

# Designing Networks with Good Equilibria under Uncertainty

George Christodoulou\*

Alkmini Sgouritsa†

## 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 are interested in 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 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 metric help in the design?*

We first demonstrate that there exist classes of graphs where knowledge of the underlying 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 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 each player is independently 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. 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 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*  $s_1, \dots, s_k$  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 analyze the quality of the equilibrium solutions; this can be measured with the Price of Anarchy (PoA) [43] (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.

\*Department of Computer Science, University of Liverpool, UK. Email: [gchristo@liverpool.ac.uk](mailto:gchristo@liverpool.ac.uk). This work was supported by EP/M008118/1, EP/K01000X/1 and LT140046.

†Department of Computer Science, University of Liverpool, UK. Email: [salkmini@liverpool.ac.uk](mailto:salkmini@liverpool.ac.uk)

<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  $O(\log / \log \log k)$  [45].

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 [22] 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 has a constant PoA; the Nash equilibria 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 (OSTP).

**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 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 of players, and we take the worst-case approach 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

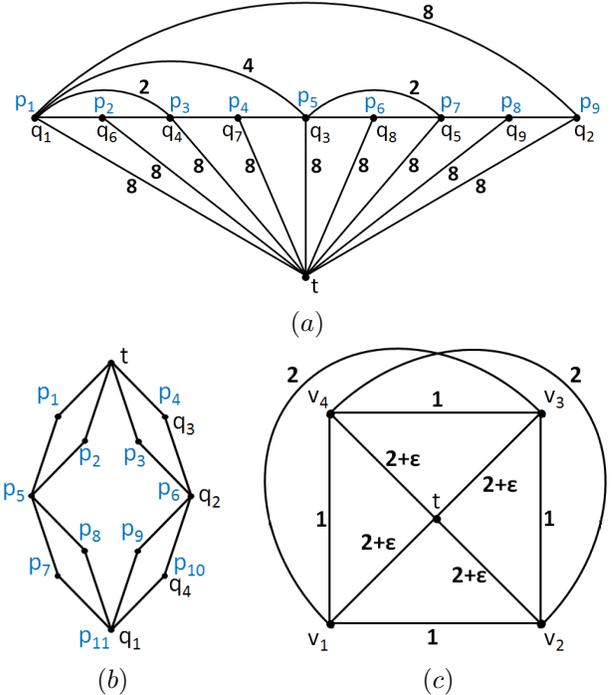


Figure 1: In (a) and (b) we assume two orders on the vertices, denoted by  $q_i$  or  $p_i$ . The  $q$ -order is adversarially chosen and simulates the adversary for the OSTP [38], that results to high PoA of  $\Omega(\log k)$ . The  $p$ -order results to constant PoA. (c) shows an example where both the *best* ordered protocol and the Shapley protocol have  $\text{PoA} \geq 5/4$ , whereas there is an *intermediate* protocol with PoA 1. In edges with no written cost, we consider the unit cost; we take  $\epsilon > 0$  arbitrarily small.

a universal protocol where the expected worst-case cost in the NE is not far from the expected optimal cost.

**EXAMPLE 1.1. (Ordered protocols).** An important special class with interesting properties is that of *ordered protocols*. 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 OSTP, where the worst-case order is considered. The following example 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 metric, that universal protocols can use in their favor 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 OSTP [38], and places the players accordingly (see for example Figure 1(a),(b), the  $q$  labels). Therefore the PoA of uniform ordered protocol is  $\Omega(\log k)$  [22].

*Universal protocols.* The designer takes into account the graph; consider the worst-case graph for the Greedy Algorithm of the OSTP (illustrated in Figure 1(a),(b) for a small number of players). For the graph of Figure 1(a), choose the linear order dictated from the path  $p_1, \dots, p_9$  (say from left to right). For the graph of Figure 1(b) order the vertices according to their distance from  $t$ ,  $p_1, \dots, p_{11}$ . 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 1.2. (Generalized weighted Shapley).** In [22], it was shown that ordered protocols are essentially optimal among uniform protocols. In our model, the choice of the optimal method may depend on the underlying graph metric. Take the example in Figure 1(c). By using Shapley cost sharing the adversary can choose  $v_1, v_2, v_3$  and in the NE  $v_1, v_3$  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, the cost of the NE is 5 and the minimum Steiner tree that connects those vertices with  $t$  has cost 4 (by ignoring  $\varepsilon$ ) and therefore,  $\text{PoA} \geq 5/4$ .

However the following (generalized Shapley) protocol has  $\text{PoA} = 1$ . 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 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 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). The central question we address is: *to what extent does prior knowledge of the 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 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 metrics.

Our main technical result shows that there exist graph metrics, for which knowing the underlying 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 metric [38, 22].

Then we switch to the stochastic model, where each player is independently activated according to some probability that is known to the designer. We show that there exists a *randomized* ordered protocol that achieves constant PoA. By using standard derandomization techniques [52, 48], 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 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 3) 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 OSTP (see the labeled path  $(q_1, q_6, q_4, \dots, q_2)$  in Figure 1(a)). The challenge is to show that high dimensional hypercubes exhibit such a distance preserving structure *no matter how the vertices are ordered*. Section 3 is devoted to this 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 [35]. 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 4), we extend the lower bound to *all universal* cost-sharing protocols, by using the characterization of [22]. At a high level, we use as basis the lower bound 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 are 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 [52, 48].

**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 [14, 16, 21, 23, 28, 29, 45, 15] 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 [45], while for broadcast games, a series of work [29, 44] lead finally to a constant due to Bilò et al. [16]. The PoA of some special equilibria has been also studied in [19, 20].

Chen, Roughgarden and Valiant [22] 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

[33]), and they thoroughly studied PoA and PoS for several cases. Falkenhausen and Harks [51] studied singleton and matroid games with weighted players, while Gkatzelis, Kollias and Roughgarden [31] focused on weighted congestion games with polynomial cost functions.

Close in spirit to universal cost-sharing protocols is the notion of Coordination Mechanisms [24] 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 [24, 39, 42, 8, 18, 26, 12, 1, 2] as well as for simple routing games [25, 13].

As discussed in Example 1.1, the analysis of the equilibria induced by ordered protocols corresponds to the analysis of Greedy Algorithm for the MST. In the uniform model, this corresponds to the analysis of the Greedy Algorithm [38, 7] for the (Generalized) OSTP [3, 9, 50], which was shown to be  $\Theta(\log k)$ -competitive by Imase and Waxman [38] ( $O(\log^2 k)$ -competitive for the Generalized OSTP by [7]). The universal model is closely related to universal network design problems [40], 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 [46, 10, 37, 34, 40].

Much work has been done in stochastic models and we only mention the most related to our work. Karger and Minkoff [41] 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. [30] gave bounds on the approximation of the stochastic online Steiner tree problem. A line of works [11, 34, 47, 48] has studied the *a priori TSP*. Shmoys and Talwar [48] 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*

<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, 22]).

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

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*  $\xi : 2^N \rightarrow \mathbb{R}_+^{|N|}$  decides, for every subset  $R \subseteq N$ , the cost-share  $\xi(i, R)$  for each player  $i \in R$ . A natural rule is that the shares for players not included in  $R$  should always be 0, i.e. if  $i \notin R$ ,  $\xi(i, R) = 0$ . W.l.o.g. each player is associated with a distinct vertex<sup>4</sup>. 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  $\xi_e$  on  $N$ .

Following previous work [22, 51], 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.

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  $\xi_e : 2^V \rightarrow \mathbb{R}_+^{|V|}$  to edge  $e$  based only on knowledge of  $G$  (without knowledge of  $S^5$ ) for the adversarial model, while in the stochastic model the method can in addition depend on  $\Pi$ . Due to the characterization in [22], we restrict ourselves to the family of generalized weighted Shapley protocols<sup>6</sup>.

**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  $U_i$  "arrives" before players from  $U_j$ . 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  $\pi$ .

**(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 the 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

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

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

<sup>4</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 Nash equilibrium are affected by this modification.

<sup>5</sup>The methods should be defined on  $V$ , since every vertex is potentially associated with some player.

<sup>6</sup>[22] characterizes the linear protocols (for every edge  $e$  of cost  $c_e \geq 0$ , it assigns the method  $c_e \cdot \xi$ , where  $\xi$  is the method it assigns to any edge of unit cost) to be the generalized weighted Shapley protocols. They further showed that for any non-linear protocol, there exists a linear one with at most the same PoA.

**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 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. We formally define (Definition 3.2) the ‘zig-zag’ pattern of the lower bounds of the Greedy Algorithm of the OSTP (see Example 1.1(a) and Figure 2). 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(a), we get the final lower bound construction.

**DEFINITION 3.1. (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  $j \in [r]$ ,  $D_j = \{v_i \mid (i \bmod 2^{r-j}) = 0 \text{ and } (i \bmod 2^{r-j+1}) \neq 0\}$ . For  $v \in D_j$ ,  $w \in D_{j'}$  and  $j < j'$ , we say that  $v$  belongs to a lower class than  $w$  (and  $w$  belongs to a higher class than  $v$ ).

As an example, consider the path  $P = (v_0, v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8)$ , 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  $\Pi(v_i) = \{w \mid d_P(v_i, w) = 2^{r-j}\}$ , i.e. the closest vertices that belong to lower classes. Remark that for all  $v \notin \{v_0, v_{2^r}\}$  i) the cardinality of  $\Pi(v)$  is 2, ii) the vertices of  $\Pi(v)$  belong to lower classes than  $v$ , iii) all vertices between  $v$  and any vertex of  $\Pi(v)$  belong to higher classes than  $v$ . We are now ready to define the ‘zig-zag’ pattern.

**DEFINITION 3.2. (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  $i \notin \{0, 2^r\}$ ,  $\pi(w) < \pi(v_i)$  for all  $w \in \Pi(v_i)$ .

An example of such a path for  $r = 3$  is shown in Figure 2. 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 2: An example of a  $P_3(\pi)$  path. The numbers correspond to the labels.

**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>7</sup>.

*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 3(a) shows such a labeling).

Next, we contract all such 2-edge paths of  $Q_m^2$  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,  $m < n$ ; 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.

*Coloring Scheme:* Consider  $Q_n$  as a bipartite  $Q_n = (A, B, E)$ . 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

<sup>7</sup>See  $Q_m^2$  of Definition 3.3 and Figure 3(a) for an illustration

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>8</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 3.1 we give the formal definition of the subgraph  $G_m$  of a hypercube  $Q_n$ . Section 3.2 is devoted to show that every 2-edge coloring of a (suitably) high dimensional hypercube contains a monochromatic copy of  $G_m$  (Lemma 3.2), by using Ramsey theory. Then, in Section 3.3 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 3.3). At last, in Section 3.4, 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 3.1). We then show how to construct a graph that serves as lower bound for all ordered protocols (Theorem 3.2). This is done by connecting any two edges of the hypercube with a direct edge of appropriate cost, similar to the example in Figure 1(a).

**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  $x = \langle x_1 \dots x_n \rangle$ , 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.  $xy = \langle x_1 \dots x_r y_1 \dots y_s \rangle$ .  $x = \langle x_j \rangle_{j=1}^r$  is the concatenation of its  $r$  bits. 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,

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

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  $\pi$ .

### 3.1 Description of $G_m$

For a positive integer  $m$ , we define a graph  $G_m = (V_m, E_m)$  that is a restriction 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:

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

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

<sup>8</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.

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  $x, x', y$ .  $\square$

### 3.2 Ramsey-type Theorem

LEMMA 3.2. *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>9</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>10</sup> using  $2^{4m}$  colors.

By Ramsey's Theorem for hypergraphs [35], 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  $R_1, R_2 \subset U$  and  $p \in [0, 4m - 1]$ , it is  $\chi(R_1, p) = \chi(R_2, p) = c_p$ . 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 with the edges that are determined by those prefix sums, contains a monochromatic copy of  $G_m$ . We will show this by constructing those edges from  $E_m$  (the set of edges of  $G_m$ ). By inserting blocks of 1's of suitable length among the bits of the edges of  $E_m$ , we construct the bits 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  $b_j$ . Each edge of  $E_m$  contains exactly  $2m$  zero bits. Also notice that  $\sum_j |\beta_j| = 2m$ . Therefore, in total we have  $6m$  bits (same as  $|U|$ ) and  $4m$  non-zero bits (same as  $|R|$ ). These  $6m$  bits are put at the coordinates of  $U$ . The rest  $n - 6m$  of the coordinates are filled with zeros. It is easy to check that, for every such edge, the prefix of the flip-bit is *always* one of the  $p_0, \dots, p_{2m-1}$ , implying that they are monochromatic. 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  $G_m$ .  $\square$

### 3.3 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 3.3. *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  $P(x, x')$ .*

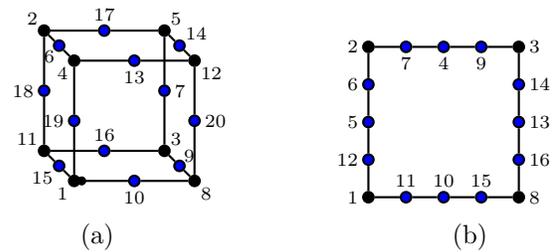


Figure 3: Examples of (a)  $Q_3^2$  and (b)  $Q_2^4$ . The labels on the nodes are examples of the labeling property, (a) after one inductive step, (b) after two inductive steps.

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).

LEMMA 3.3. *For any positive integer  $m$ , for all  $n \geq n_0 = g(m)$  and for any ordering  $\pi$  of  $V(Q_n)$ , there*

<sup>9</sup>The result could be extended to any (fixed) number of colors, but we need only two for our application.

<sup>10</sup>A  $k$ -uniform hypergraph is a hypergraph such that all its hyperedges have size  $k$ .

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 3.2. 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 3.2 implies that  $Q_n$  contains a monochromatic copy (blue or red) of  $G_m$ . Recall that  $G_m$  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  $G_m$  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 3.1). 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 1 and scaling factor  $1/2$ , by mapping the first  $m$  pairs of bits into single bits; map  $\langle 01 \rangle$  to 0 and  $\langle 10 \rangle$  to 1. 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  $d(x, x') = 2$ , Lemma 3.1 implies that there exists  $y = \theta(x, x') \in V_3$ , such that  $d(x, y) = d(x', y) = 1$ . 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 vertices  $V_2$ . 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 3.1 implies that there exists  $y = \theta(x, x') \in V_1$ , such that  $d(x, y) = d(x', y) = 1$ . 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.  $\square$

### 3.4 Lower Bound Construction

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

**THEOREM 3.1.** *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 3.2. 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 3.1.** *For every  $i \in [0, r]$ , and for any vertex ordering  $\pi$  of  $Q_{n_0}$ , it contains a subgraph isomorphic to  $Q_{n_i}^{2^i}$ , such that for every  $(x, x') \in Z(Q_{n_i}^{2^i})$ ,  $P(x, x')$  is a distance preserving path isomorphic to  $P_i(\pi)$ .*

*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 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 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 3.3 on  $Q_{n_i}$ , guarantees the existence of a subgraph  $W$  isomorphic to  $Q_{n_{i+1}}^{2^{i+1}}$  ( $n_i = g(n_{i+1})$ ), 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  $Q_{n_{i+1}}^{2^{i+1}}$  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 path  $P_{i+1}(\pi)$ .  $\square$

We now argue that the resulting  $P_r(\pi)$  is a distance preserving path. Our analysis indicate a sequence of hypercubes  $Q_{n_0}, Q_{n_1}, \dots, Q_{n_r}$ . Recall that in Lemma 3.3, in order to get  $Q_{n_{i+1}}$  from  $Q_{n_i}$  we mapped  $\langle 01 \rangle$  to 0 and  $\langle 10 \rangle$  to 1 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  $P_r(\pi)$ . 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 as shown in the example in Figure 1(a).

**THEOREM 3.2.** *Any ordered universal cost-sharing 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 3.1, 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(a)).

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 3 for the term “parent”).  $\square$

#### 4 Lower Bound for all universal protocols

In this section, we exhibit metric spaces for which no universal cost-sharing protocol admits a PoA better than  $\Omega(\log k)$ . Due to the characterization of [22], we can restrict ourselves in generalized weighted Shapley protocols (GWSPs). We follow the notation of [22], and for the sake of self-containment we include here the most related definitions and lemmas.

##### 4.1 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.

**LEMMA 4.1.** *(Lemma 4.10 of [22]). 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 4.2.** *(Lemma 4.11 of [22]). 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  $S \supseteq \{i, j\}$ ,  $\xi(j, S) \leq 1/(\beta + 1)$ .*

##### 4.2 Lower Bound

The following two technical lemmas will be used in our main theorem. The proofs are provided in the full version.

**LEMMA 4.3.** *Let  $X$  be a finite set of size  $m$ , 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  $i \in [m]$ .*

**LEMMA 4.4.** *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 no two distinct  $x, y \in S$  are consecutive, i.e. for every  $x, y \in S$ ,  $|x - y| \geq 2$ .*

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 4). 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 4.3 and 4.4 we show that there exists a set of potential players  $B$ , with *ordered-like* values, one at each vertex of the hypercube. Then, by using the results of Section 3, there exists a path where the vertices are *zig-zag-ordered*.

The separation into these two extreme cases was first used in [22]. 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 4.1.** *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 3.1. 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 4).

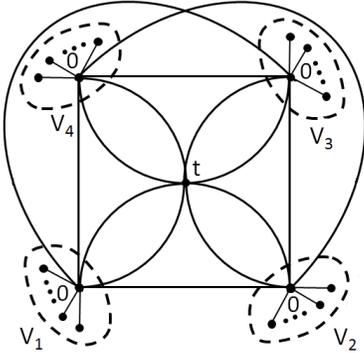


Figure 4: An example of  $Q^*$  for  $Q_2$  as the base hypercube.

We will show that any GWSP for  $Q^*$  has PoA  $\Omega(\log k)$ . 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 [22], 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  $R$ . Budget-balance implies that there exists some vertex  $i^* \in R$  which is charged at most  $1/k$  proportion of the cost. Moreover, Lemma 4.1 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.

Note that there is a NE where all players follow the edge  $e_{q2}$ , with cost  $2k \cdot k/6$ ; 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  $2k$ . Therefore, the PoA is  $\Omega(k)$  for this case.

**Ordered-like cost-sharing.** If there is no such  $R$  with at least  $k$  vertices, then  $|A_{js} \cap V_q| \leq k$  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 4.3 on the set  $N$ , for  $r = 4k^6 2^n$  and  $s = m = 2^n$ , by considering the subsets  $V_q$  as the partition of  $N$  (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 4.3 guarantees that for each  $V_q$  there exists  $V'_q \subset V_q$  of size  $2^n$ , such that every  $v \in V' = \cup_q V'_q$  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 4.4 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  $A_{js}$  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 3.1 guarantees that there always exists at least one distance preserving path  $P_r(\pi)$  (see Definition 3.2). Let  $S$  be the vertices of  $P_r(\pi)$  excluding the last class  $D_r$  (see Definition 3.1). The adversary will activate

this 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  $s_\ell \in \Pi(s_i)$  be one of  $s_i$ 's parents in the class hierarchy (we refer the reader to the beginning of Section 3); 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 4.1 that there exists a strategy profile  $\mathbf{P}^* \in \mathcal{P}'$  which is a NE.  $\mathbf{P}^*$  has cost:

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

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 4.1. *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 [22, 33]). 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  $P_2$  (arguments (b)), respectively. We further show that, given any strategy profile  $\hat{\mathbf{P}}$ , there exists some  $\mathbf{P} \in \mathcal{G}$  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 as follows:  $P_1 = e_{11}$ ,  $P_2 = (s_1, s_2) \cup P_1$  and from  $i = 3$  to  $k$ ,  $P_i \in A = \arg \min_{P'_i} \{c_i(\mathbf{P}^i, P'_i) | \exists (P'_{i+1}, \dots, P'_k) \text{ s.t. } (P_1, \dots, P_{i-1}, P'_i, \dots, P'_k) \in \mathcal{P}'\}$ . If  $\hat{P}_i \in A$  then we set  $P_i = \hat{P}_i$ , otherwise we choose a path from  $A$  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 4.2). Moreover,  $i$ 's share is at least  $\frac{8k-1}{8k} c_e$ .
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 class that  $s_{i'}$  belongs to. 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  $\mathbf{P}$  is at most  $2k$  and any other path would incur a cost strictly greater than  $2k$ .
- (b) The share of player 2 under  $\mathbf{P}$  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 rather than the shortcut  $(s_1, s_2)$ , the total cost of that path is at least  $2^r \left(\frac{k-1}{k}\right)^{r-1}$ . Player 2 is first according to  $\pi$  at that path 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 2 (the minimum distance between two activated vertices). By argument 1, her share is at least  $2 \frac{8k-1}{8k} > 1$ , where 1 is at most her share for  $P_\ell$  (argument 3).

- (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). However, all edges that are used by players that precedes  $i$  in  $\pi$  have cost at least  $\hat{c}_{r-j}$ . Therefore, in  $\hat{P}_i$ , player  $i$  is the first according to  $\pi$  for edges of total cost at least  $\hat{c}_{r-j+1}$ . This implies a cost-share of at least  $\frac{8k-1}{8k} \hat{c}_{r-j+1}$  (argument 1). But for  $k \geq 6$  and  $j < r$ ,  $\frac{8k-1}{8k} \hat{c}_{r-j+1} > \hat{c}_{r-j} + 1$ .

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  $\mathbf{P} \in \mathcal{P}'$ .  $\square$

## 5 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 OSTP can be embedded in an outerplanar graph (see Figure 5(a) for an illustration).

We next define an ordered universal cost-sharing protocol  $\Xi_{tour}$ , and we show that it has constant PoA. W.l.o.g. we assume that the metric space is defined by a given *biconnected* outerplanar graph<sup>11</sup>. Every biconnected graph admits a *unique* Hamiltonian cycle [49] that can be found in linear time [27].  $\Xi_{tour}$

<sup>11</sup>If it is not already biconnected, we turn it into an *equivalent* biconnected graph, by appropriately adding edges of infinity cost. By *equivalent* we mean that any NE outcome and the minimum Steiner tree solution remain unchanged after the transformation. Equivalence is obvious since we only add edges of infinity costs that cannot be used in neither any NE nor the minimum Steiner tree outcome.  $\square$

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 5(a),  $\pi(q_8) < \pi(q_4) < \pi(q_9) < \dots < \pi(q_{15})$ .

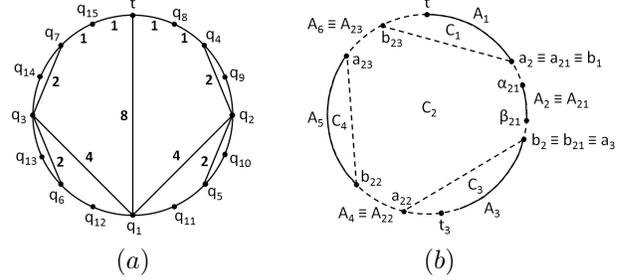


Figure 5: (a) shows an example of an outerplanar graph where the order  $q_i < q_{i+1}$  gives PoA of  $\Omega(\log k)$ . (b) illustrates some elements from the proof of Theorem 5.1, focusing on cycle  $C_2$ . The dashed components represent the optimum tree  $T^*$ .

As a warm-up, we first bound from above the PoA of  $\Xi_{tour}$  for cycle graphs, and then extend it to all outerplanar graphs.

LEMMA 5.1. *The PoA of  $\Xi_{tour}$  in cycle graphs is at most 2.*

*Proof.* Consider a cycle graph  $C = (V, E, t)$  and let  $S \subseteq V$  be the set of the activated vertices. Let  $T^*$  be the minimum Steiner tree (path) that connects  $S \cup \{t\}$ , and  $a, b$  be its two endpoints. Note that minimality of  $T^*$  implies that  $a, b \in S \cup \{t\}$ .  $a$  and  $b$  partition  $C$  into two paths ( $T^*, C \setminus T^*$ ) and  $t$  divides further  $T^*$  into two paths  $P'_1, P'_2$ . Let  $S_1 = \{u_1, \dots, u_r = a\}$  and  $S_2 = \{w_1, \dots, w_s = b\}$  be the activated vertices of  $P'_1$  and  $P'_2$ , respectively. W.l.o.g., assume that  $\pi(u_i) < \pi(u_{i+1})$  and  $\pi(w_{j+1}) < \pi(w_j)$ , for all  $i, j$ .

Consider any NE,  $\mathbf{P} = (P_i)_{i \in N}$ . We bound from above the share of each player  $v \neq w_s$ , by its distance from their immediate predecessor in  $\pi$ , as follows. By adopting the convention that  $u_0 = t$ ,

$$\begin{aligned} c_{u_i}(\mathbf{P}) &\leq d(u_i, u_{i-1}), \quad \forall i \in [r] \\ c_{w_j}(\mathbf{P}) &\leq d(w_j, w_{j+1}), \quad \forall j \in [s-1]. \end{aligned}$$

Also  $c_{w_s}(\mathbf{P}) \leq d(w_s, t)$ . Overall,

$$\begin{aligned} c(\mathbf{P}) &= \sum_{v \in S} c_v(\mathbf{P}) \\ &\leq \sum_{u_i \in S_1} d(u_i, u_{i-1}) + \sum_{w_j \in S_2 \setminus \{w_s\}} d(w_j, w_{j+1}) + d(w_s, t) \\ &\leq c(P'_1) + c(P'_2) + c(P'_2) \leq 2c(T^*). \end{aligned}$$

$\square$

**THEOREM 5.1.** *The PoA of  $\Xi_{tour}$  in outerplanar graphs is at most 8.*

*Proof.* Based on the previous discussion, it is sufficient to consider only biconnected outerplanar graphs with non-negative costs, including infinity. Let  $G = (V, E, t)$  be any such graph with  $S$  being the set of activated vertices.

Let  $T^*$  be the minimum Steiner tree that connects  $S \cup \{t\}$ , and  $C$  be the unique Hamiltonian tour of  $G$ , forming its outer face. Let  $E^* = E(T^*) \setminus E(C)$  be the set of non-crossing chords of  $C$  that belong to  $T^*$ . Then  $C \cup E^*$  forms  $|E^*| + 1 = r$  cycles  $C_1, \dots, C_r$ , where every pair  $C_i, C_j$  are either edge-disjoint or they have a single common edge belonging to  $E^*$ . On the other hand, each edge of  $C$  belongs to exactly one  $C_i$  and each edge of  $E^*$  belongs to exactly two  $C_i$ 's. Figure 5(b) provides an illustration.

For every  $i \in [r]$ , let  $S_i = (S \cup \{t\}) \cap V(C_i)$  be the activated vertices that lie in  $C_i$  and  $t_i$  be the vertex that is first in  $\pi$  among  $S_i$ . W.l.o.g. assume that, for all  $i \in [r-1]$ ,  $\pi(t_i) \leq \pi(t_{i+1})$  (then  $t_1 = t$ ). Also let  $T_i^*$  be the subgraph of  $T^*$  that intersects with  $C_i$ . Then  $T_i^*$  should be a path connecting  $S_i$ .

Consider any NE,  $\mathbf{P} = (P_i)_{i \in S}$ . We show separately that the shares of all  $S_i \setminus \{t_i\}$  are bounded by  $4c(T^*)$  and the shares of all  $t_i$ 's are bounded by  $4c(T^*)$ .

For the first case we use Lemma 5.1. For any cycle  $C_i$ , by considering  $t_i$  as the root, Lemma 5.1 provides a bound on the shares of  $S_i \setminus \{t_i\}$ . So,  $\sum_{v \in S_i \setminus \{t_i\}} c_v(\mathbf{P}) \leq 2c(T_i^*)$ . Recall, that each edge of  $E(T^*)$  belongs to at most two  $C_i$ 's, so by summing over all  $i \in [r]$ ,

$$\sum_{i \in [r]} \sum_{v \in S_i \setminus \{t_i\}} c_v(\mathbf{P}) \leq 2 \sum_{i \in [r]} c(T_i^*) \leq 4c(T^*).$$

The second case requires more careful treatment. The endpoints of the edges of  $E^*$  divide  $C$  into a partition of nonzero-length arcs,  $A_1, \dots, A_n$ , named based on their clockwise appearance in  $C$ , starting from an arc containing  $t$ . For every  $j \in [n]$ , let  $a_j$  and  $b_j$  be the two endpoints of  $A_j$ . The share of each  $t_i$  can be bounded by its distance from  $t_{i-1}$ , for  $i > 1$  (recall that  $t_1 = t$ ). Let  $A_{s_i}$  be an arc that  $t_i$  lies, then

$$\begin{aligned} \sum_{i \in [2, r]} c_{t_i}(\mathbf{P}) &\leq \sum_{i \in [2, r]} d(t_i, t_{i-1}) \\ &\leq \sum_{i \in [r]} (d(a_{s_i}, t_i) + d(t_i, b_{s_i})) + \sum_{j \in [n] \setminus \{s_1, \dots, s_r\}} d(a_j, b_j) \\ &= F. \end{aligned}$$

We next upper bound  $F$  by  $\sum_{i \in [r]} 2c(T_i^*)$ . Note that each arc  $A_j$  belongs to exactly one  $C_i$  and every

$C_i$  contains at least one such arc (otherwise  $T^*$  would have a cycle). We concentrate to a specific  $C_i$  and show that the portion of  $F$  associated with  $C_i$ 's arcs is upper bounded by  $2c(T_i^*)$ .

Let  $A_{i_1}, \dots, A_{i_m}$  be the arcs belonging to  $C_i$  and  $a_{ij}, b_{ij}$  be the endpoints of  $A_{ij}$ . Also let  $A_{i_s}$  be the arc containing  $t_i$ . Recall that  $T_i^*$  is a path and every edge of  $E(C_i) \cap E^*$  belongs to  $T_i^*$ . Therefore,  $T_i^*$  contains entirely all but one  $A_{ij}$ , say  $A_{i_m}$  (see Figure 5(b)). We examine the two cases of  $m = s$  and  $m \neq s$  separately.

**Case 1:  $m = s$ .**  $a_{i_s}, b_{i_s}$  (as endpoints of edges of  $E^*$ ) and  $t_i$  are vertices of the path  $T_i^*$ . Therefore, either some path from  $t_i$  to  $a_{i_s}$  or some path from  $t_i$  to  $b_{i_s}$  belongs to  $T_i^*$ ; w.l.o.g. assume that it is some path from  $t_i$  to  $a_{i_s}$ . Then  $\sum_{j \in [n_i], j \neq s} d(a_{ij}, b_{ij}) + d(t_i, a_{i_s}) \leq c(T_i^*)$ . Moreover, since  $b_{i_s}$  and  $t_i$  are vertices of  $T_i^*$ ,  $d(t_i, b_{i_s}) \leq c(T_i^*)$ .

**Case 2:  $m \neq s$ .** Similarly,  $\sum_{j \in [n_i], j \neq m, s} d(a_{ij}, b_{ij}) + d(t_i, a_{i_s}) + d(t_i, b_{i_s}) \leq c(T_i^*)$ . Also  $a_{i_m}$  and  $b_{i_m}$  are vertices of  $T_i^*$  and hence,  $d(a_{i_m}, b_{i_m}) \leq c(T_i^*)$ .

To sum up, in both cases it holds that

$$\sum_{j \in [n_i], j \neq s} d(a_{ij}, b_{ij}) + d(t_i, a_{i_s}) + d(t_i, b_{i_s}) \leq 2c(T_i^*).$$

By summing over all  $i$ ,  $F \leq \sum_{i \in [r]} 2c(T_i^*) \leq 4c(T^*)$ . Finally, by summing over the whole  $S$ ,  $c(\mathbf{P}) = \sum_{v \in S} c_v(\mathbf{P}) \leq 8c(T^*)$ .  $\square$

## 6 Stochastic Network Design

In this section we study the *stochastic* model, where the set of active vertices is drawn from some probability distribution  $\Pi$ . Each vertex  $v$  is activated independently with probability  $p_v$ ; the set of the activated vertices are no longer picked adversarially, but it is sampled based on the probabilities  $p_v$ 's, i.e., the probability that set  $S$  is active is  $\Pi(S) = \prod_{v \in S} p_v \cdot \prod_{v \notin S} (1 - p_v)$ . On the other hand, the probabilities  $p_v$ 's (and therefore  $\Pi$ ), are 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 show that there exists a *randomized* ordered protocol that achieves constant PoA. This result holds even for the *black-box* model [48], meaning that the probabilities are not known to the designer, however she is allowed to draw independent (polynomially many) samples. On the other hand, if we assume that the probabilities  $p_v$ 's are known to the designer, there exists a *deterministic* ordered protocol that achieves constant PoA. We note that both protocols can be determined in polynomial time.

The result for the randomized protocol depends on

approximation ratios of the minimum Steiner tree problem. More precisely, given an  $\alpha$ -approximate minimum Steiner tree, we show an upper bound of  $2(\alpha + 2)$ . 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 6.1. This algorithm and its slight variants have been used in different contexts: rent-or-buy problem [36], a priori TSP [48] and, stochastic Steiner tree problem [30].

**ALGORITHM 6.1. Randomized protocol  $\Xi_{rand}$**

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

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

- 1) Choose a random set of vertices  $R$  by drawing from distribution  $\Pi$  and construct an  $\alpha$ -approximate minimum Steiner tree,  $T_\alpha(R)$ , over  $R \cup \{t\}$ .
- 2) Connect all other vertices  $V \setminus V(T_\alpha(R))$  with their nearest neighbor in  $V(T_\alpha(R))$  (by breaking ties arbitrarily).
- 3) 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 6.1.** *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>12</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 6.1, 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 6.1). 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  $S$ . 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,

$$(6.1) \quad c(\mathbf{P}_R(S)) = \sum_{s_i \in S} c_{s_i}(\mathbf{P}_R(S)) \leq \sum_{s_i \in S} d_G(s_i, s_{i-1}) \\ \leq \sum_{s_i \in S} d_{T_{R,S}}(s_i, s_{i-1}) \leq 2c(T_{R,S}).$$

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,

$$(6.2) \quad c(T_{R,S}) = c(T_\alpha(R)) + \sum_{v \in S \setminus V(T_\alpha(R))} D_v(V(T_\alpha(R))) \\ \leq c(T_\alpha(R)) + \sum_{v \in S} D_v(R).$$

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:

$$(6.3) \quad \mathbb{E}_R[\mathbb{E}_S[\sum_{v \in V} I(v \in S) D_v(R)]] = \sum_{v \in V} \mathbb{E}_S[I(v \in S)] \mathbb{E}_R[D_v(R)] \\ = \sum_{v \in V} \mathbb{E}_S[I(v \in S)] \mathbb{E}_S[D_v(S)] = \mathbb{E}_S[\sum_{v \in V} I(v \in S) D_v(S)] \\ = \mathbb{E}_S[\sum_{v \in S} D_v(S)] \leq \mathbb{E}_S[c(MST(S))].$$

The last 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,  $D_v(S)$  is upper bounded by the minimum 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,

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

□

<sup>12</sup>The metric closure of a graph  $G$  is the complete undirected graph on the vertex set  $V(G)$ , where the edge costs equal the shortest paths in  $G$ .

By applying the 1.39-approximation algorithm of [17] we get the following.

COROLLARY 6.1.  $\Xi_{rand}$  has PoA at most 6.78.

THEOREM 6.2. *There exists a deterministic ordered protocol with PoA at most 16.*

We use derandomization techniques similar to [52, 48]. We provide the proof in the full version of this paper.

## 7 Conclusion

An interesting direction would be to get a deeper understanding of settings for which universal protocols provide better guarantees comparing to uniform protocols. On the one hand, our negative result for the adversarial setting shows that there exist graphs for which universal protocols behave no better than uniform protocols. On the other hand, our positive results demonstrate a separation between uniform and universal protocols. There exist classes of graphs, such as the outerplanar graphs, where universal protocols clearly outperform uniform protocols, with respect to price of anarchy bounds. A natural question that arises is what structural graph properties can be useful to the designer? We leave as an open question the characterization of the graphs, where such a separation exists.

## 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, Wroclaw, Poland, September 8-10, 2014. Proceedings*, pages 13–24, 2014.
- [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*, pages 12–23, 2012.
- [3] N. Alon and Y. Azar. On-line Steiner trees in the euclidean plane. *Discrete & Computational Geometry*, 10:113–121, 1993.
- [4] N. Alon, R. Radoicic, B. Sudakov, and J. Vondrák. A ramsey-type result for the hypercube. *Journal of Graph Theory*, 53(3):196–208, 2006.
- [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(4):1602–1623, 2008.
- [6] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.
- [7] B. Awerbuch, Y. Azar, and Y. Bartal. On-line generalized Steiner problem. *Theor. Comput. Sci.*, 324(2-3):313–324, 2004.
- [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*, pages 323–332, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [9] P. Berman and C. Coulston. On-line algorithms for Steiner tree problems (extended abstract). In *STOC*, pages 344–353, 1997.
- [10] D. Bertsimas and M. Grigni. On the spacefilling curve heuristic for the euclidean traveling salesman problem, 1988.
- [11] D. J. Bertsimas, P. Jaillet, and A. R. Odoni. A priori optimization. *Operations Research*, 38(6):pp. 1019–1033, 1990.
- [12] 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*, pages 121–134, 2014.
- [13] 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*, pages 186–197, 2014.
- [14] V. Bilò and R. Bove. Bounds on the price of stability of undirected network design games with three players. *Journal of Interconnection Networks*, 12(1-2):1–17, 2011.
- [15] 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(4):668–686, 2013.
- [16] 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*, pages 638–647, 2013.
- [17] 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*, pages 583–592, 2010.
- [18] I. Caragiannis. Efficient coordination mechanisms for unrelated machines scheduling. In *SODA*, pages 815–824, 2009.
- [19] M. Charikar, H. J. Karloff, C. Mathieu, J. Naor, and M. E. Saks. Online multicast with egalitarian cost sharing. In *SPAA*, pages 70–76, 2008.
- [20] C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications*, 25(6):1193–1206, 2007.
- [21] H. Chen and T. Roughgarden. Network design with weighted players. *Theory Comput. Syst.*, 45(2):302–324, 2009.
- [22] H. Chen, T. Roughgarden, and G. Valiant. Designing network protocols for good equilibria. *SIAM J. Comput.*, 39(5):1799–1832, 2010.

- [23] 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*, pages 86–97, 2009.
- [24] G. Christodoulou, E. Koutsoupias, and A. Nana-vati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.
- [25] G. Christodoulou, K. Mehlhorn, and E. Pyrga. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica*, 69:619–640, 2014.
- [26] 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*, pages 539–548, 2011.
- [27] G. Cornuéjols, J. Fonlupt, and D. Naddef. The traveling salesman problem on a graph and some related integer polyhedra. *Mathematical Programming*, 33(1):1–27, 1985.
- [28] 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*, pages 158–169, 2013.
- [29] 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*, pages 608–618, 2006.
- [30] 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*, pages 942–951, 2008.
- [31] 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*, pages 72–88, 2014.
- [32] M. X. Goemans and D. P. Williamson. A general approximation technique for constrained forest problems. *SIAM J. Comput.*, 24(2):296–317, Apr. 1995.
- [33] R. Gopalakrishnan, J. R. Marden, and A. Wierman. Potential games are necessary to ensure pure nash equilibria in cost sharing games. In *EC*, pages 563–564, 2013.
- [34] 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*, pages 178–191, 2010.
- [35] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey Theory, 2nd Edition*. Wiley Series in Discrete Mathematics and Optimization, 1990.
- [36] A. Gupta, A. Kumar, M. Pál, and T. Roughgarden. Approximation via cost sharing: Simpler and better approximation algorithms for network design. *J. ACM*, 54(3):11, 2007.
- [37] 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*, pages 649–658, 2006.
- [38] M. Imase and B. M. Waxman. Dynamic Steiner tree problem. *SIAM J. Discrete Math.*, 4(3):369–384, 1991.
- [39] N. Immorlica, L. Li, V. S. Mirrokni, and A. Schulz. Coordination mechanisms for selfish scheduling. In *Internet and Network Economics, First International Workshop (WINE)*, pages 55–69, 2005.
- [40] 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*, pages 386–395, 2005.
- [41] 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*, pages 613–623, 2000.
- [42] K. Kollias. Non-preemptive coordination mechanisms for identical machine scheduling games. In *SIROCCO*, pages 197–208, 2008.
- [43] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
- [44] 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*, pages 607–620, 2013.
- [45] 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(15):876–878, 2009.
- [46] L. K. Platzman and J. J. B. III. Spacefilling curves and the planar travelling salesman problem. *J. ACM*, 36(4):719–737, 1989.
- [47] F. Schalekamp and D. B. Shmoys. Algorithms for the universal and a priori TSP. *Oper. Res. Lett.*, 36(1):1–3, 2008.
- [48] D. B. Shmoys and K. Talwar. A constant approximation algorithm for the a priori traveling salesman problem. In *Integer Programming and Combinatorial Optimization, 13th International Conference, IPCO 2008, Bertinoro, Italy, May 26-28, 2008, Proceedings*, pages 331–343, 2008.
- [49] M. M. Syslo. Characterizations of outerplanar graphs. *Discrete Mathematics*, 26(1):47 – 53, 1979.

- [50] 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*, pages 1373–1387, 2015.
- [51] P. von Falkenhausen and T. Harks. Optimal cost sharing for resource selection games. *Math. Oper. Res.*, 38(1):184–208, 2013.
- [52] 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(6):707–712, 2007.