Manipulation-resistant facility location mechanisms for ZV-line graphs

Ilan Nehama*
Bar-Ilan University
Ramat-Gan, Israel
ilan.nehama@mail.huji.ac.il

Taiki Todo
Kyushu University
Fukuoka, Japan
RIKEN AIP
Tokyo, Japan
todo@inf.kyushu-u.ac.jp

Makoto Yokoo
Kyushu University
Fukuoka, Japan
RIKEN AIP
Tokyo, Japan
yokoo@inf.kyushu-u.ac.jp

ABSTRACT

In many real-life scenarios, a group of agents needs to agree on a common action, e.g., on a location for a public facility, while there is some consistency between their preferences, e.g., all preferences are derived from a common metric space. The *facility location* problem models such scenarios and it is a well-studied problem in social choice. We study mechanisms for facility location on unweighted undirected graphs, which are resistant to manipulations (*strategy-proof*, *abstention-proof*, and *false-name-proof*) by both individuals and coalitions and are efficient (*Pareto optimal*). We define a family of graphs, *ZV-line graphs*, and show a general facility location mechanism for these graphs which satisfies all these desired properties. Our result unifies the few works in the literature of false-name-proof facility location on discrete graphs including the preliminary (unpublished) works we are aware of.

KEYWORDS

Facility location; Strategy-proofness; False-name-proofness; ZV-line graphs

ACM Reference Format:
Ilan Nehama, Taiki Todo, and Makoto Yokoo. 2019. Manipulation-resistant facility location mechanisms for ZV-line graphs. In *Proc.of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

Reaching an agreement could be hard. The seminal works of Gibbard [10] and Satterthwaite [23] show that one cannot devise a general procedure for aggregating the preferences of strategic agents to a single outcome, besides trivial procedures that a-priori ignore all agents except one (that is, the outcome is based on the preference of a predefined agent) or a-priori rule out all outcomes except two (that is, regardless of the agents' preferences, the outcome is one of two predefined outcomes). The problem is that agents might act strategically aiming to get an outcome which they prefer, so there might be scenarios in which for any profile of actions (a possible agreement) at least one of the agents will prefer changing his action. Note that while we refer to a *procedure* and later to a *mechanism*, this impossibility is not technical but conceptual. We identify a procedure with the conceptual mapping induced by the procedure from the opinions of the agents to an agreement, while the procedure itself could be complex and abstract, e.g., to have several rounds or include a deliberation process between the agents (cheap-talk). For simplicity of terms, we refer to the *direct mechanism* which implements this mapping. That is, we think of an exogenous entity, the *designer*, who receives as input the opinions of the agents and returns as output the aggregated decision. This assumption does not hurt the generality, as according to the revelation principle [19] any general procedure is equivalent (w.r.t. the properties we study) to such a direct mechanism.

But in many natural scenarios, it is exogenously given that the preferences satisfy some additional rationality property, i.e., the mechanism should not be defined for any profile of preferences, giving rise to mechanisms that are not prone to the above drawbacks. Two prominent examples are *VCG mechanisms* and *generalized-median mechanisms*. VCG mechanisms [4, 11, 22, 27] are the mechanisms which are resistant to manipulations like the ones described above for scenarios in which the agents' preferences are quasi-linear with respect to money [15, Def. 3.b.7], and monetary transfers are allowed . The second example, *Generalized-median mechanisms*, do not include monetary transfers and have more of an ordinal flavor. Generalized-median mechanisms [16] are the mechanisms which are resistant to manipulations like above when it is known that the preferences are *single-peaked* w.r.t. the real line [3]. That is, the outcomes are *locations* on the real line, each agent has a unique optimal location, $\ell^\star$, and her preference over the locations to the right of $\ell^\star$ is derived by the proximity to $\ell^\star$, and similarly for the locations to the left of $\ell^\star$. For example, in the Euclidean single-peaked case, the preferences for all agents are minimizing the distance to their respective optimal locations.

The facility location problem

A natural generalization of the second scenario is the *facility location* problem. In this problem, we are given a metric space over the outcomes (that is, a distance function between outcomes) and it is assumed that the preference of each of the agents is defined by the distance to her optimal outcome: An agent with an optimal outcome $\ell^\star$ prefers outcome a over outcome b if and only if a is closer to $\ell^\star$ than b. For ease of presentation, throughout this paper we assume that there are finitely many agents and finitely many locations, and in Section 5 discuss the extension to the infinite case. In the finite case, a natural way to represent the common metric space is using

*Corresponding author, Main contributor.

Proc.of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

a weighted undirected graph. That is, having a vertex (location) for each outcome and weighted edges between vertices s.t. the distance between any two outcomes is equal to the distance between the two respective vertices. Roughly speaking, given such a graph one seeks to find a mechanism that on one hand will not a-priori ignore some of the agents or rule-out some of the locations, and on the other hand will be resistant to manipulations of the agents. Facility location problems, and moreover facility location problems for complex combinatorial structures, model many real-life scenarios of group decision making in which it is natural to assume some homogeneity between the different agents' preferences (e.g., an additional rationality assumption). These examples include not only locating a common facility, like a school, a bus-stop, or a library, but also more general agreement scenarios with a common metric, e.g., partition of a common budget to several tasks, committee selection, and group decision making with a multi-dimensional criteria. Following the common facility problem, we sometimes refer to the outcome of the mechanism as *the facility*. In this work, we look for mechanisms which satisfy the following desired properties:

Anonymity: The mechanism should not a-priori ignore agents and moreover we desire it to be a function of the agents' votes (which we refer to as *ballots*) but not their identities. Formally, the outcome of the mechanism should be invariant to any permutation of the ballots.

Citizen Sovereignty/Non-imposition *[1, 18]:* The mechanism should not a-priori rule out a location. Formally, the mapping to a facility location should be onto. Moreover, the mechanism should respect the preferences of the agents and aim to optimize the aggregated welfare of the agents.

Pareto optimality: The mechanism should not return a location ℓ if there exists a location ℓ' s.t. switching from ℓ to ℓ' will benefit one of the agents (move the facility closer to her) while not hurting any of the other agents. Note that any reasonable (monotone) notion of aggregated welfare optimization entails Pareto optimality.

Strategy-proofness: An agent should not be able to change the outcome to a location she strictly prefers by reporting a location different than her true location.

Abstention-proofness:[1] An agent should not be able to change the outcome to a location she strictly prefers by not casting a ballot.

False-name-proofness: An agent should not be able to change the outcome to a location she strictly prefers by casting more than one ballot.

False-name-manipulations received less attention in the classic social choice literature, since in most voting scenarios there exists a central authority that can enforce a 'one person, one vote' principle (but cannot enforce participation or sincere voting). In contrast, many of the voting and aggregation scenarios nowadays are run in a distributed manner on some network and include virtual identities or avatars, which can be easily generated, so a manipulation of an agent pretending to represent many voters is eminent.

Resistance to group manipulations: We also consider a generalization of the above three properties dealing with manipulations of coalitions of agents. We define the *preference of a coalition* as the unanimous preference of its members. That is, a coalition C weakly prefers an outcome a over an outcome b if all the members of C weakly prefer a over b.[2] We desire that a coalition should not be able to change the outcome to a location it strictly prefers by its members casting insincere ballots, abstaining, or casting more than one ballot. We note that for onto mechanisms this property entails Pareto optimality. Nevertheless, we prefer to think of Pareto optimality apart from this property due to the different motivations.

Our contribution

Besides the work of Todo et al. [25], who characterized the false-name-proof mechanisms for facility location on the continuous line and on continuous trees, we are not aware of other works dealing with characterizing false-name-proof mechanisms on a graph. Moreover, as far as we know, a false-name-proof mechanism is known to the community only for very few simple graphs, and the current knowledge is still highly preliminary. (When starting to work on this problem, we initially devised mechanisms for few of the examples we describe below - cycles, cliques, and the $2 \times n$ grid. We are not aware of any other previously-known positive results besides these graphs or small perturbations of them.)

In this paper we present a family of unweighted undirected graphs, which we name ZV-line graphs, and show a general mechanism for facility location over these graphs which satisfies the desired properties. To the best of our knowledge, this is the first work to show a general false-name-proof mechanism for a general family of graphs. Our mechanism for the ZV-line graphs family unifies the few mechanisms that are known and induces mechanisms for many other graphs. The mechanism is Pareto optimal and in particular satisfies citizen sovereignty and does not a-priori rule out any location; It is anonymous, so in particular no agent is ignored; But on the other hand, it is resistant to all the above manipulations.

Roughly speaking, in a ZV-line graph there are two types of locations Z and V (and we refer to them as Z-vertices and V-vertices, respectively), and the facility is 'commonly' (except if all agents unanimously agree differently) located on a Z-vertex. For instance, the Z-vertices could represent commercial locations for locating a public mall, or a set of status-quo outcomes.

In order to demonstrate the richness and naturality of this family, we give a series of common simple graphs which are studied in the literature. The full formal definition of ZV-line graphs is given in Section 3. Consider the following family of graphs (which is a subfamily of ZV-line graphs and captures the gist of our mechanism). Let $G = \langle \mathcal{V}, E \rangle$ be a bipartite unweighted undirected graph with vertex set $\mathcal{V}$ and edge set E. That is, there exists a partition of the vertices $\mathcal{V} = V \mathbin{\dot\cup} Z$ s.t. there are no edges between V-vertices and no edges between Z-vertices. In addition, we require that ***(a)*** there exists a predefined order over the Z-vertices, which we refer to as left-to-right order, and that ***(b)*** any of the V-vertices is connected to an interval (according to the order) of Z-vertices. Similarly to the

[1] In the voting literature (e.g., [5, 9, 17]) this property is also referred to as **voluntary participation** and the **no-show paradox**. This property is also equivalent to **individual-rationality** which takes a different point of view of mechanism design.

[2] Hence, C strictly prefers a over b if all the members of C weakly prefer a over b, and at least one member of C strictly prefers a over b.

single-peaked consistency case [3], one can think of this constraint as a homogeneity constraint over the agents' preferences. Our mechanism for such graphs:

- ▶ The mechanism returns the leftmost Pareto optimal Z-vertex,[3] if one exists.
- ▶ If no location in Z is Pareto optimal, then necessarily all agents voted for the same location, and the mechanism returns this location.

For example, bi-cliques (full bipartite graphs) can be represented as a ZV-line graph in which each V-vertex is connected to all the Z-vertices as follows (and we use below ○ for Z-vertices and ♦ for V-vertices):

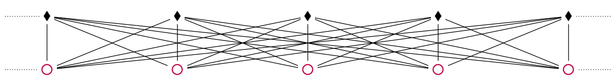

Our mechanism for this case:

- ▶ If all agents voted unanimously for the same location, the mechanism returns this location.
- ▶ If all agents voted for V-vertices, the mechanism returns the leftmost Z-vertex.
- ▶ Otherwise, the mechanism returns the leftmost Z-vertex that was voted for.

Notice that in this case the order over the Z-vertices is arbitrary (as well as the choice of one of the sides to be the Z-vertices) in the sense that it is not derived from the graph but a parameter of the mechanism. For instance, the order might represent the social norm of the society.

A second example is the discrete line graph, which can be represented as a ZV-line graph in which every two consecutive Z-vertices are connected by a unique V-vertex, 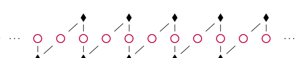 .

In particular, we show strategy-proof, false-name-proof, Pareto optimal mechanisms which are far from *generalized-median mechanisms* (for instance, in the common case the output of the mechanism belongs to a subset consisting of only half of the locations), in contrary to the characterization of these mechanisms for the continuous line [25, Thm. 2]. Dokow et al. [6, Thm. 3.4] characterized the strategy-proof mechanisms for the discrete line as a superset of generalized-median mechanisms, hence we get a strict subset of their characterization (and actually a small fraction of their characterization) due to requiring also false-name-proofness.

Two simple graphs that are generalizations of (the ZV-line graph representation of) the discrete line graph are 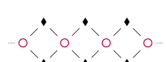 , in which every two consecutive Z-vertices are connected by two V-vertices, and the $2\times n$ grid  which can be represented as a ZV-line graph in which every three consecutive Z-vertices are connected by a unique V-vertex, i.e., 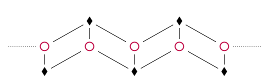 .

A common property to all the above examples is their regularity: All the V-vertices have the same degree and similarly all the Z-vertices have the same degree. An example we encountered of a non-regular graph for which a mechanism exists is 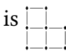, which can be represented as a non-regular ZV-line graph as 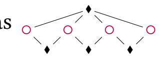 .

In the definition of the ZV-line graphs family (Def. 3.3) we extend the above family (and extend the mechanism accordingly) in two different ways: allowing edges between the Z-vertices (under a similar interval constraint), and replacing vertices by a tree, a clique, or any other ZV-line graph. For example,

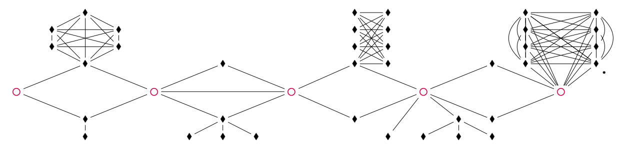

In particular, the ZV-line graphs family includes cycles of size up to 4 (note that there is no manipulation-resistant Pareto optimal anonymous mechanism for cycles of size larger than 5), trees, cliques, block graphs [12], and all graphs for which (as far as we found) a false-name-proof mechanism is known to the community. In Section 3.1, we show that for recursive sub-families like trees, cliques, or block graphs, a recursive mechanism is easily derived.

Related work

Problems of facility location on discrete graphs were also studied by Dokow et al. [6], who characterized the strategy-proof mechanisms for the discrete line and discrete cycle. Other variants of the facility location problem were also considered in the literature. For instance, Schummer and Vohra [24] considered the case of continuous graphs, Lu et al. [13, 14] studied variants in which several facilities need to be located and scenarios in which an agent is located on several locations, and Feldman et al. [8] studied the impact of constraining the input language of the agents.

False-name-proofness was first introduced by Yokoo et al. [28, (based on a series of previous conference papers)] in the framework of combinatorial auctions. In this work, the authors showed that the VCG mechanism does not satisfy false-name-proofness in the general case, and they proposed a property of the preferences under which this mechanism becomes false-name-proof. A similar concept was also studied in the framework of peer-to-peer systems by Douceur [7] under the name sybil attacks. Later, Conitzer [5] analyzed false-name-proof mechanisms in voting scenarios, Todo et al. [26] characterized other false-name-proof mechanisms for combinatorial auctions, and Todo et al. [25] characterized the false-name-proof mechanisms for facility location on the continuous line and on continuous trees. In a recent work, Ono et al. [20] showed, in the framework of facility location on the discrete line, a relation between false-name-proofness and the property of *population monotonicity.*

The characterization of manipulation-resistant mechanisms for facility location is highly related to problems in *Approximate mechanism design without money* [21]. In these problems, agents are characterized using cardinal utilities and the designer seeks to find an outcome maximizing a desired target function (e.g., sum of utilities, product of utilities, or minimal utility). These works bound the trade-of between the target function and manipulation-resistance, that is, they bound the loss to the target function due to

[3] That is, there exists no other location ℓ in the graph s.t. switching the outcome to ℓ benefits one of the agents while not hurting any of the other agents.

manipulation-resistance constraints. Similar bounds were derived for false-name-proof facility location mechanisms on the continuous line and tree by Todo et al. [25], strategy-proof facility location on the continuous cycle by Alon et al. [2], and for strategy-proof facility location on the discrete line and cycle by Dokow et al. [6].

Approximate mechanism design. In this work we do not analyze the approximation implications of the characterization and in particular we do not assume a specific cardinal representation of the agents' preferences. Yet, we claim that for most natural representations and target functions the approximation ratio is expected to be bad. For example, recall the above bi-clique example. In this mechanism, the facility might be located on an 'extremely' left Z-vertex. Moreover, the facility might be very far from the vast majority of the agents, resulting in a very bad approximation ratio for most reasonable target functions. This phenomenon is not specific for the bi-clique graphs. For most ZV-line graphs, (due to the false-name-proof requirement) the mechanism might be located on a location extremely far from almost all agents, resulting in a very bad approximation ratio (roughly, the number of agents times the girth of the graph) for most reasonable target functions.

2 MODEL

Consider a graph $G = \langle \mathcal{V}, E \rangle$ with a set of vertices $\mathcal{V}$ and a set of, neither weighted nor directed, edges $E \subseteq \binom{\mathcal{V}}{2}$, and we refer to the vertices $v \in \mathcal{V}$ also as *locations* and use the two terms interchangeably. The distance between two vertices $v, u \in \mathcal{V}$, notated $d(v, u)$, is the length of the shortest path connecting v and u, and the distance between a vertex $v \in \mathcal{V}$ and set of vertices $S \subseteq \mathcal{V}$, $d(v, S)$, is defined as the minimal distance between v and a vertex in S. For simplicity, we assume the graph is connected so the distance is finite, and in Section 5 discuss the extension to unconnected graphs. We define $B(v, d)$, the *ball* of radius $d \geqslant 0$ around a vertex $v \in \mathcal{V}$, to be the set of vertices of distance at most d from v, $B(v, d) = \{u \in \mathcal{V} \mid d(v, u) \leqslant d\}$. We say that two vertices are *neighbors* if there is an edge connecting them and notate by $N(v)$ the set of neighbors of a vertex v.

An instance of the *facility location problem over* G is comprised of n agents who are located on vertices of $\mathcal{V}$; Formally, we represent it by a *location profile* $\boldsymbol{x} \in \mathcal{V}^n$ where x_i is the location of Agent i. Given an instance $\boldsymbol{x}$, we would like to locate a facility on a vertex of the graph while taking into account the preferences of the agents over the locations. In this work, we assume the preference of an agent is defined by her distance to the facility: An agent located on $x \in \mathcal{V}$ strictly prefers the facility being located on $v \in \mathcal{V}$ over it being located on $u \in \mathcal{V}$ iff $d(x, v) < d(x, u)$.

A *general facility location mechanism* (or shortly a *mechanism*) defines for any profile of agents' locations a location for the facility. We require the mechanism to assign a location for the facility for any profile and any number of agents. Hence, we represent the mechanism by a function $F\colon \bigcup_{t \geqslant 0} \mathcal{V}^t \to \mathcal{V}$. We also think on F as a voting procedure: Each agent votes (and we also refer to his vote as a *ballot*) for a location, and based on the ballots F returns a location for the facility. We say that a mechanism is **anonymous** if the outcome $F(\boldsymbol{x})$ does not depend on the identities of the agents, i.e., it can be defined as a function of the ballot tally, the number of votes for each of locations.

Manipulation-resistance

A strategic agent might act untruthfully if she thinks it might cause the mechanism to return a location she prefers (that is, a location closer to her). In this work we consider the following manipulations: **Misreport**: An agent might report to the mechanism a location different from her real location; **False-name-report**: An agent might pretend to be several agents and submit several (not necessarily identical) ballots; **Abstention**: An agent might choose not to participate in the mechanism at all. A mechanism in which no agent benefits from these manipulations, regardless to the ballots of the other agents, is said to be **strategy-proof**, **false-name-proof**, and **abstention-proof**, respectively. We also consider a generalization of these manipulations to manipulations of a coalition, and say a mechanism is **group-manipulation-resistant** (shortly manipulation-resistant) if no coalition can change the outcome, by misreporting, false-name-reporting, or abstaining, to a different location which they unanimously agree is no worse than the original outcome (i.e., if they vote sincerely) and at least one of the coalition's members strictly prefers the new location.

Definition 2.1 (Group-manipulation-resistant). An anonymous mechanism F is group-manipulation-resistant if there exists no coalition of agents $C \subseteq \{1, \ldots, n\}$, a vector of locations $\boldsymbol{x} \in \mathcal{V}^n$, and a set of ballots[4] $\boldsymbol{A} \in \bigcup_{t \geqslant 0} \mathcal{V}^t$ s.t. $\boldsymbol{(i)}$ all the members of C weakly prefer $F(\boldsymbol{A}, \boldsymbol{x}_{-C})$, that is, the outcome when the agents outside of C do not change their vote and the agents of C replace their ballots by $\boldsymbol{A}$, over $F(\boldsymbol{x})$ and $\boldsymbol{(ii)}$ at least one of C's members strictly prefers $F(\boldsymbol{A}, \boldsymbol{x}_{-C})$ over $F(\boldsymbol{x})$.

We note that for $C = \{i\}$ being a singleton, this general manipulation coincides with misreport for $|\boldsymbol{A}| = 1$, with false-name-report for $|\boldsymbol{A}| > 1$, and with abstention for $\boldsymbol{A} = \varnothing$.

The revelation principle. One could consider more general mechanisms in which the agents vote using more abstract ballots, and define similar manipulation-resistance terms for the general framework. Applying a simple direct revelation principle [19] shows that any such general manipulation-resistant mechanism is equivalent to a manipulation-resistant mechanism in our framework: The two mechanisms implement the same mapping of the agents private preferences to a location for the facility, and since the above properties are defined for the mapping they are invariant to this transformation. That is, given some general mechanism M that maps abstract actions to a location for the facility and a behavior protocol D that maps types of the agents (i.e., locations) to actions of M, if D satisfies the generalized desiderata, then the direct mechanism $M \circ D$ satisfies our desiderata (w.r.t. truth-telling).

Efficiency

So far, we defined the desired manipulation-resistance properties for a mechanism. On the other hand, we would also like the mechanism to respect the preferences of the agents. We would like to avoid a scenario in which, after the mechanism has been used, the agents can agree that a different location is preferable. Given a location profile $\boldsymbol{x} \in \mathcal{V}^n$, the set of *Pareto optimal* locations, $PO(\boldsymbol{x})$, is the set of all locations which the agents cannot agree to rule out. Formally,

[4]Since F is an anonymous mechanisms, we define $\boldsymbol{A}$ as a set of ballots ignoring identities.

given two locations $v, u \in \mathcal{V}$, we say that u *Pareto dominates* v (w.r.t. a location profile $\boldsymbol{x}$) if $(\boldsymbol{i})$ all agents weakly prefer u over v and $(\boldsymbol{ii})$ at least one agent strictly prefers u over v. We say that v is *Pareto optimal* ($v \in PO(\boldsymbol{x})$) if it is not Pareto dominated by any other location. We say a mechanism is **Pareto optimal** if for any report profile $\boldsymbol{x}$ (and assuming truthful reporting) $F(\boldsymbol{x}) \in PO(\boldsymbol{x})$. In particular, Pareto optimality entails **unanimity**, if all the agents unanimously vote for the same location then the mechanism outputs this location, and **citizen sovereignty**, the mechanism is onto and does not a-priori rule out any location.

3 MAIN RESULT

In this work, we define a family of graphs, *ZV-line graphs*, and present a general mechanism for this family. This family is defined by introducing a simple combinatorial structure - a partition to two types of vertices and a connectivity constraint. One could think of the partition as representing a social agreement according to which the mechanism is defined, e.g., a subset of status-quo locations or an a-priori priority hierarchy over the locations. The connectivity constraint (as the graph in general) represents homogeneity over the agents' preferences, and allows us to construct a manipulation-resistant mechanism.

Definition 3.1 (ZV-ordered partition). Given an unweighted undirected connected graph $G = (\mathcal{V}, E)$, we say that a sequence of non-empty sets of vertices $Z, V_1, \dots, V_k \subseteq \mathcal{V}$ ($k \geqslant 0$) is a ZV-ordered partition if the following holds.

1. The sets V_i are disjoint, $V_i \cap V_j = \varnothing \quad$ for $i \neq j$.
2. The sequence is a cover of $\mathcal{V}$, $Z \cup V_1 \cup \dots \cup V_k = \mathcal{V}$, and no sub-sequence of it is a cover of $\mathcal{V}$.
3. For $i = 1, \dots, k$, there is a unique vertex in V_i which is closest to Z. We refer to it as *the root of* V_i and denote it by $\mathcal{R}(V_i)$,
$$\mathcal{R}(V_i) = \operatorname*{argmin}_{v \in V_i} d(v, Z).$$
4. All paths between vertices of V_i and vertices outside of V_i pass through the root $\mathcal{R}(V_i)$ and through Z.
5. Z is equipped with an order (that is, an injective mapping from Z to $\mathfrak{R}$). For simplicity of description, we refer to this order as an order from left to right. We call a subset A of Z an *interval* if it is a sequence of vertices according to the order, i.e., if A is the preimage of an interval in $\mathfrak{R}$.

We use the notions V_i-subgraphs, V-vertices, and Z-vertices for the respective sets of vertices. Note that we do not require the sets of Z-vertices and V-vertices to be disjoint. For instance, in the last example in the introduction the 9-clique includes the rightmost Z-vertex. Notice that from the third condition it is clear that for all i the intersection $V_i \cap Z$ is of size at most one.

For instance, in all examples in the introduction but the last one: The V_i-subgraphs consist of single vertices (the ⧫ vertices), which are also the roots of the respective subgraphs; The Z-vertices are not connected to each other and are ordered on a horizontal line. In the last example in the introduction, there are 10 disjoint V_i-subgraphs[5] and for two of them the root is a Z-vertex.

[5]Note that one could also define the tree of V-vertices in the bottom right of the figure as two disjoint V_i-subgraphs.

Given a graph $G = (\mathcal{V}, E)$ with a ZV-ordered partition, $Z, V_1, \dots, V_k \subseteq \mathcal{V}$, and mechanisms $F_i \colon \bigcup_{t \geqslant 0} (V_i)^t \to V_i$ for $i = 1, \dots, k$, we define the following mechanism $\mathcal{F}^\star \colon \bigcup_{t \geqslant 0} \mathcal{V}^t \to \mathcal{V}$.

Definition 3.2 ($\mathcal{F}^\star$). Given a vector of reports $\boldsymbol{x} \in \bigcup_{t \geqslant 0} \mathcal{V}^t$

- ▶ If all ballots belong to the same V_i-subgraph, return $F_i(\boldsymbol{x})$.
- ▶ Otherwise, return the leftmost Pareto optimal location in Z.

Notice that $\mathcal{F}^\star$ is defined w.r.t. a specific ZV-ordered partition, so when G can be represented as a ZV-line graph w.r.t. several ZV-ordered partitions, e.g., when G is a bi-clique, different mechanisms could arise. It is also important to note that there is no assumption that the agents 'know' the ZV-ordered partition of the graph (but they know the mechanism $\mathcal{F}^\star$). In other words, we see this structure as a combinatorial property of a graph which derives the agents' preferences, and could represent some homogeneity of the preferences or a social norm which motivates giving priority to the Z-vertices.

It is not hard to see the following: • $\mathcal{F}^\star$ is well defined: Unless all ballots belong to the same V_i-subgraph, there exist two locations, which belong to two different V_i-subgraphs, and a shortest path between them s.t. all its vertices are in $PO(\boldsymbol{x})$, so $PO(\boldsymbol{x}) \cap Z \neq \varnothing$, • $\mathcal{F}^\star$ runs in polynomial time (since finding $PO(\boldsymbol{x})$ can be done in time $|\mathcal{V}|^2 \cdot |\mathcal{V}|$ by iterating over all pairs to find the Pareto-dominated locations, and • If $F_1, \dots, F_k$ can be defined as the first Pareto optimal location according to some order, then an equivalent way to define $\mathcal{F}^\star$ is as the first Pareto optimal location in the following order: First, go over the Z-vertices from left to right, and then on the V-vertices in some order s.t. for each subgraph the order over its vertices matches the order of F_i.

Next, we define ZV-line graphs by introducing a connectivity constraint.

Definition 3.3 (ZV-line graph). An unweighted undirected connected graph $G = (\mathcal{V}, E)$ is a ZV-line graph w.r.t. $\mathcal{V} = Z \cup (V_1 \dot\cup \cdots \dot\cup V_k)$, if $(\boldsymbol{a})$ $\langle Z, V_1, \dots, V_k \rangle$ is a ZV-ordered partition of G, $(\boldsymbol{b})$ for any vertex $z \in Z$, $B(z, 1) \cap Z$ is an interval in Z, and if $k > 0$ then for $i = 1, \dots, k$ $(\boldsymbol{c})$ the induced graph $G_i = \langle V_i, E \cap (V_i \times V_i) \rangle$ is a ZV-line graph, $(\boldsymbol{d})$ $\mathcal{R}(V_i)$ is a Z-vertex of G_i (that is a Z-vertex in the representation of G_i), and it is the leftmost Z-vertex of G_i, and last $(\boldsymbol{e})$ $B(\mathcal{R}(V_i), 1) \cap Z$ is an interval in Z.

For example, in all examples in the introduction but the last one: • The Z-vertices are not connected to each other so $B(z, 1) = \{z\}$ for all $z \in Z$; • For any V_i-subgraph, V_i is a singleton, so $\langle V_i, E \cap (V_i \times V_i) \rangle$ is the trivial ZV-line graph which consists of a single vertex (so also Prop. $(\boldsymbol{d})$ holds); and • each V-vertex is connected to an interval of Z-vertices. Note that the first example, the bi-clique, is a ZV-line graph w.r.t. any order on the Z-vertices. Similarly, for any $\ell \geqslant 1$ the clique over ℓ vertices, K_ℓ, is a ZV-line graph w.r.t. $Z = \mathcal{V}$ and any order over the vertices, so the mechanism $\mathcal{F}^\star$ for K_ℓ returns the first location which was voted for.

Given a ZV-line graph $G = \langle \mathcal{V}, E \rangle$, applying Def. 3.2 recursively on G and its V_i-subgraphs gives us a mechanism $\mathcal{F}^\star \colon \bigcup_{t \geqslant 0} \mathcal{V}^t \to \mathcal{V}$. Our main result shows that $\mathcal{F}^\star$ satisfies the desired properties.

Theorem 3.4 (Main result). *Let $G = (\mathcal{V}, E)$ be a ZV-line graph w.r.t. $\mathcal{V} = Z \cup (V_1 \dot\cup \cdots \dot\cup V_k)$ and let $\mathcal{F}^\star \colon \bigcup_{t \geqslant 0} \mathcal{V}^t \to \mathcal{V}$ be the result of applying Definition 3.2. recursively on G. Then $\mathcal{F}^\star$ is an anonymous Pareto optimal mechanism and $\mathcal{F}^\star$ satisfies:*

For any vector of locations $\boldsymbol{x} \in \mathcal{V}^n$, a coalition of agents C, and a set of ballots $\boldsymbol{A} \in \bigcup_{t \geqslant 0} \mathcal{V}^t$, $\boldsymbol{A}$ is not a beneficial deviation for C.

Note that the theorem does not hold for weighted graphs. Consider the following weighted graph and a profile in which Alice is located on z_r and Bob on v. Then, the outcome of $\mathcal{F}^\star$ is z_r, but Bob can move the facility to a preferred location z_ℓ both $\boldsymbol{(i)}$ by misreporting z_ℓ, hence $\mathcal{F}^\star$ is not strategy-proof, and $\boldsymbol{(ii)}$ by false-name-reporting z_ℓ in addition to his sincere report, hence $\mathcal{F}^\star$ is not false-name-proof.[6]

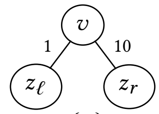

$V = \{v\}$
$Z = \{z_\ell, z_r\}$

There are trivial mechanisms which satisfy subsets of these properties: • The *fixed mechanism*, which locates the facility on a pre-defined location ignoring the votes, is trivially manipulation-resistant and anonymous, but it is not onto and hence not Pareto optimal. • A *dictatorship* of of the first agent, i.e., a mechanism which always locates the facility on the location reported by the first agent, is not anonymous but clearly it is manipulation-resistant.[7] • The *median mechanism*, which minimizes the sum of distances between the facility and the ballots, is anonymous and Pareto optimal, and it is not hard to see that for the discrete line it satisfies strategy-proofness and abstention-proofness both against one manipulator and against a coalition but an agent will benefit by casting multiple identical ballots. • The *mean mechanism*, which minimizes the sum of squares of the distances between the facility and the ballots, is anonymous and Pareto optimal but might not be strategy-proof or false-name-proof even against one agent, e.g., for the discrete line graph (it is abstention-proof, though).

3.1 Implications: Mechanisms for recursive graph families

By applying the main result to a recursive graph family, we can generate a recursive (and hence commonly simple) mechanism which satisfies our desiderata. For instance, a corollary of our result is a manipulation-resistant mechanism for the following family of rooted graphs (that is, $\langle \mathcal{V}, E, r\rangle$ s.t. $E \subseteq \binom{\mathcal{V}}{2}$ and $r \in \mathcal{V}$).

Definition 3.5 ($\mathcal{F}$).

- $\langle \{v\}, \emptyset, v\rangle \in \mathcal{F}$.
- For any $k, \ell \geqslant 1$: If $\{\langle \mathcal{V}_i, E_i, r_i\rangle\}_{i=1}^k$ are in $\mathcal{F}$ (and the $\mathcal{V}_i$ are disjoint), then also the following graph is in $\mathcal{F}$.

$$\left(\left\{\widehat{r}_j\right\}_{j=1}^{\ell} \dot\cup \left(\dot{\bigcup_{i=1}^{k}} \mathcal{V}_i\right), \left\{\left\langle \widehat{r}_j, r_i\right\rangle\right\}_{i=1\ldots k, j=1\ldots\ell} \dot\cup \left(\dot{\bigcup_{i=1}^{k}} E_i\right), \widehat{r}_1\right)$$

I.e., adding a new layer of pre-roots, a bi-clique between them and the roots of the graphs of the previous stage, and defining one of the pre-roots to be the new root.

Claim 3.6. The anonymous Pareto optimal mechanism $F(\boldsymbol{x}) = \operatorname{argmin}_{v \in PO(\boldsymbol{x})} d(v, r)$, which returns the Pareto optimal location closest to the root and breaks ties according to a predefined order, is manipulation-resistant.

Note that by setting $\ell = 1$ in the second step of the definition we get a recursive definition of *rooted trees*. Hence, we get that for any tree G the mechanism that returns the lowest common ancestor of the ballots (with regard to some root) is a manipulation-resistant mechanism (These are also the mechanisms which Todo et al. [25] characterized as the false-name-proof, anonymous, and Pareto optimal mechanisms for the continuous tree.).

Proof. We prove the claim by induction over, $h(G)$, the number of steps needed to generate G.

If $h(G) = 0$, i.e., $G = \langle \{v\}, \emptyset, v\rangle$ consists of a single vertex and the trivial mechanism satisfies all the desired properties.

If $h(G) \geqslant 1$, then G is a ZV-line graph w.r.t. $Z = \left\{\widehat{r}_j\right\}_{j=1}^{\ell}$ and $V_i = \mathcal{V}_i$. Note that for all V_i-subgraphs $h(\langle \mathcal{V}_i, E_i, r_i\rangle) \leqslant h(G) - 1$. Hence, our recursive mechanism returns one of the pre-roots of the 'lowest' subgraph which includes $\boldsymbol{x}$ when ties are broken according to the (arbitrary) order over the pre-roots. □

A second example is *connected block graphs [12]*.[8] A connected graph $G = \langle \mathcal{V}, E\rangle$ is a block graph if the following equivalent conditions hold:

- Every biconnected component of G is a clique. (Since for any graph the structure of its biconnected components is described by a block-cut tree,[9] connected block graphs are also called *clique trees.*)
- The intersection of any two connected subgraphs of G is either empty or connected.
- For every four vertices $u, v, w, x \in \mathcal{V}$, the larger two of the distance sums $d(u,v) + d(w,x)$, $d(u,w) + d(v,x)$, and $d(u,x) + d(v,w)$ are equal.

For a connected block graph, $\mathcal{F}^\star$ returns the Pareto optimal location closest to an arbitrarily predefined location, breaking ties according to an arbitrarily predefined order over the locations.

Proof sketch. G is connected block graph and hence all biconnected components of G are cliques. The *block-cut tree* of G is a tree $\mathcal{T}(G)$ which is defined in the following way. In $\mathcal{T}(G)$ there is a vertex (*component-vertex*) for each maximal biconnected component of G and a vertex (*intersection-vertex*) for each vertex in G which belongs to more than one maximal biconnected component. There is an edge in $\mathcal{T}(G)$ between each component-vertex and the intersection-vertices belonging to this component.

Following the inductive structure of $\mathcal{T}(G)$, and recalling that a clique is a ZV-line graph w.r.t. all vertices of the clique being Z-vertices and any order over them, we get that our mechanism is defined by an arbitrary predefined component-vertex of $\mathcal{T}(G)$, $\mathcal{R}$, and a series of arbitrary predefined orders over the locations of each of the components. The mechanism is:

- ► If all ballots belong to the same component, return the first location (according to the order) that was voted for.

[6]The mechanism which returns the leftmost ballot according to the order $z_\ell - v - z_r$ satisfies the desiderata. Notice that this mechanism can be defined as $\mathcal{F}^\star$ w.r.t. a ZV-ordered partition with $Z = \{z_\ell, v, z_r\}$.

[7]While we did not formally define false-name-proofness for non-anonymous mechanisms, assuming a false-name-ballot cannot be counted as the vote of the first agent, no agent can benefit from casting additional ballots.

[8]We thank Ayumi Igarashi for suggesting us this family as an example.

[9]The *block-cut tree* of a graph G is a tree $\mathcal{T}(G)$ which is defined in the following way. In $\mathcal{T}(G)$ there is a vertex (*component-vertex*) for each maximal biconnected component of G and a vertex (*intersection-vertex*) for each vertex in G which belongs to more than one maximal biconnected component. There is an edge in $\mathcal{T}(G)$ between each component-vertex and the intersection-vertices belonging to this component.

- ► Otherwise, choose the component closest to $\mathcal{R}$ s.t. one of the locations of the component is Pareto optimal, and return the first location (according to the order) in this component.

An equivalent definition of this mechanism is returning the closest Pareto optimal location to some location $v \in \mathcal{R}$, breaking ties according to a concatenation of the orders over the components. □

4 PROOF OF MAIN RESULT (THM. 3.4)

We prove a stronger result which shows a general method for generating a mechanism $\mathcal{F}^\star$ (satisfying the desired properties) for a given graph from mechanisms for its V_i-subgraphs. Theorem 3.4 is an immediate special case of this theorem. The same proof shows that also for weaker manipulation-resistance properties, e.g., against individual agents, against misreporting, or against abstentions, manipulation-resistance of the mechanisms for the V_i-subgraphs, result in the same manipulation-resistance notion for $\mathcal{F}^\star$, the mechanism of the graph.

Theorem 4.1. *Let $G = (\mathcal{V}, E)$ be a graph with a ZV-ordered partition $\mathcal{V} = Z \cup (V_1 \dot\cup \cdots \dot\cup V_k)$ and let $F_i\colon \bigcup_{t\geqslant 0} (V_i)^t \to V_i$ be a sequence of mechanisms s.t. for $i = 1, \ldots, k$*

- *F_i is anonymous and Pareto optimal;*
- *For an infinite number of $\tau \in \mathbb{N}$, there exists a profile $\boldsymbol{x} \in \bigcup_{t\geqslant 0} (V_i)^t$ in which all locations in V_i were voted for at least τ times and $F_i(\boldsymbol{x}) = \mathcal{R}(V_i)$; and*
- *For any vector of locations $\boldsymbol{x} \in (V_i)^n$, a coalition of agents C, and a set of ballots $\boldsymbol{A} \in \bigcup_{t\geqslant 0} (V_i)^t$, $\boldsymbol{A}$ is not a beneficial deviation for C.* (★)

Then, for $\mathcal{F}^\star\colon \bigcup_{t\geqslant 0} \mathcal{V}^t \to \mathcal{V}$ as defined in Definition 3.2, $\mathcal{F}^\star$ is an anonymous and Pareto optimal mechanism and

- *(I) If G is a ZV-line graph w.r.t. $\mathcal{V} = Z \cup (V_1 \dot\cup \cdots \dot\cup V_k)$, then $\mathcal{F}^\star$ satisfies (★).*
- *(II) If $\mathcal{R}(V_i) \in Z$ for all $i = 1, \ldots, k$, and the mechanism $F_Z\colon \bigcup_{t\geqslant 0} Z^t \to Z$ which returns the leftmost Pareto optimal location satisfies (★), then also $\mathcal{F}^\star$ satisfies (★).*

Proof. The anonymity of $\mathcal{F}^\star$ is an immediate corollary of the mechanisms F_i and F_Z being anonymous mechanisms. Notice that if all agents are in the same V_i-subgraph, then all of them strictly prefer $\mathcal{R}(V_i)$ over any location outside of V_i, so $PO(\boldsymbol{x}) \subseteq V_i$. Moreover, any location $v \in V_i \backslash PO(\boldsymbol{x})$ is Pareto dominated by a location $y \in PO(\boldsymbol{x}) \subseteq V_i$. Hence, the Pareto optimal set when considering only the locations in V_i equals to the Pareto optimal set when considering all locations. Since, the mechanisms F_i are Pareto optimal mechanisms we get that also $\mathcal{F}^\star$ is Pareto optimal.

In order to prove the main part of the theorem, we assume towards a contradiction that there exists a vector of locations $\boldsymbol{x} \in \mathcal{V}^n$, a coalition of agents C, and a set of ballots $\boldsymbol{A} \in \bigcup_{t\geqslant 0} \mathcal{V}^t$, s.t. C can, by voting $\boldsymbol{A}$, get an outcome $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$ which it strictly prefers, that is, all of its members weakly prefer $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$ over $\mathcal{F}^\star(\boldsymbol{x}) = \mathcal{F}^\star(\boldsymbol{x}_C, \boldsymbol{x}_{-C})$, and at least one of C's members, Agent i for $i \in C$, strictly prefers $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$ over $\mathcal{F}^\star(\boldsymbol{x})$. $\mathcal{F}^\star(\boldsymbol{x}) \in PO(\boldsymbol{x})$ and in particular the coalition of all agents does not strictly prefer $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$ over $\mathcal{F}^\star(\boldsymbol{x})$. Hence, there exists an Agent j, for $j \notin C$, who strictly prefers $\mathcal{F}^\star(\boldsymbol{x})$ over $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$.

If $\mathcal{F}^\star(\boldsymbol{x})$ is not in Z: Then necessarily, all the locations in $\boldsymbol{x}$ and $\mathcal{F}^\star(\boldsymbol{x})$ belong to the same V_i-subgraph, w.l.o.g. V_1, so $\mathcal{F}^\star(\boldsymbol{x}) = F_1(\boldsymbol{x})$. Since F_1 is resistant to false-name manipulations of Agent i and since Agent i can achieve $\mathcal{R}(V_1)$ by casting enough false ballots, we get that Agent i weakly prefers $\mathcal{F}^\star(\boldsymbol{x})$ over $\mathcal{R}(V_1)$ and hence Agent i strictly prefers $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$ over $\mathcal{R}(V_1)$. Since for any u outside of V_1 it holds that $d(x_i, \mathcal{R}(V_1)) < d(x_i, u)$, we get that $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C}) \in V_1 \backslash \mathcal{R}(V_1) \subseteq V_1 \backslash Z$. Hence, $\boldsymbol{A} \subseteq V_1$ and $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C}) = F_1(\boldsymbol{A}, \boldsymbol{x}_{-C})$, and we get a contradiction to the false-name-proofness of F_1. The dual case, $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$ is not in Z, is symmetric to the above.

If both $\mathcal{F}^\star(\boldsymbol{x})$ and $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$ are in Z: We deal with the two cases separately.

(*I*) G is a ZV-line graph w.r.t. $\mathcal{V} = Z \cup (V_1 \dot\cup \cdots \dot\cup V_k)$: We first prove the following two auxiliary lemmas.

Lemma i. *For any $v \in \mathcal{V}$ and $d \geqslant 0$, $B(v, d) \cap Z$ is an interval in Z.*

Proof. We prove the lemma by induction over d.
$d = 0$: $B(v, 0) \cap Z$ is either the empty set or $\{v\}$.
$d = 1$: $B(v, 1) \cap Z$ is either the empty set or an interval in Z.
$d \geqslant 2$: If $d < d(v, Z)$, $B(v, d) \cap Z = \emptyset$. If $d \geqslant d(v, Z) > 1$ (in particular, $v \notin Z$ and is not a root), then we take u to be the root of v's V_i-subgraph and note that all paths from v to locations in Z pass through u, $1 \leqslant d(v, u) \leqslant d(v, Z) \leqslant d$ and

$$B(v, d) \cap Z = B(u, d - d(v, u)) \cap Z$$

which is an interval by the induction hypothesis. Otherwise, $d(v, Z) \leqslant 1 < d$ and in particular $B(v, d) \cap Z \neq \emptyset$, and hence

$$B(v, d) \cap Z = (B(v, 1) \cap Z) \cup \left(\bigcup_{\substack{u \in N(v) \text{ s.t.} \\ d(u, Z) \leqslant 1}} B(u, d - 1) \cap Z \right).$$

For any $u \in N(v)$ s.t. $d(u, Z) \leqslant 1$ we claim that $B(u, d-1) \cap Z$ and $B(v, 1) \cap Z$ intersect.

- If $u \in Z$: $u \in (B(u, d - 1) \cap Z) \cap (B(v, 1) \cap Z)$.
- If $u \notin Z$: then $v \in Z$ and $v \in (B(u, d - 1) \cap Z) \cap (B(v, 1) \cap Z)$.

Hence, for any $u \in N(v)$ s.t. $d(u, Z) \leqslant 1$, $B(u, d - 1) \cap Z$ and $B(v, 1) \cap Z$ are intersecting intervals in Z. So $B(v, d) \cap Z$ is an interval as the union of intersecting intervals. □

Lemma ii. *Let $\boldsymbol{x}$ be a vector of locations s.t. $\mathcal{F}^\star(\boldsymbol{x}) \in Z$ and let $v \in Z$ be a location s.t. Agent i strictly prefers v over $\mathcal{F}^\star(\boldsymbol{x})$. Then $\mathcal{F}^\star(\boldsymbol{x})$ is to the left of v.*

Proof. If $x_i \in Z$ then $x_i \in PO(\boldsymbol{x}) \cap Z$ and by the definition of $\mathcal{F}^\star$, $\mathcal{F}^\star(\boldsymbol{x})$ is to the left of x_i. Since $\mathcal{F}^\star(\boldsymbol{x}) \notin B(x_i, d(x_i, v)) \cap Z$ and since this set is an interval which includes x_i, we get that $\mathcal{F}^\star(\boldsymbol{x})$ is to the left of the interval and in particular to the left of v.

Otherwise, $x_i \notin Z$ and there exists an Agent k for which x_k is not in the same V_i-subgraph as x_i. Hence, there exists a location $u \in Z$ s.t. u is on a shortest-path from x_i to x_k, $u \in Z$, and $u \in PO(\boldsymbol{x})$. Hence, $d(x_i, u) \leqslant d(x_i, v)$ and so

$$u \in B(x_i, d(x_i, u)) \cap Z \subseteq B(x_i, d(x_i, v)) \cap Z.$$

The two sets are intervals in Z, $\mathcal{F}^\star(\boldsymbol{x})$ is to the left of u (or equal to it), and $\mathcal{F}^\star(\boldsymbol{x}) \notin B(x_i, d(x_i, v)) \cap Z$. Hence, $\mathcal{F}^\star(\boldsymbol{x})$ is to the left of v. □

By applying Lemma ii for the profile $\boldsymbol{x}$ and Agent i, we get that $\mathcal{F}^\star(\boldsymbol{x})$ is to the left of $\mathcal{F}^\star(\boldsymbol{A}, \boldsymbol{x}_{-C})$; and by applying Lemma ii for the profile $(\boldsymbol{A}, \boldsymbol{x}_{-C})$ and Agent j, we get that $\mathcal{F}^{\star}(\boldsymbol{A}, \boldsymbol{x}_{-C})$ is to the left of $\mathcal{F}^{\star}(\boldsymbol{x})$. Hence, we get a contradiction.

$\underline{(\boldsymbol{II})\ \mathcal{R}(V_i) \in Z \text{ for } i = 1, \ldots, k \text{ and } F_Z \text{ satisfies } (\bigstar)}$: Then the preference of an agent which is located in a V_i-subgraph over the locations in Z is identical to the preference of an agent which is located on $\mathcal{R}(V_i)$. Hence, for any profile $\boldsymbol{y}$ if $\mathcal{F}^{\star}(\boldsymbol{y}) \in Z$ then $\mathcal{F}^{\star}(\boldsymbol{y}) = F_Z(\widehat{\boldsymbol{y}})$ for $\widehat{\boldsymbol{y}}$ being the profile generated from $\boldsymbol{y}$ by replacing each ballot outside of Z with the root of its V_i-subgraph. Therefore, for the profile $\widehat{\boldsymbol{x}} \in Z^n$ the coalition C can, by voting $\widehat{\boldsymbol{A}}$, get an outcome $F_Z\left(\widehat{\boldsymbol{A}}, \widehat{\boldsymbol{x}}_{-C}\right)$ which it strictly prefers over $F_Z(\widehat{\boldsymbol{x}})$, in contradiction to F_Z satisfying $(\bigstar)$. □□

5 SUMMARY & FUTURE WORK

In this work, we presented a new family of graphs, ZV-line graphs, and a generic anonymous Pareto optimal manipulation-resistant mechanism for the facility location problem on these graphs. To the best of our knowledge, the (very few) false-name-proof mechanisms which are currently known are for specific graphs and this work is the first to show a generic false-name-proof mechanism for a large family, utilizing a broad graph property and unifying all existence results which we are aware of. The construction of the mechanism is inductive: We derive a mechanism for a given ZV-line graph from mechanisms for its subgraphs (which might not be ZV-line graphs). Hence, it is straightforward to derive from our construction general mechanisms for recursive graph families.

Two technical assumptions we had are connectivity of the graph and finiteness of the number of agents and locations. Our results can be extended to the case of an infinite number of agents and locations under common natural constraints like finite diameter of the graph, measurability of $N(v)$ and of coalitions, and the order over Z-vertices being a well-order. It is also not hard to see that the following extension for graphs in which the connected components are ZV-line graphs will satisfy the same desiderata.

- ▶ At the first stage, choose the first connected component according to some predefined order s.t. at least one agent voted for a location in this component.
- ▶ At the second stage, run $\mathcal{F}^{\star}$ taking into account only agents who voted for locations in the chosen component.

The mechanism we presented is not the only mechanism satisfying the desired properties. Also taking any other order over the Z-vertices s.t. the constraints of Def. 3.3 hold and defining $\mathcal{F}^{\star}$ accordingly will satisfy them. For instance, a mechanism which takes at the second stage of Def. 3.2 the rightmost Pareto optimal Z-vertex also satisfies the desiderata. We did not find any mechanism satisfying the desiderata which is not of this template, and we conjecture that these are the only anonymous Pareto optimal manipulation-resistant mechanisms.

Conjecture 5.1. *Let $G = (\mathcal{V}, E)$ be a ZV-line graph w.r.t. $\mathcal{V} = Z \cup (V_1 \dot{\cup} \cdots \dot{\cup} V_k)$ and let $F \colon \bigcup_{t \geqslant 0} \mathcal{V}^t \to \mathcal{V}$ be a mechanism s.t.*

- *F is anonymous and Pareto optimal; and*
- *For any vector of locations $\boldsymbol{x} \in \mathcal{V}^n$, a coalition C, and a set of ballots $\boldsymbol{A} \in \bigcup_{t \geqslant 0} \mathcal{V}^t$, $\boldsymbol{A}$ is not a beneficial deviation for C.*

Then, for $i = 1, \ldots, k$: Whenever $\boldsymbol{x} \in (V_i)^n$, also $F(\boldsymbol{x}) \in V_i$. Moreover, F is the outcome of applying Def. 3.2 for some order over Z which satisfies the constraints of Def. 3.3 and mechanisms F_i which are defined by $\boldsymbol{x} \in (V_i)^n \mapsto F(\boldsymbol{x})$.

Furthermore, unifying non-existence results for specific graphs we've found so far, we think that the partition to Z-vertices and V-vertices is a fundamental property of a false-name-proof mechanism. Consequentially, showing that a given graph does not have such structure could be an easy and efficient way to prove non-existence of a desired mechanism.

Conjecture 5.2. *For almost any graph $G = \langle \mathcal{V}, E \rangle$, if there exists an anonymous and Pareto optimal mechanism $F \colon \bigcup_{t \geqslant 0} \mathcal{V}^t \to \mathcal{V}$ s.t. "For any vector of locations $\boldsymbol{x} \in \mathcal{V}^n$, a coalition C, and a set of ballots $\boldsymbol{A} \in \bigcup_{t \geqslant 0} \mathcal{V}^t$, $\boldsymbol{A}$ is not a beneficial deviation for C." then there exists a sequence of non-empty sets of vertices $Z, V_1, \ldots, V_k \subseteq \mathcal{V}$ s.t. G is a ZV-line graph w.r.t. $\mathcal{V} = Z \cup (V_1 \dot{\cup} \cdots \dot{\cup} V_k)$.*

The only counter example we've found to the conjecture is the cycle of size 5 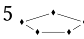, (and graphs derived from it, e.g., 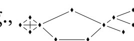). It is not hard to verify that • a mechanism which returns the first Pareto optimal location according to one of the following orders - 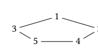

, 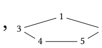

, 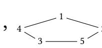

, or their rotations and reflections - is a manipulation-resistant mechanism and that • while all these mechanisms are of the template of Def. 3.2 (for all vertices being Z-vertices), these representations do not satisfy the connectivity constraints of Def. 3.3 and the cycle of size 5 is not a ZV-line graph. We conjecture that this is a representative extreme exception and intend to characterize the exception and replace 'almost' in Conjecture 5.2 with an exact statement.

Last, an important continuation of this work is analyzing the implications for *approximate mechanism design without money* [21]. That is, assuming the agents are accurately represented by a cost function (e.g., the distance to the facility or a monotone function of the distance) and analyzing implications of manipulation-resistance on the approximability of the minimization problem of natural social cost functions, e.g., the average cost (Harsanyi's social welfare), the geometric mean of the costs (Nash's social welfare), or the maximal cost (Rawls' criterion). For instance, assuming the two conjectures above, one gets that whenever there is a large disagreement in the population (i.e., the agents are dispersed over many V_i-subgraphs) an extreme status-quo alternative must be chosen by the mechanism, which results in a bad *price of false-name-proofness*. Nowadays, many aggregation mechanisms are highly susceptible to double voting and to false-name manipulations in general (e.g., mechanisms over huge anonymous networks like the internet, but also other scenarios in which vote frauds are known to be easy). We think that such results should open a discussion on the costs of these protocols (since the benefits are clear).

ACKNOWLEDGMENTS

We would like to thank Kentaro Yahiro, Nathanaël Barrot, and Tomer Siedner for their comments, and the anonymous reviewers for their detailed reviews, which helped us to improve the presentation of this work.

This work was supported by Israel Science Foundation Grant 1626/18, JSPS KAKENHI Grants JP17H00761 and JP17H04695, and JST Strategic International Collaborative Research Program (SICORP).

REFERENCES

[1] Fuad Aleskerov. 2002. Chapter 2 Categories of arrovian voting schemes. In *Handbook of Social Choice and Welfare*, Amartya K. Sen Kenneth J. Arrow and Kotaro Suzumura (Eds.). Handbook of Social Choice and Welfare, Vol. 1. Elsevier, 95 – 129.

[2] Noga Alon, Michal Feldman, Ariel D. Procaccia, and Moshe Tennenholtz. 2010. Strategyproof Approximation of the Minimax on Networks. *Mathematics of Operations Research* 35, 3 (2010), 513–526.

[3] Duncan Black. 1948. On the Rationale of Group Decision-making. *Journal of Political Economy* 56, 1 (1948), 23 – 34.

[4] E. H. Clarke. 1971. Multipart Pricing of Public Goods. *Public Choice* 1 (1971), 17–33.

[5] Vincent Conitzer. 2008. Anonymity-Proof Voting Rules. In *Proceedings of* $4^{\underline{th}}$ *Workshop on Internet and Network Economics (WINE)*. 295–306.

[6] Elad Dokow, Michal Feldman, Reshef Meir, and Ilan Nehama. 2012. Mechanism design on discrete lines and cycles. In *Proceedings of* $13^{\underline{th}}$ *ACM Conference on Economics and Computation (EC)*. 423–440.

[7] John R. Douceur. 2002. The Sybil Attack. In *Proceedings of* $1^{\underline{st}}$ *International Workshop Peer-to-Peer Systems IPTPS*. 251–260.

[8] Michal Feldman, Amos Fiat, and Iddan Golomb. 2016. On Voting and Facility Location. In *Proceedings of* $16^{\underline{th}}$ *ACM Conference on Economics and Computation (EC)*. 269–286.

[9] Peter C. Fishburn and Steven J. Brams. 1983. Paradoxes of Preferential Voting. *Mathematics Magazine* 56, 4 (1983), 207–214.

[10] Allan Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41 (1973), 587–601.

[11] Theodore Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–631.

[12] Frank Harary. 1963. A characterization of block-graphs. *Canadian Mathematical Bulletin (CMB)* 6 (1963), 1–6.

[13] Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan Allen Zhu. 2010. Asymptotically Optimal Strategy-Proof Mechanisms for Two-Facility Games. In *Proceedings of* $11^{\underline{th}}$ *ACM Conference on Economics and Computation (EC)*. 315–324.

[14] Pinyan Lu, Yajun Wang, and Yuan Zhou. 2009. Tighter bounds for facility games. In *Proceedings of* $5^{\underline{th}}$ *Workshop on Internet and Network Economics (WINE)*. 137–148.

[15] Andreu Mas-Colell, Michael Whinston, and Jerry Green. 1995. *Microeconomic Theory*. Oxford University Press.

[16] Hervé Moulin. 1980. On Strategy-Proofness and Single Peakedness. *Public Choice* 35, 4 (1980), 437–455.

[17] Hervé Moulin. 1988. Condorcet's principle implies the no show paradox. *Journal of Economic Theory* 45, 1 (1988), 53 – 64.

[18] Hervé Moulin. 1995. *Cooperative Microeconomics: A Game-Theoretic Introduction*. Princeton University Press.

[19] Roger B. Myerson. 1979. Incentive Compatibility and the Bargaining Problem. *Econometrica* 47 (1979), 61–74.

[20] Tomohiro Ono, Taiki Todo, and Makoto Yokoo. 2017. Rename and False-Name Manipulations in Discrete Facility Location with Optional Preferences. In *Proceedings of* $20^{\underline{th}}$ *Principles and Practice of Multi-Agent Systems (PRIMA)*. 163–179.

[21] Ariel D. Procaccia and Moshe Tennenholtz. 2013. Approximate Mechanism Design without Money. *ACM Transactions on Economics and Computation (TEAC)* 1, 4, Article 18 (2013).

[22] Kevin Roberts. 1979. The characterization of implementable choice rules. In *Aggregation and Revelation of Preferences*, Jean-Jacques Laffont (Ed.). North-Holland, 321 – 349.

[23] Mark A. Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10 (1975), 187–217.

[24] James Schummer and Rakesh V. Vohra. 2002. Strategy-proof Location on a Network. *Journal of Economic Theory* 104, 2 (2002), 405–428.

[25] Taiki Todo, Atsushi Iwasaki, and Makoto Yokoo. 2011. False-name-proof mechanism design without money. In *Proceedings of* $10^{\underline{th}}$ *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 651–658.

[26] Taiki Todo, Atsushi Iwasaki, Makoto Yokoo, and Yuko Sakurai. 2009. Characterizing false-name-proof allocation rules in combinatorial auctions. In *Proceedings of* $8^{\underline{th}}$ *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 265–272.

[27] William Vickrey. 1960. Utility, Strategy, and Social Decision Rules. *The Quarterly Journal of Economics* 74, 4 (1960), 507–535.

[28] Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. 2004. The effect of false-name bids in combinatorial auctions: new fraud in internet auctions. *Games and Economic Behavior* 46, 1 (2004), 174–188.