $i$  and  $j$  in  $G$  and  $c_e$  is the cost of edge  $e$ .

<b>LP1: CFLP</b>		
	$\min$	$B + C$
subject to	$\sum_{i \in V} x_{ij} = 1$	$\forall j \in V$
	$\sum_{e \in \delta(U)} z_e \geq \sum_{i \in U} x_{ij}$	$\forall j \in V, \forall U \subseteq V \setminus \{t\}$
	$B = \sum_{e \in E} c_e z_e$	
	$C = \sum_{j \in V} p_j \sum_{i \in V} d(i, j) x_{ij}$	
	$z_e, x_{ij} \geq 0$	$\forall i, j \in V$ and $\forall e \in E$

Let  $(\mathbf{z}^* = (z_e^*)_e, \mathbf{x}^* = (x_{ij}^*)_{ij}, B^*, C^*)$  be the optimum solution of LP1.

*Claim 31.*  $B^* + C^* \leq 3 \mathbb{E}_S[c(T_S^*)]$ , where  $T_S^*$  is the minimum Steiner tree over  $S \cup \{t\}$ .

*Proof.* Given a set  $S \subseteq V$ , for every edge  $e \in T_S^*$ , let  $z_e = 1$  and, for  $e \notin T_S^*$ , let  $z_e = 0$ . Moreover, for every  $j \in V$ , let  $x_{ij} = 1$ , if  $i$  is  $j$ 's nearest neighbor in  $(S \cup \{t\}) \setminus \{j\}$ . Set the rest of  $x_{ij}$  equal to 0. Note that this is a feasible solution of LP1 with objective value  $B_S + C_S \leq c(T_S^*) + \sum_{v \in V} p_v D_v(S)$ . By taking the expectation over  $S$ ,

$$\begin{aligned}
B^* + C^* &\leq \mathbb{E}_S[B_S + C_S] \leq \mathbb{E}_S[c(T_S^*)] + \sum_{v \in V} \mathbb{E}_S[I(v \in S)] \mathbb{E}_S[D_v(S)] \\
&= \mathbb{E}_S[c(T_S^*)] + \mathbb{E}_S[\sum_{v \in S} D_v(S)] \leq \mathbb{E}_S[c(T_S^*)] + \mathbb{E}_S[c(MST(S))] \leq 3 \mathbb{E}_S[c(T_S^*)].
\end{aligned}$$

□

By using the solution  $(\mathbf{z}^* = (z_e^*)_e, \mathbf{x}^* = (x_{ij}^*)_{ij}, B^*, C^*)$ , we construct a feasible solution for the following LP relaxation of the STP over some set  $R \cup \{t\}$ .

<b>LP2: STP over <math>R \cup \{t\}</math></b>		
	$\min$	$\sum_{e \in E} c_e y_e$
subject to	$\sum_{e \in \delta(U)} y_e \geq 1$	$\forall U \subseteq V \setminus \{t\} : R \cap U \neq \emptyset$
	$y_e \geq 0$	$\forall e \in E$

We define  $a_{ij}(e) = 1$  if  $e$  lies in the shortest path between  $i$  and  $j$  and it is 0 otherwise. For every edge  $e$  we set  $\bar{y}_e = z_e^* + \sum_{j \in R} \sum_{i \in V} a_{ij}(e) x_{ij}^*$ .

*Claim 32.*  $\bar{\mathbf{y}} = (\bar{y}_e)_e$  is a feasible solution for LP2.

*Proof.* The proof is identical with the one in [58] but we give it here for completeness. Consider any set  $U \subseteq V \setminus \{t\}$  such that  $R \cap U \neq \emptyset$  and let  $\ell \in R \cap U$ . It follows that

$$\begin{aligned}
\sum_{e \in \delta(U)} \bar{y}_e &\geq \sum_{e \in \delta(U)} z_e^* + \sum_{e \in \delta(U)} \sum_{j \in R} \sum_{i \in V} a_{ij}(e) x_{ij}^* \geq \sum_{i \in U} x_{i\ell}^* + \sum_{e \in \delta(U)} \sum_{i \in V} a_{i\ell}(e) x_{i\ell}^* \\
&\geq \sum_{i \in U} x_{i\ell}^* + \sum_{i \notin U} x_{i\ell}^* \sum_{e \in \delta(U)} a_{i\ell}(e) \geq \sum_{i \in U} x_{i\ell}^* + \sum_{i \notin U} x_{i\ell}^* = 1.
\end{aligned}$$

For the last inequality, note that  $a_{i\ell}(e)$  should be 1 for at least one  $e \in \delta(U)$  since  $i \notin U$  and  $\ell \in U$ . □

*Claim 33.* Let  $\bar{c}_{ST}(R)$  be the cost of the objective of LP2 induced by the solution  $\bar{\mathbf{y}}$ . Then  $\mathbb{E}_R[\bar{c}_{ST}(R)] = B^* + C^*$ .

*Proof.*

$$\begin{aligned}\mathbb{E}_R[\bar{c}_{ST}(R)] &= \mathbb{E}_R\left[\sum_{e \in E} c_e(z_e^* + \sum_{j \in R} \sum_{i \in V} a_{ij}(e)x_{ij}^*)\right] = B^* + \mathbb{E}_R\left[\sum_{j \in R} \sum_{i \in V} \sum_{e \in E} c_e a_{ij}(e)x_{ij}^*\right] \\ &= B^* + \mathbb{E}_R\left[\sum_{j \in R} \sum_{i \in V} d(i, j)x_{ij}^*\right] = B^* + \sum_{j \in V} p_j \sum_{i \in V} d(i, j)x_{ij}^* = B^* + C^*.\end{aligned}$$

□

Observe that due to the expression of  $\bar{\mathbf{y}}$  we can efficiently compute any conditional expectation

$$\mathbb{E}_R[\bar{c}_{ST}(R)|Q_1 \subseteq R, Q_2 \cap R = \emptyset];$$

this is because

$$\mathbb{E}_R\left[\sum_{j \in R} \sum_{i \in V} a_{ij}(e)x_{ij}^*|Q_1 \subseteq R, Q_2 \cap R = \emptyset\right] = \sum_{j \in Q_1} \sum_{i \in V} a_{ij}(e)x_{ij}^* + \sum_{j \notin Q_1 \cup Q_2} p_j \sum_{i \in V} a_{ij}(e)x_{ij}^*.$$

We further define  $c_C(R) = \sum_{v \in V} p_v D_v(R)$ . We can also efficiently compute any conditional expectation  $\mathbb{E}[c_C(R)|Q_1 \subseteq R, Q_2 \cap R = \emptyset]$  (Claim 2.1 of [58]). We are ready to define our estimator:

$$EST(Q_1, Q_2) = 2 \mathbb{E}_R[\bar{c}_{ST}(R)|Q_1 \subseteq R, Q_2 \cap R = \emptyset] + \mathbb{E}_R[\bar{c}_C(R)|Q_1 \subseteq R, Q_2 \cap R = \emptyset].$$

Our goal is to define a deterministic set  $\hat{R}$  to replace the sampled  $R$  of Algorithm 1. We process the vertices one by one and we decide if they belong to  $\hat{R}$  by using the model conditional expectations on  $EST(Q_1, Q_2)$ . More specifically, assume that we have already processed the sets  $Q_1$  and  $Q_2$  (starting from  $Q_1 = \{t\}$  and  $Q_2 = \emptyset$ ) such that  $Q_1 \subseteq \hat{R}$  and  $Q_2 \cap \hat{R} = \emptyset$ . Let  $v$  be the next vertex to be processed. From the conditional expectations and the independent activations we know that  $EST(Q_1, Q_2) = p_v EST(Q_1 \cup \{v\}, Q_2) + (1 - p_v) EST(Q_1, Q_2 \cup \{v\})$ . If  $EST(Q_1 \cup \{v\}, Q_2) \leq EST(Q_1, Q_2)$  we add  $v$  to  $Q_1$ , otherwise we add  $v$  to  $Q_2$ . After processing all vertices and by using Claims 31 and 33,

$$\begin{aligned}EST(\hat{R}, V \setminus \hat{R}) &\leq EST(\{t\}, \emptyset) = 2 \mathbb{E}_R[\bar{c}_{ST}(R)] + \mathbb{E}_R[\bar{c}_C(R)] \\ &\leq 6 \mathbb{E}_S[c(T_S^*)] + \sum_{v \in V} p_v \mathbb{E}_R[D_v(R)] = 6 \mathbb{E}_S[c(T_S^*)] + \mathbb{E}_R\left[\sum_{v \in V} I(v \in R) D_v(R)\right] \\ &\leq 6 \mathbb{E}_S[c(T_S^*)] + \mathbb{E}_R[c(MST(R))] \leq 8 \mathbb{E}_S[c(T_S^*)].\end{aligned}$$

Let  $T_{PD}(\hat{R})$  be the Steiner tree over  $\hat{R} \cup \{t\}$  computed by the primal-dual algorithm [36]. Then,

$$EST(\hat{R}, V \setminus \hat{R}) = 2\bar{c}_{ST}(\hat{R}) + \sum_{v \in V} p_v D_v(\hat{R}) \geq c(T_{PD}(\hat{R})) + \mathbb{E}_S\left[\sum_{v \in S} D_v(\hat{R})\right].$$

By combining inequalities (1) and (2) (after replacing  $R$  by  $\hat{R}$  and  $T_\alpha(\hat{R})$  by  $T_{PD}(\hat{R})$ ) with all the above, we have that

$$\mathbb{E}_S[c(\mathbf{P}_{\hat{R}}(S))] \leq 2 \left( c(T_{PD}(\hat{R})) + \mathbb{E}_S\left[\sum_{v \in S} D_v(\hat{R})\right] \right) \leq 2EST(\hat{R}, V \setminus \hat{R}) \leq 16 \mathbb{E}_S[c(T_S^*)].$$

□

## References

- [1] F. ABED, J. R. CORREA, AND C. HUANG, *Optimal coordination mechanisms for multi-job scheduling games*, in Algorithms - ESA 2014 - 22th Annual European Symposium, Wrocław, Poland, September 8-10, 2014. Proceedings, 2014, pp. 13–24.
- [2] F. ABED AND C. HUANG, *Preemptive coordination mechanisms for unrelated machines*, in Algorithms - ESA 2012 - 20th Annual European Symposium, Ljubljana, Slovenia, September 10-12, 2012. Proceedings, 2012, pp. 12–23.
- [3] N. ALON AND Y. AZAR, *On-line Steiner trees in the euclidean plane*, Discrete Comput. Geom., 10 (1993), pp. 113–121.
- [4] N. ALON, R. RADOICIC, B. SUDAKOV, AND J. VONDRÁK, *A ramsey-type result for the hypercube*, J. Graph Theory, 53 (2006), pp. 196–208.
- [5] E. ANSHELEVICH, A. DASGUPTA, J. M. KLEINBERG, É. TARDOS, T. WEXLER, AND T. ROUGHGARDEN, *The price of stability for network design with fair cost allocation*, SIAM J. Comput., 38 (2008), pp. 1602–1623.
- [6] E. ANSHELEVICH, A. DASGUPTA, É. TARDOS, AND T. WEXLER, *Near-optimal network design with selfish agents*, Theory Comput., 4 (2008), pp. 77–109.
- [7] B. AWERBUCH, Y. AZAR, AND Y. BARTAL, *On-line generalized Steiner problem*, Theoret. Comput. Sci., 324 (2004), pp. 313–324.
- [8] Y. AZAR, K. JAIN, AND V. MIRROKNI, *(Almost) optimal coordination mechanisms for unrelated machine scheduling*, in Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '08, Philadelphia, PA, USA, 2008, Society for Industrial and Applied Mathematics, pp. 323–332.
- [9] P. BERMAN AND C. COULSTON, *On-line algorithms for steiner tree problems (extended abstract)*, in Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997, 1997, pp. 344–353.
- [10] D. BERTSIMAS AND M. GRIGNI, *On the spacefilling curve heuristic for the euclidean traveling salesman problem*, Technical report, Massachusetts Institute of Technology, Cambridge, MA, (1988).
- [11] D. J. BERTSIMAS, P. JAILLET, AND A. R. ODONI, *A priori optimization*, Oper. Res., 38 (1990), pp. 1019–1033.
- [12] A. BHALGAT, D. CHAKRABARTY, AND S. KHANNA, *Optimal lower bounds for universal and differentially private Steiner trees and TSPs*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011, Princeton, NJ, USA, August 17-19, 2011. Proceedings, 2011, pp. 75–86.
- [13] S. BHATTACHARYA, S. IM, J. KULKARNI, AND K. MUNAGALA, *Coordination mechanisms from (almost) all scheduling policies*, in Innovations in Theoretical Computer Science, ITCS'14, Princeton, NJ, USA, January 12-14, 2014, 2014, pp. 121–134.

- [14] S. BHATTACHARYA, J. KULKARNI, AND V. S. MIRROKNI, *Coordination mechanisms for selfish routing over time on a tree*, in Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I, 2014, pp. 186–197.
- [15] V. BILÒ AND R. BOVE, *Bounds on the price of stability of undirected network design games with three players*, Journal of Interconnection Networks, 12 (2011), pp. 1–17.
- [16] V. BILÒ, I. CARAGIANNIS, A. FANELLI, AND G. MONACO, *Improved lower bounds on the price of stability of undirected network design games*, Theory Comput. Syst., 52 (2013), pp. 668–686.
- [17] V. BILÒ, M. FLAMMINI, AND L. MOSCARDELLI, *The price of stability for undirected broadcast network design with fair cost allocation is constant*, in 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA, 2013, pp. 638–647.
- [18] J. BYRKA, F. GRANDONI, T. ROTHVOSS, AND L. SANITÀ, *An improved  $lp$ -based approximation for steiner tree*, in Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010, 2010, pp. 583–592.
- [19] I. CARAGIANNIS, *Efficient coordination mechanisms for unrelated machine scheduling*, Algorithmica, 66 (2013), pp. 512–540.
- [20] M. CHARIKAR, H. J. KARLOFF, C. MATHIEU, J. NAOR, AND M. E. SAKS, *Online multicast with egalitarian cost sharing*, in SPAA 2008: Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures, Munich, Germany, June 14-16, 2008, 2008, pp. 70–76.
- [21] C. CHEKURI, J. CHUZHUY, L. LEWIN-EYTAN, J. NAOR, AND A. ORDA, *Non-cooperative multicast and facility location games*, IEEE Journal on Selected Areas in Communications, 25 (2007), pp. 1193–1206.
- [22] H. CHEN AND T. ROUGHGARDEN, *Network design with weighted players*, Theory Comput. Syst., 45 (2009), pp. 302–324.
- [23] H. CHEN, T. ROUGHGARDEN, AND G. VALIANT, *Designing network protocols for good equilibria*, SIAM J. Comput., 39 (2010), pp. 1799–1832.
- [24] G. CHRISTODOULOU, C. CHUNG, K. LIGETT, E. PYRGA, AND R. VAN STEE, *On the price of stability for undirected network design*, in Approximation and Online Algorithms, 7th International Workshop, WAOA 2009, Copenhagen, Denmark, September 10-11, 2009. Revised Papers, 2009, pp. 86–97.
- [25] G. CHRISTODOULOU, E. KOUTSOPIAS, AND A. NANAVATI, *Coordination mechanisms*, Theoret. Comput. Sci., 410 (2009), pp. 3327–3336.
- [26] G. CHRISTODOULOU, S. LEONARDI, AND A. SGOURITSA, *Designing cost-sharing methods for bayesian games*, Theory Comput. Syst., 63 (2019), pp. 4–25.
- [27] G. CHRISTODOULOU, K. MEHLHORN, AND E. PYRGA, *Improving the price of anarchy for selfish routing via coordination mechanisms*, Algorithmica, 69 (2014), pp. 619–640.

- [28] G. CHRISTODOULOU AND A. SGOURITSA, *Designing networks with good equilibria under uncertainty*, in Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016, 2016, pp. 72–89.
- [29] G. CHRISTODOULOU AND A. SGOURITSA, *An improved upper bound for the universal TSP on the grid*, in Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19, 2017, pp. 1006–1021.
- [30] R. COLE, J. R. CORREA, V. GKATZELIS, V. S. MIRROKNI, AND N. OLVER, *Inner product spaces for minsum coordination mechanisms*, in Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011, 2011, pp. 539–548.
- [31] G. CORNUÉJOLS, J. FONLUPT, AND D. NADDEF, *The traveling salesman problem on a graph and some related integer polyhedra*, Math. Program., 33 (1985), pp. 1–27.
- [32] Y. DISSER, A. E. FELDMANN, M. KLIMM, AND M. MIHALÁK, *Improving the  $H_k$ -bound on the price of stability in undirected shapley network design games*, in Algorithms and Complexity, 8th International Conference, CIAC 2013, Barcelona, Spain, May 22-24, 2013. Proceedings, 2013, pp. 158–169.
- [33] A. FIAT, H. KAPLAN, M. LEVY, S. OLONETSKY, AND R. SHABO, *On the price of stability for designing undirected networks with fair cost allocations*, in Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part I, 2006, pp. 608–618.
- [34] N. GARG, A. GUPTA, S. LEONARDI, AND P. SANKOWSKI, *Stochastic analyses for online combinatorial optimization problems*, in Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008, 2008, pp. 942–951.
- [35] V. GKATZELIS, K. KOLLIAS, AND T. ROUGHGARDEN, *Optimal cost-sharing in weighted congestion games*, in Web and Internet Economics - 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings, 2014, pp. 72–88.
- [36] M. X. GOEMANS AND D. P. WILLIAMSON, *A general approximation technique for constrained forest problems*, SIAM J. Comput., 24 (1995), pp. 296–317.
- [37] R. GOPALAKRISHNAN, J. R. MARDEN, AND A. WIERMAN, *Potential games are Necessary to ensure pure nash equilibria in cost sharing games*, Math. Oper. Res., 39 (2014), pp. 1252–1296.
- [38] I. GORODEZKY, R. D. KLEINBERG, D. B. SHMOYS, AND G. SPENCER, *Improved lower bounds for the universal and a priori TSP*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010, Barcelona, Spain, September 1-3, 2010. Proceedings, 2010, pp. 178–191.
- [39] R. L. GRAHAM, B. L. ROTHSCHILD, AND J. H. SPENCER, *Ramsey Theory, 2nd Edition*, Wiley Ser. Discrete Math. Optim., New York, 1990.
- [40] A. GUPTA, A. KUMAR, M. PÁL, AND T. ROUGHGARDEN, *Approximation via cost sharing: Simpler and better approximation algorithms for network design*, J. ACM, 54 (2007), p. 11.

- [41] M. T. HAJIAGHAYI, R. D. KLEINBERG, AND F. T. LEIGHTON, *Improved lower and upper bounds for universal TSP in planar metrics*, in Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006, 2006, pp. 649–658.
- [42] T. HARKS, A. HUBER, AND M. SUREK, *A characterization of undirected graphs admitting optimal cost shares*, in Web and Internet Economics - 13th International Conference, WINE 2017, Bangalore, India, December 17-20, 2017, Proceedings, 2017, pp. 237–251.
- [43] M. IMASE AND B. M. WAXMAN, *Dynamic Steiner tree problem*, SIAM J. Discrete Math., 4 (1991), pp. 369–384.
- [44] N. IMMORLICA, E. L. LI, V. S. MIRROKNI, AND A. S. SCHULZ, *Coordination mechanisms for selfish scheduling*, in Internet and Network Economics, First International Workshop, WINE 2005, Hong Kong, China, December 15-17, 2005, Proceedings, 2005, pp. 55–69.
- [45] L. JIA, G. LIN, G. NOUBIR, R. RAJARAMAN, AND R. SUNDARAM, *Universal approximations for TSP, Steiner tree, and set cover*, in Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005, 2005, pp. 386–395.
- [46] D. R. KARGER AND M. MINKOFF, *Building steiner trees with incomplete global knowledge*, in 41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA, 2000, pp. 613–623.
- [47] K. KOLLIAS, *Nonpreemptive coordination mechanisms for identical machines*, Theory Comput. Syst., 53 (2013), pp. 424–440.
- [48] E. KOUTSOUPIAS AND C. H. PAPADIMITRIOU, *Worst-case equilibria*, in STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings, 1999, pp. 404–413.
- [49] E. LEE AND K. LIGETT, *Improved bounds on the price of stability in network cost sharing games*, in ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013, 2013, pp. 607–620.
- [50] J. LI, *An  $o(\log(n)/\log(\log(n)))$  upper bound on the price of stability for undirected shapley network design games*, Inf. Process. Lett., 109 (2009), pp. 876–878.
- [51] J. R. MARDEN AND A. WIERMAN, *Distributed welfare games*, Oper. Res., 61 (2013), pp. 155–168.
- [52] L. K. PLATZMAN AND J. J. BARTHOLDI III, *Spacefilling curves and the planar travelling salesman problem*, J. ACM, 36 (1989), pp. 719–737.
- [53] F. SCHALEKAMP AND D. B. SHMOYS, *Algorithms for the universal and a priori TSP*, Oper. Res. Lett., 36 (2008), pp. 1–3.
- [54] D. B. SHMOYS AND K. TALWAR, *A constant approximation algorithm for the a priori travelling salesman problem*, in Integer Programming and Combinatorial Optimization, 13th International Conference, IPCO 2008, Bertinoro, Italy, May 26-28, 2008, Proceedings, 2008, pp. 331–343.

- [55] M. M. SYSLO, *Characterizations of outerplanar graphs*, Discrete Math., 26 (1979), pp. 47 – 53.
- [56] S. UMBOH, *Online network design algorithms via hierarchical decompositions*, in Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015, 2015, pp. 1373–1387.
- [57] P. VON FALKENHAUSEN AND T. HARKS, *Optimal cost sharing for resource selection games*, Math. Oper. Res., 38 (2013), pp. 184–208.
- [58] D. P. WILLIAMSON AND A. VAN ZUYLEN, *A simpler and better derandomization of an approximation algorithm for single source rent-or-buy*, Oper. Res. Lett., 35 (2007), pp. 707–712.