An Analysis Framework for Metric Voting based on LP Duality

David Kempe
University of Southern California

Abstract

Distortion-based analysis has established itself as a fruitful framework for comparing voting mechanisms. m voters and n candidates are jointly embedded in an (unknown) metric space, and the voters submit rankings of candidates by non-decreasing distance from themselves. Based on the submitted rankings, the social choice rule chooses a winning candidate; the quality of the winner is the sum of the (unknown) distances to the voters. The rule's choice will in general be suboptimal, and the worst-case ratio between the cost of its chosen candidate and the optimal candidate is called the rule's *distortion*. It was shown in prior work that every deterministic rule has distortion at least 3, while the Copeland rule and related rules guarantee distortion at most 5; a very recent result gave a rule with distortion $2+\sqrt{5} \approx 4.236$.

We provide a framework based on LP-duality and flow interpretations of the dual which provides a simpler and more unified way for proving upper bounds on the distortion of social choice rules. We illustrate the utility of this approach with three examples. First, we show that the Ranked Pairs and Schulze rules have distortion $\Theta(\sqrt{n})$. Second, we give a fairly simple proof of a strong generalization of the upper bound of 5 on the distortion of Copeland, to social choice rules with short paths from the winning candidate to the optimal candidate in generalized weak preference graphs. A special case of this result recovers the recent $2+\sqrt{5}$ guarantee. Finally, our framework naturally suggests a combinatorial rule that is a strong candidate for achieving distortion 3, which had also been proposed in recent work. We prove that the distortion bound of 3 would follow from any of three combinatorial conjectures we formulate.

1 Introduction

Voting is an important way for a group to choose one out of multiple available candidate options. The group could be a country, academic department, or other organization, and the n candidate options they choose from could be courses of action or human candidates. Typically, each voter submits a total order of all options, called a *ranking* or *preference order*. Based on all the submitted rankings, a *social choice rule* (or *mechanism*) determines the winning option.

Complementing the "traditional" axiomatic approach for comparing social choice rules (see, e.g., (Brandt et al. 2016) for a fairly recent overview), an alternative approach is to view social choice through the lens of *optimization and approximation*. In this line of work (e.g., (Boutilier et al. 2015; Caragiannis and Procaccia 2011; Procaccia 2010; Procaccia and Rosenschein 2006)), each voter has a cost/utility for each candidate, and the social welfare of a candidate is the sum/average of the costs/utilities of all voters.

While the social choice rule might aim to optimize social welfare, it does not have access to the actual costs/utilities; rather, it only receives the voters' rankings, which convey partial information. In other words, even though the voting mechanism must optimize a *cardinal* objective function, it only receives *ordinal* information as input, namely, for each voter, whether her[1] utility/cost for candidate x is larger or smaller than that for candidate y. As a result, the mechanism will typically select a suboptimal candidate. The *distortion* of a mechanism is the worst-case ratio between the welfare/cost of the mechanism's selected (based only on ordinal information) candidate and the optimum (with full knowledge of the cardinal values) candidate, over all possible inputs. (Formal definitions of this concept and all other terms can be found in Section 2.)

A particularly natural way of defining costs is via a joint *metric space* on candidates and voters. The distance $d(v, x)$ between voter v and candidate x captures their difference in opinion and hence the cost; therefore, voters rank candidates by non-decreasing distance from themselves. This modeling approach was proposed in (Anshelevich, Bhardwaj, and Postl 2015); see also (Anshelevich et al. 2018) for an expanded/improved journal version, and (Anshelevich 2016) for a broader overview of the area and its results.

The main result of (Anshelevich, Bhardwaj, and Postl 2015; Anshelevich et al. 2018) is that under the model of metric costs, many widely used voting rules (including Plurality, Veto, Borda count, and others) have distortion linear in the number of candidates or worse. Furthermore, even with just 2 candidates and a 1-dimensional metric space, *every* deterministic voting mechanism has distortion at least 3.

[1]For ease of presentation, we use female pronouns for voters and male pronouns for candidates throughout.

On the positive side, any rule which always outputs a candidate from the uncovered set of candidates has distortion at most 5, for all metric spaces and numbers of candidates. Uncovered sets are defined in terms of the tournament graph G on n candidates in which the directed edge (x, y) is present iff a (weak) majority of voters prefer x to y. The uncovered set is the set of candidates that have a directed path of length at most 2 in G to every other candidate (see (Moulin 1986)). Very recently, (Munagala and Wang 2019) gave a voting rule based on uncovered sets in a weighted tournament graph which improves the upper bound from 5 to $2+\sqrt{5} \approx 4.236$.

There is an obvious gap between the lower bound of 3 for the distortion of every mechanism, and the upper bound of $2 + \sqrt{5}$. In the original version of (Anshelevich, Bhardwaj, and Postl 2015), it was conjectured that a mechanism called *Ranked Pairs* (defined in Section 2) achieves a distortion of 3. This conjecture was disproved by (Goel, Krishnaswamy, and Munagala 2017), who showed a lower bound of 5 on the distortion of Ranked Pairs (and the Schulze rule, also defined in Section 2).

The proof of the upper bound of 5, the recent upper bound of $2 + \sqrt{5}$, and many other proofs in the literature are based on reasoning about all metric spaces that are consistent with assumed rankings. They often involve intricate case distinctions and rather ad hoc arguments. So far, a more solid foundation and framework for distortion proofs has been missing from the literature.

1.1 Our Contribution

Our main contribution, presented in Section 3, is an analysis framework based on LP duality and flows for proving upper bounds on the metric distortion of voting mechanisms. Our point of departure is a linear program for the following problem: given the rankings of all voters, a winning candidate (presumably selected by a mechanism) and an "optimum" candidate, find a metric space maximizing the distortion of this choice.[2] We show that the dual of the cost maximization LP can be interpreted as a flow problem with an unusual objective function. Using this framework, in order to show an upper bound on the metric distortion of a particular mechanism, rather than having to explicitly consider all possible metric spaces, it is enough to exhibit a flow of small cost meeting certain demands. We illustrate the power of this analysis framework with three applications.

First, in Section 4, we resolve the distortion of the Ranked Pairs and Schulze rules (defined in Section 2): we show that both have distortion $\Theta(\sqrt{n})$. The upper bound is a clean application of the duality framework, while the lower bound is obtained with a generalization of the example which (Goel, Krishnaswamy, and Munagala 2017) used to lower-bound the distortion of both rules by 5. The distortion of both rules is thus significantly higher than the distortions of 5 and $2 + \sqrt{5}$ achieved by the uncovered set mechanisms. Understanding the distortion of the Schulze rule in particular is of importance because it is widely used in practice.

[2] This approach can of course immediately be leveraged into an optimal polynomial-time voting mechanism; we discuss this more in Section 3.

Then, in Section 5, we give a strong generalization of the key lemmas from (Anshelevich, Bhardwaj, and Postl 2015) (Theorem 7) and (Munagala and Wang 2019) (Lemma 3.7), used to prove distortions of 5 and $2 + \sqrt{5}$ for the respective mechanisms. The common idea of both is that when a large enough fraction of voters prefer x to y, and a large enough fraction prefer y to z, then the cost of x can be bounded in terms of the cost of z. These bounds immediately imply the upper bounds on the distortion for any candidate in the uncovered set of a suitably defined tournament graph. We give a generalization to arbitrary chains of preferences, and upper-bound the cost of x_1 in terms of the cost of x_ℓ when a p_i fraction of voters prefers x_i over x_{i+1}, for each $i = 1, \ldots, \ell - 1$. For the specific case when all $p_i \geq p$, the bound can be stated very cleanly: the cost of x_1 is at most $\frac{\ell}{p} - 1$ times that of x_ℓ if ℓ is even, and at most $\frac{\ell-1}{p} + 1$ times that of x_ℓ if ℓ is odd. Our results fully recover and generalize the bounds of (Anshelevich, Bhardwaj, and Postl 2015) and (Munagala and Wang 2019). The generalization to longer path lengths can be useful in analyzing voting mechanisms that are missing information. This can happen if the environment restricts the communication between voters and the mechanism, so that parts of the ranking remain unknown, as in (Kempe 2020). In fact, the results of Section 5 can be used to significantly improve the upper bounds on the performance of "Copeland-like" mechanisms with missing information, compared to the bounds in (Kempe 2020).

As a third application, the flow interpretation naturally suggests a candidate mechanism that might achieve distortion 3, which we present in Section 6. The analysis points to a sufficient condition for distortion 3: that for every given preference profile of the voters, there be a candidate x such that for all other candidates y, a certain bipartite graph on the voters have a perfect matching. In fact, the mechanism itself can be phrased in this terminology, leading to a purely combinatorial polynomial-time mechanism. This mechanism was independently discovered and presented in (Munagala and Wang 2019). There, it is also shown — again with a case distinction proof over metric spaces — that if such a candidate x exists, the mechanism guarantees distortion 3. Our duality framework gives a cleaner and simpler proof of this fact. The main question is then whether the desired candidate x always exists.

(Munagala and Wang 2019) conjecture — as do we — that it does. In Section 6, we present several very different-looking conjectures, each of which would resolve the question positively, i.e., establish a distortion of 3. One of the conjectures is phrased in terms of certain preferences between candidates and sets under randomly drawn preference orders, while another talks about cycles in certain induced subsets of a type of graph we define. The fact that they are sufficient to establish distortion 3 is based on Hall's Marriage Theorem for bipartite graphs. We have verified the conjecture for $n \leq 7$ using exhaustive computer search.

Due to space constraints, many proofs, as well as a more in-depth discussion of related work, are omitted from this conference version. A full version is available on the arXiv (Kempe 2019).

2 Preliminaries

2.1 Voters, Candidates, and Social Choice Rules

An *instance* $(X, \mathcal{P})$ consists of a set of n candidates X, and the voters' preferences $\mathcal{P}$ among these candidates. Candidates will always be denoted by lowercase letters w, x, y, z, with w specifically reserved for a candidate chosen as winner by a mechanism. Sets of candidates are denoted by uppercase letters X, Y, Z. The m voters are denoted by v, v' and variants thereof, and the set of all voters is V.

Each voter v has a *total order* (or *preference order* or *ranking*) $\succ_v$ over the n candidates. $x \succ_v y$ denotes that v (strictly) prefers x over y, and $x \succeq_v y$ denotes that v weakly prefers x over y (allowing $x = y$). We extend this notation to sets, writing, for instance, $Y \succ_v Z$ to denote that v (strictly) prefers all candidates in Y over all candidates in Z. We write $[x \succ Y] = \{v \in V \mid x \succ_v Y\}$ for the set of voters who rank x strictly ahead of all candidates in Y, and $[Y \succ x] = \{v \in V \mid Y \succ_v x\}$ for the set of voters who rank x strictly behind all candidates in Y.

A *vote profile* $\mathcal{P}$ is the vector of the rankings of all voters $\mathcal{P} = (\succ_v)_{v \in V}$. A *social choice rule* (or *mechanism*) $f : (X, \mathcal{P}) \mapsto w$ is given the rankings of all voters, i.e., $\mathcal{P}$, and deterministically produces as output one *winning* candidate $w = f(X, \mathcal{P}) \in X$.

2.2 (Pseudo-)Metric Space and Distortion

The voter preferences are assumed to be derived from distances between voters and candidates. The distance $d(v, x)$ between voter v and candidate x captures how similar their positions on key issues are. The distances d form a *pseudo-metric*, i.e., they are non-negative and satisfy the triangle inequality $d(v, x) \leq d(v, y) + d(v', y) + d(v', x)$ for all voters v, v' and candidates x, y.

A vote profile $\mathcal{P}$ is *consistent* with the pseudo-metric d if and only if each voter ranks the candidates by non-decreasing distance from herself; that is, if $x \succ_v y$ whenever $d(v, x) < d(v, y)$. When $\mathcal{P}$ is consistent with d, we write $d \sim \mathcal{P}$. If there are ties among distances, several vote profiles will be consistent with d.

Definition 2.1 (Social Cost, Distortion) *1. The* social cost *of candidate x is the sum of distances from x to all voters: $C(x) = \sum_v d(v, x)$.*

2. A candidate is an optimum candidate *iff he minimizes the social cost: $x^*_d \in \operatorname{argmin}_{x \in X} C(x)$.*

3. The distortion *of a mechanism f is the largest possible ratio between the cost of the candidate chosen by f, and the optimal (with respect to the pseudo-metric d, which f does not know) candidate x^*_d:*

$$\rho(f) \;=\; \max_{\mathcal{P}} \sup_{d : d \sim \mathcal{P}} \frac{C(f(X, \mathcal{P}))}{C(x^*_d)}.$$

2.3 Ranked Pairs and the Schulze Rule

Both the Ranked Pairs and Schulze Rules are based on a weighted directed graph on the set of candidates X. The weight $p_{x,y}$ of the edge from candidate x to y is the fraction of voters who have $x \succ y$.

In Ranked Pairs (Tideman 1987), the (ordered) pairs (x, y) are considered in non-increasing order of $p_{x,y}$. When the pair (x, y) is considered, the directed edge (x, y) is inserted into the graph if and only if doing so creates no cycle. When the insertion process terminates, the graph has a unique source node, which is returned as the winner.

In the Schulze Method (Schulze 2011), a directed weighted graph is created in which each ordered pair (x, y) has an edge with weight $p_{x,y}$. Then, for each pair (x, y), let $s_{x,y}$ be the width of the widest path from x to y, that is, the largest p such that there is a path from x to y on which all edges (x', y') have $p_{x',y'} \geq p$. It has been shown (Schulze 2011) that there is a candidate node x such that $s_{x,y} \geq s_{y,x}$ for all other candidates y. Any such candidate x is returned as the winner.

For the purposes of our analysis, the only property of these methods that matters is captured by the following lemma, which is well known (and proved in the full version for completeness).

Lemma 2.2 *Let w be the candidate selected by the rule (either Ranked Pairs or Schulze), and y any other candidate. Then, there exists a p and a sequence of (distinct) candidates $x_1 = w, x_2, \ldots, x_\ell = y$ with the property that at least a p fraction of voters prefer x_i over x_{i+1} (for each i), and at most a p fraction of voters prefer y over w.*

3 The LP Duality Approach and Flows

In this section, we develop the key tool for our analysis: the dual linear program for distortion in metric voting.

The voters' preferences $\mathcal{P} = (\succ_v)_v$ are given. Let w be a candidate that the mechanism is considering as a potential winner, and x^* the optimal candidate. Following (Anshelevich et al. 2018; Goel, Krishnaswamy, and Munagala 2017), we phrase the problem of finding a distortion-maximizing metric as a linear program whose variables $d_{v,x}$ denote distances between voters v and candidates x. These distances must be non-negative, obey the triangle inequality, and be consistent with the reported preferences of the voters. The objective is to maximize the distortion, i.e., the ratio between the cost of w and the cost of x^*.

$$\begin{array}{ll} \text{Maximize } \sum_v d_{v,w} \text{ subject to} & \\ \quad d_{v,x} \leq d_{v',x} + d_{v',y} + d_{v,y} \quad \text{for all } x, y, v, v' & \\ \qquad\qquad (\triangle \text{ Inequality}) & \\ \quad d_{v,x} \leq d_{v,y} \quad \text{for all } x, y, v \text{ such that } x \succ_v y & \\ \qquad\qquad \text{(consistency)} & (1) \\ \quad \sum_v d_{v,x^*} = 1 \quad \text{(normalization)} & \\ \quad \sum_v d_{v,x} \geq 1 \quad \text{for all } x \text{ (optimality of } x^*) & \\ \quad d_{v,x} \geq 0 \quad \text{for all } x, v. & \end{array}$$

As is standard in the use of LPs for optimizing a ratio, the normalization side-steps the issue of having to write a ratio: for any worst-case metric, one could simply rescale all distances by a constant so that the normalization holds — this does not change any ratios, and thus also not the distortion.

As discussed in (Anshelevich et al. 2018; Goel, Krishnaswamy, and Munagala 2017), the LP can be immediately used to define an instance-optimal mechanism. However, analyzing such a mechanism appears difficult. We will show how to use the dual of the LP to obtain a general approach for bounding the metric distortion of several different voting rules.

3.1 The Dual Linear Program

Rearranging the primal LP into normal form, taking the dual, and switching the signs of the α_x variables yields the following dual LP (2). In this LP, the $\psi_{x,y}^{(v,v')}$ are the dual variables for the triangle inequality constraints, $\phi_{x,y}^{(v)}$ are the dual variables for the consistency constraints, and the α_x are the dual variables for the normalization/optimality constraints.

$$\begin{aligned}
&\text{Minimize} \quad \textstyle\sum_x \alpha_x \quad \text{subject to} \\
&\quad \alpha_x + \textstyle\sum_{y:x \succ_v y} \phi_{x,y}^{(v)} - \sum_{y:y \succ_v x} \phi_{y,x}^{(v)} \\
&\quad + \textstyle\sum_{y,v'} \left(\psi_{x,y}^{(v,v')} - \psi_{y,x}^{(v,v')} - \psi_{x,y}^{(v',v)} - \psi_{y,x}^{(v',v)} \right) \\
&\quad \geq \begin{cases} 1 \text{ if } x = w \\ 0 \text{ if } x \neq w \end{cases} \qquad \text{for all } v, x \qquad\qquad (2) \\
&\quad \psi_{x,y}^{(v,v')} \geq 0 \quad \text{for all } v, v', x, y \\
&\quad \phi_{x,y}^{(v)} \geq 0 \quad \text{for all } v, x, y \\
&\quad \alpha_x \leq 0 \quad \text{for all } x \neq x^*.
\end{aligned}$$

Notice that α_{x^*} is in fact unconstrained, due to the equality constraint in the normalization.

The advantage of studying the dual linear program instead of the primal (or reasoning about the distortion directly) is that it omits any reference to any metric space. The goal in analyzing a mechanism is to show that for any voter preferences $\mathcal{P}$, with a suitably chosen winner w, there is a setting of the dual variables giving small objective value.

3.2 Using the Dual by Exhibiting Flows

The LP (2) looks rather unwieldy, mostly due to the terms involving the $\psi_{x,y}^{(v,v')}$ variables. However, by making some specific choices for these variables, it can be interpreted as a flow problem on a suitably defined graph, with a somewhat unusual objective function. This is captured by the following lemma:

Lemma 3.1 *Let $H = (U, E)$ be a directed graph with vertex set $U = V \times X$, and edges defined as follows:*

- *Whenever $x \succ_v y$, E contains the directed edge $(v, x) \to (v, y)$. We call such edges* preference edges*.*
- *For all x and $v \neq v'$, E contains the directed edge $(v, x) \to (v', x)$. We call such edges* sideways edges*.*

Let f be a flow on H such that exactly one unit of flow originates at the node (v, w) for each voter v, and flow is only absorbed at nodes (v, x^*) for voters v. Define the cost of f at voter v to be $\gamma_v^{(f)} = \sum_{e \text{ into } (v,x^*)} f_e + \sum_{x \neq x^*} \sum_{v' \neq v} (f_{(v',x)\to(v,x)} + f_{(v,x)\to(v',x)})$.

Then, $C(w) \leq C(x^*) \cdot \max_v \gamma_v^{(f)}$.

The graph H has two types of edges. For any fixed voter v, the preference edges $(v, x) \to (v, y)$ (over all candidate pairs x, y) exactly correspond to v's preference order. For any fixed candidate x, the sideways edges $(v, x) \to (v', x)$ (over all voter pairs v, v') form a complete directed graph.

The flow's cost function has two terms for each voter v. The first is fairly standard in the study of multi-commodity flows: the capacity required at the sink node (v, x^*) to be able to absorb all of the flow. The second one is rather nonstandard: for each voter v, there is an additional penalty term for all incoming and outgoing flows of nodes (v, x) for $x \neq x^*$ along sideways edges. In other words, using preference edges is much less costly than using sideways edges: the former just route flow, while the latter route the flow, but also incur a cost penalty at both endpoints.

4 Distortion of Ranked Pairs and Schulze

As a first application of Lemma 3.1, we pin down the distortion of the Ranked Pairs and Schulze rules to within constant factors.

Corollary 4.1 *Both the Ranked Pairs mechanism and the Schulze rule asymptotically have distortion at most $5\sqrt{n} + o(\sqrt{n})$ and at least $\frac{\sqrt{2}}{2}\sqrt{n}$.*

Proof. The lower bound is based on a straightforward generalization of the construction of (Goel, Krishnaswamy, and Munagala 2017), which showed a lower bound of 5 on the distortion. It is given in the full version. Here, we prove the upper bounds. Let w be the candidate selected by the rule, and x^* the optimum candidate. By Lemma 2.2, applied with $y = x^*$, there exists a p and a sequence of distinct candidates $x_1 = w, x_2, \ldots, x_\ell = x^*$ with the property that for each i, at least a p fraction of voters prefer x_i over x_{i+1}, and at most a p fraction of voters prefer x^* over w. The existence of $x_1, \ldots, x_\ell$ with these properties is all that we need from the specific voting rules. The rest of the proof will be completely generic, and would thus also apply to any other voting rule satisfying Lemma 2.2.

We consider two cases, depending on the value of p. The case $p \leq 1 - \frac{2}{5\sqrt{n}}$ is easy. In this case, at least a $1 - p \geq \frac{2}{5\sqrt{n}}$ fraction of voters prefer w over x^*. Lemma 6 from (Anshelevich et al. 2018) states that if at least a q fraction of voters prefer x over x', then $C(x) \leq (1 + \frac{2(1-q)}{q}) \cdot C(x')$. Applying this lemma to w and x^*, the distortion of w is at most $\frac{2}{1-p} - 1 \leq 5\sqrt{n}$.

When $p > 1 - \frac{2}{5\sqrt{n}}$, we use Lemma 3.1, and define a flow. Let $\lambda = \lfloor \sqrt{n}/2 \rfloor$.[3] Let $B = \lceil \ell/\lambda \rceil$. Because $\ell \leq n$, we get that $B + 1 \leq 2\sqrt{n} + o(\sqrt{n})$. Consider the $B + 1$ candidates $y_j := x_{j\lambda+1}$ for $j = 0, 1, \ldots, B - 1$, and $y_B := x_\ell$. For each $j < B$, let A_j be the set of voters who prefer candidate y_j to y_{j+1}. Because for each i, at least a $1 - \frac{2}{5\sqrt{n}}$ fraction of voters prefer x_i to x_{i+1}, a union bound over the candidates

[3] To ensure that $\lambda \geq 1$, we may assume that $n \geq 4$. For smaller n, it is easy to see that the distortion of both rules is at most a constant, which of course can be absorbed in the $o(\sqrt{n})$ term.

$x_{j\lambda+1}, x_{j\lambda+2}, \ldots, x_{(j+1)\lambda}$ shows that for each $j < B$, at least a $1 - \lambda \cdot \frac{2}{5\sqrt{n}} \geq \frac{4}{5}$ fraction of voters prefer y_j over y_{j+1}; that is, $|A_j| \geq \frac{4}{5} \cdot m$.

We are now ready to construct the flow, which we will do by increasing j. Initially, each node $(v, x_1) = (v, y_0)$ has one unit of incoming flow. Each node (v, y_0) with $v \notin A_0$ distributes its unit of flow evenly over the nodes (v', y_0) with $v' \in A_0$. Then, for each $j = 0, \ldots, B-1$ and each voter $v \in A_j$, the node (v, y_j) routes all its flow to the node (v, y_{j+1}) along the preference edge $(v, y_j) \to (v, y_{j+1})$; this preference edge $(v, y_j) \to (v, y_{j+1})$ exists because $v \in A_j$. Subsequently, the flow into (v, y_{j+1}) gets distributed uniformly to nodes (v', y_{j+1}) with $v' \in A_{j+1}$ along sideways edges. This concludes the description of the flow.

We now analyze the flow's cost, according to the cost metric of Lemma 3.1. Because $|A_j| \geq \frac{4}{5} \cdot m$ for all j, no node ever has more than $\frac{5}{4}$ units of flow. Now focus on any voter v. Cost for v is incurred by incoming flow into nodes (v, y_j) along sideways edges, outgoing flow from nodes (v, y_j) along sideways edges, and flow into (v, y_B). Each of these cost terms is bounded by $\frac{5}{4}$ by the preceding observation. Thus, the total cost for node v associated with one particular y_j (for $j < B$) is at most $\frac{5}{2}$, while the cost associated with y_B is at most $\frac{5}{4}$. Adding these terms for all $j = 0, \ldots, B$ gives an upper bound of $\frac{5}{4} + 5\sqrt{n} + o(\sqrt{n})$. Since this holds for each v, Lemma 3.1 implies that the distortion is at most $5\sqrt{n} + o(\sqrt{n})$. ■

Remark 4.2 The upper bound in Corollary 4.1 was a direct application of our flow-based framework. While the lower bound did not explicitly use the framework, the counterexample was in fact discovered after failed attempts to improve the upper bound. The failure to find ways to route flow very clearly suggested the types of rankings that were obstacles (i.e., reversed block structures).

5 Generalization of Distortion Bounds for Undominated Nodes

As a second corollary of Lemma 3.1, we obtain a strong generalization of Theorem 7 in (Anshelevich, Bhardwaj, and Postl 2015) and Lemma 3.7 of (Munagala and Wang 2019) (which are given below for comparison). The most general version can be stated as follows:

Corollary 5.1 *Let $x_1, x_2, \ldots, x_\ell$ be (distinct) candidates such that for each $i = 2, \ldots, \ell$, at least a $p_i > 0$ fraction of voters prefer candidate x_{i-1} over candidate x_i. Define $\lambda_1 = 1, \lambda_2 = \frac{2}{p_2} - 1$, and $\lambda_i = \frac{2}{p_i}$ for $2 < i \leq \ell$. Let $\Lambda = \max_{S \subseteq \{1,\ldots,\ell\}, S \text{ indep.}} \sum_{i \in S} \lambda_i$. (Here, independence of a set S of natural numbers means that the set contains no two consecutive numbers.) Then, $C(x_1) \leq \Lambda \cdot C(x_\ell)$.*

Proof. We define a flow f and analyze its cost. For each i, we call the nodes (v, x_i) (for all voters v) *layer* i. Let V_i be the set of voters v with $x_{i-1} \succ_v x_i$, with $V_1 := V$ for notational simplicity.

We construct the flow layer by layer; our construction will ensure that each node (v, x_i) with $v \in V_i$ has exactly $\frac{m}{|V_i|}$ units of flow entering. This holds in the base case $i = 1$, because each node in layer 1 is the source node of one unit of flow.

For the i^{th} step of the construction, we first route all the flow within layer i using sideways edges, from nodes (v, x_i) with $v \in V_i$ to nodes (v', x_i) with $v' \in V_{i+1}$. We then route it to nodes (v', x_{i+1}) in layer $i+1$ using preference edges. More specifically, to route the flow within layer i, we first consider voters $v \in V_i \cap V_{i+1}$. For those voters, $\min(\frac{m}{|V_i|}, \frac{m}{|V_{i+1}|})$ units of flow simply stay at (v, x_i). The node (v, x_i) for such v will have additional incoming flow from other nodes (if $\frac{m}{|V_{i+1}|} > \frac{m}{|V_i|}$) or additional outgoing flow to other nodes (if $\frac{m}{|V_{i+1}|} < \frac{m}{|V_i|}$). The remaining flow is routed arbitrarily using sideways edges from nodes (v, x_i) with $v \in V_i$ to nodes (v', x_i) with $v' \in V_{i+1}$, of course ensuring that each such node (v', x_i) has in total $\frac{m}{|V_{i+1}|}$ units of flow entering.

After this redistribution within layer i, each (v, x_i) routes its flow to (v, x_{i+1}). Notice that this is always possible, because $x_i \succ_v x_{i+1}$ for all $v \in V_{i+1}$. The construction is illustrated with an example in Figure 1.

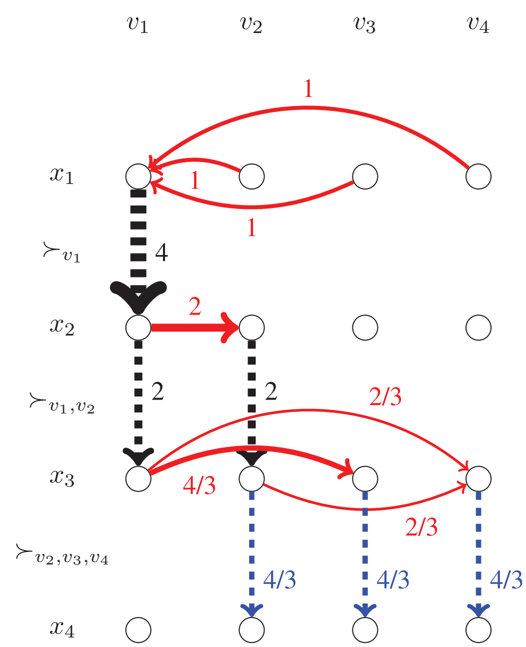

Figure 1: An illustration of the flow construction. In the example, there are 4 voters and 4 relevant candidates, with voter preferences shown on the left. The preference fractions are $p_1 = 1/4, p_2 = 1/2, p_3 = 3/4$. Sideways flows are shown in solid red, while flow along preference edges is shown in dashed lines. The dashed lines into nodes for candidate x_4 are shown in blue (instead of black), to emphasize that they contribute to the objective function. The amount of flow is given numerically, and also shown using the width of the lines/arcs. Edges that are not used by the flow are not shown.

We now analyze the cost associated with any fixed voter v. The cost has two components: the incoming flow at (v, x_ℓ) (shown in blue in Figure 1), and the cost associated with incoming/outgoing flow using sideways edges incident on (v, x_i) for $i < \ell$ (shown in red in Figure 1). We begin with the incoming flow at (v, x_ℓ): if $v \in V_{\ell-1}$, the incoming flow is $\frac{m}{|V_\ell|} = \frac{1}{p_\ell}$; otherwise, it is 0.

Next, we consider the cost associated with sideways edges. As a general guideline (subtleties will be discussed momentarily), when $v \in V_i$, the node (v, x_i) has $\frac{m}{|V_i|} = \frac{1}{p_i}$ units of flow coming in along sideways edges, and the node (v, x_{i+1}) has the same amount of flow leaving along sideways edges. The associated cost of both together is $\frac{2}{p_i}$. Two obvious exceptions are layers $i = 1$ and $i = \ell - 1$. For $i = 1$, one unit of flow simply originates with (v, x_1), resulting in no cost. For $i = \ell - 1$, no sideways edge is used to route outgoing flow; however, this is compensated by the incoming flow at (v, x_ℓ) (discussed in the preceding paragraph), which adds the same cost term.

However, simply adding up the bounds from the preceding paragraph over all steps i with $v \in V_i$ is too pessimistic, because our flow construction avoids routing more flow than necessary when $v \in V_i \cap V_{i+1}$. A tighter bound is captured by the following lemma:

Lemma 5.2 *Let I be the set of all indices i with $v \in V_i$. I can be partitioned into disjoint intervals of integers $\{L_1, L_1 + 1, \ldots, R_1\}, \{L_2, L_2 + 1, \ldots, R_2\}, \ldots, \{L_K, L_K + 1, \ldots, R_K\}$ (for some $K \geq 1$) such that:*

1. *For each k, there exists an index $M_k \in \{L_k, \ldots, R_k\}$ such that $p_{L_k} \geq p_{L_k+1} \geq \cdots \geq p_{M_k} > 0$ and $p_{M_k} \leq p_{M_k+1} \leq \cdots \leq p_{R_k}$; that is, the p_i are monotone non-increasing from L_k to M_k, and monotone non-decreasing from M_k to R_k.*
2. *The total cost of flow (both sideways flow and flow into (v, x_ℓ) in case $R_k = \ell$) associated with nodes (v, x_i) with $L_k \leq i \leq R_k$ is at most λ_{M_k}.*

To apply Lemma 5.2, the key observation is that the set $\{M_1, M_2, \ldots, M_K\}$ is independent, i.e., contains no two consecutive integers. If it did — say, $i = M_k$ and $i + 1 = M_{k'}$ — then both $i, i+1 \in I$. If $p_{i+1} \leq p_i$, this would contradict the maximality of i in the definition of M_k; on the other hand, if $p_{i+1} > p_i$, then $i + 1 \leq R_k$ by the definition of R_k, so it is impossible that $i + 1 = M_{k'}$.

Now, summing up the costs for each of the disjoint intervals, we obtain that the total cost of the flow at nodes associated with v (both sideways flow and flow into (v, x_ℓ)) is at most $\sum_{k=1}^{K} \lambda_{M_k}$; because the set of M_k is independent, this sum is at most Λ. Using Lemma 3.1, this completes the proof. ■

5.1 Special Cases

Lemma 3.7 of (Munagala and Wang 2019) is the special case of Corollary 5.1 with $\ell = 3, x_1 = w, x_3 = x^*$, and $p_1 = \frac{3-\sqrt{5}}{2}, p_2 = \frac{\sqrt{5}-1}{2}$. Our Corollary 5.1 then exactly recovers the bound of $2 + \sqrt{5}$.

When we have a uniform lower bound on the p_i, Corollary 5.1 can be simplified significantly.

Corollary 5.3 *Let $x_1, x_2, \ldots, x_\ell$ be (distinct) candidates such that for each $i = 2, \ldots, \ell$, at least a $p > 0$ fraction of the voters prefer candidate x_{i-1} over candidate x_i. Then, if ℓ is even, $C(x_1) \leq (\frac{\ell}{p} - 1) \cdot C(x_\ell)$; if ℓ is odd, $C(x_1) \leq (\frac{\ell-1}{p} + 1) \cdot C(x_\ell)$.*

Proof. We substitute $p_i \geq p$ for all i in Corollary 5.1; then, we observe that for even ℓ, the independent set of integers giving the largest sum is $\{2, 4, \ldots, \ell\}$, while for odd ℓ, it is $\{1, 3, \ldots, \ell\}$. ■

When $p = \frac{1}{2}$ (i.e., in the case of the majority graph), the bound simply becomes $2\ell - 1$. The result thus strongly generalizes Theorem 7 in (Anshelevich, Bhardwaj, and Postl 2015), which is the special case of $p = \frac{1}{2}$ and $\ell = 3$.

6 A Candidate Algorithm for Distortion 3

As a third application, we derive a purely combinatorial (i.e., not LP-based) voting mechanism, which we conjecture to have distortion 3. We show that this conjecture would follow from any of three different-looking combinatorial conjectures we will formulate.

The point of departure for the derivation of the mechanism is Corollary 6.1, which simplifies Lemma 3.1, reducing it to a purely combinatorial property of a certain graph. Corollary 6.1 was proved as Theorem 4.4 in (Munagala and Wang 2019), using a significantly more complex proof.

For any two candidates x, y, we consider the following bipartite graph $H_{x,y}$ on the node set (V, V); that is, there is one node on the "left" for each voter v, and one node on the "right" for each voter v'. (We will use "left" and "right" to distinguish the two vertex sets.) There is an edge (v, v') if and only if there exists a candidate $z \in X$ ($z = x$ or $z = y$ are explicitly allowed) such that $x \succeq_v z$ and $z \succeq_{v'} y$.

Corollary 6.1 *Let $x \neq y$ be two candidates. If $H_{x,y}$ has a perfect matching, then $C(x) \leq 3C(y)$.*

Proof. Assume that there is a perfect matching in $H_{x,y}$; for each voter v, let μ_v be the voter v is matched with. By definition of $H_{x,y}$, there is a candidate z_v such that $x \succeq_v z_v$ and $z_v \succeq_{\mu_v} y$.

We now define the flow f from each source node (v, x). We route one unit of flow along the path $(v, x) \to (v, z_v) \to (\mu_v, z_v) \to (\mu_v, y)$. Notice that by definition of z_v, the first and third edge always exist. Also, if $z_v = x$, then the first two nodes are the same, and we omit the first edge. Similarly, if $\mu_v = v$, we omit the second edge, and if $z_v = y$, we omit the third edge.

This construction defines a valid flow, routing one unit of flow from each (v, x) to some (v', y) (for some v'). So it only remains to bound $\gamma_v^{(f)} \leq 3$ for all v.

Because μ is a matching, there is exactly one unit of flow arriving at each node (v, y). For a given voter v, let v' be the unique voter with $\mu_{v'} = v$. Then, the only two edges of the form $(v, z) \to (v', z)$ or $(v', z) \to (v, z)$ that can be used by f are $(v, z_v) \to (\mu_v, z_v)$ and $(v', z_{v'}) \to (v, z_{v'})$. Hence, the second part of the cost term $\gamma_v^{(f)}$ is at most 2, meaning that $\gamma_v^{(f)} \leq 3$. The claim now follows by applying Lemma 3.1. ■

Corollary 6.1 immediately suggests a natural mechanism with distortion at most three, which was also given as MATCHINGUNCOVERED in (Munagala and Wang 2019):

ALL BIPARTITE MATCHINGS:
Find a candidate w such that for all other candidates x, the bipartite graph $H_{w,x}$ has a perfect matching.

The mechanism ALL BIPARTITE MATCHINGS sidesteps having to solve the $\Theta(n^2)$ LPs (1), instead solving $\Theta(n^2)$ bipartite matching problems. The question then is whether such a candidate w actually exists. We present three different conjectures, each of which would imply the existence of w, and thus, by Corollary 6.1, that ALL BIPARTITE MATCHINGS has distortion 3.

6.1 The Candidate Comparison Graph $\mathcal{G}$

A key analysis tool in this section is a directed graph $\mathcal{G}$ on the set of all candidates X, which we call the *Candidate Comparison Graph*. $\mathcal{G}$ contains the directed edge (y, x) if and only if the graph $H_{x,y}$ does *not* have a perfect matching. Any candidate w without incoming edges in $\mathcal{G}$ is a safe choice as a winner, because Corollary 6.1 implies a bound of 3 on its cost ratio to (the unknown) x^*.

One sufficient condition for the existence of a source node in $\mathcal{G}$ (i.e., a node without incoming edges) is for $\mathcal{G}$ to be acyclic. This gives rise to our first conjecture, which was also given as Conjecture 4.8 in (Munagala and Wang 2019):

Conjecture 1 *For every instance $(X, \mathcal{P})$, the graph $\mathcal{G} = \mathcal{G}(X, \mathcal{P})$ is non-Hamiltonian.*[4]

Despite appearances, Conjecture 1 is equivalent to the guaranteed existence of a source node, as we show in the following proposition:

Proposition 6.2 ALL BIPARTITE MATCHINGS *succeeds* on all inputs *if and only if Conjecture 1 is true.*

Notice that the proposition does not say that whenever a *specific* instance violates Conjecture 1, the algorithm will fail on *that* instance. It only implies that the algorithm fails on *some* (potentially different) instance.

6.2 Distributions of Permutations

We next derive a much simpler-looking — but actually equivalent — conjecture, which is phrased only in terms of distributions of permutations. The key lemma for deriving this equivalent conjecture is the following:

Lemma 6.3 *$\mathcal{G}$ contains the edge (y, x) if and only if there exists a set $Z_{x,y}$ of candidates with $x \in Z_{x,y}$ and $y \notin Z_{x,y}$ such that*

$$||[y \succ Z_{x,y}]|| + ||[\overline{Z_{x,y}} \succ x]|| > m. \tag{3}$$

Based on Lemma 6.3, we formulate the following conjecture, and prove it equivalent to Conjecture 1.

Conjecture 2 *Let $Z_1, Z_2, \ldots, Z_n \subseteq \{1, \ldots, n\}$ be arbitrary sets with $i \in Z_i$. Define the following two indicator functions over elements i and total orders $\succ$:*

$$\alpha(i, \succ) = \begin{cases} 1 & \text{if } i+1 \succ Z_i \\ 0 & \text{otherwise;} \end{cases} \quad \beta(i, \succ) = \begin{cases} 1 & \text{if } \overline{Z_i} \succ i \\ 0 & \text{otherwise;} \end{cases} \tag{4}$$

here all additions/subtractions are modulo n; that is, $n + 1 := 1$, and $1 - 1 := n$.

Let $\mathcal{D}$ be any distribution over total orders $\succ$ of $\{1, \ldots, n\}$. Then, there exists an i such that

$$\mathbb{E}_{\succ \sim \mathcal{D}} [\alpha(i, \succ) + \beta(i, \succ)] \leq 1. \tag{5}$$

Proposition 6.4 *Conjecture 2 is true if and only if Conjecture 1 is true.*

6.3 A Graph-Theoretic Reformulation

Our attempts to prove Conjecture 2 (so far unsuccessful) have been based on proofs by contradiction. The assumed constraints from Conjecture 2 prescribe several constraints on rankings that must hold simultaneously; using transitivity, this leads to a contradiction by forcing preferences to contain cycles. The essence of this approach is captured by another conjecture. To formulate it, we define the following class of directed graphs, which we term *Constraint-Choice Graphs*.

Definition 6.5 (Constraint-Choice Graph) *Let $Y_n = \{y_1, \ldots, y_n\}, A_n = \{a_1, \ldots, a_n\}, B_n = \{b_1, \ldots, b_n\}$ be three disjoint sets of nodes. A* constraint-choice graph *for a given n contains $3n$ nodes $U_n = Y_n \cup A_n \cup B_n$, and the following edges:*

- *For each i, it contains the directed edges*[5] *$(y_i, a_{i-1}), (y_i, b_{i-1}), (a_i, y_i), (b_i, y_i)$.*
- *For each i, j with $j \neq i, j \neq i - 1$, it contains exactly one of the two directed edges $(a_j, y_i), (y_i, b_j)$.*

An example of a constraint choice graph is shown in Figure 2. The edges listed first in Definition 6.5 are shown in solid black, while the edges listed second are shown in dashed red lines.

Conjecture 3 *For every n and every constraint choice graph G_n of $3n$ nodes, there exists a non-empty index set $S \subseteq \{1, \ldots, n\}$ with the following property: For every vertex set $T \subseteq \{a_i, b_i \mid i \in S\}$ of size $|T| > |S|$, the induced subgraph $G_n[Y_n \cup T]$ contains a directed cycle.*

[4]Recall that a directed graph is *Hamiltonian* if it contains a directed cycle of all nodes.

[5]As before, we define $1 - 1 := n$ and $n + 1 := 1$.

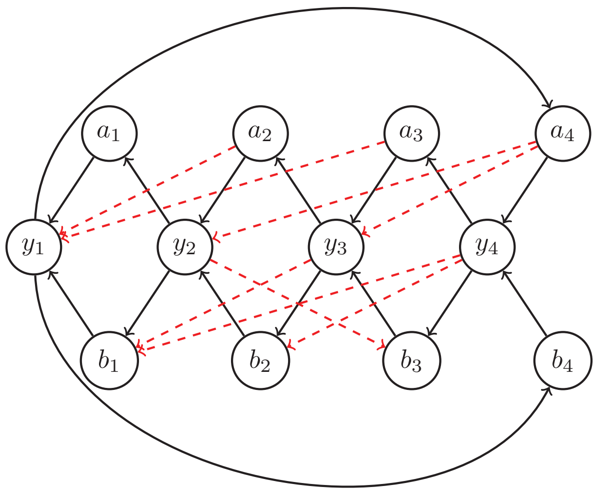

Figure 2: An illustration of a constraint-choice graph for $n = 4$ candidates.

Remark 6.6 Notice that the conjecture indeed talks about the subgraph induced by *all* nodes y_i (not just those with indices in T), in addition to at least $|S|+1$ nodes from among the a_i, b_i with $i \in S$.

Proposition 6.7 *If Conjecture 3 is true, then Conjecture 2 is true.*

7 Conclusions

Our work raises a very obvious question: prove (or possibly disprove) the conjectures stated in Section 6. Based on exhaustive computer search, it seems more likely that the conjectures are true, and the ALL BIPARTITE MATCHINGS mechanism in fact is always able to find a candidate with distortion at most 3.

Going beyond these conjectures, we believe that the duality-based framework may be useful for bounding the performance of other voting mechanisms, in particular, those that may miss information on parts of voters' ranking. For instance, such a situation can occur in the setting of (Kempe 2020), where voters can only name the candidates in a subset of positions on their ballot, rather than giving a complete ranking. The analysis of a mechanism proposed in (Kempe 2020) becomes much simpler (and tighter) using the techniques developed here.

While we have only studied deterministic mechanisms here, the framework can also be extended to randomized mechanisms. When the mechanism selects a candidate x with probability q_x, an upper bound can be obtained by bounding a flow that inserts q_x units of flow at each of the nodes (v, x), which again have to be routed to x^*.

It is conceivable that duality-based approaches similar to the one we developed could be helpful for the analysis of mechanisms for other problems in the cardinal/ordinal framework: a worst-case metric for a given input can often be characterized in terms of a linear program, and the dual may in general lead to a framework for proving upper bounds on the performance of a chosen mechanism.

Acknowledgements The author would like to thank Elliot Anshelevich, Sid Banerjee, Shaddin Dughmi, Bobby Kleinberg, and Kai Wang for useful conversations and pointers, and anonymous reviewers for useful feedback.

References

Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264:27–51.

Anshelevich, E.; Bhardwaj, O.; and Postl, J. 2015. Approximating optimal social choice under metric preferences. In *Proc. 29th AAAI Conf. on Artificial Intelligence*, 777–783.

Anshelevich, E. 2016. Ordinal approximation in matching and social choice. *ACM SIGecom Exchanges* 15(1):60–64.

Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227:190–213.

Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

Caragiannis, I., and Procaccia, A. D. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175(9):1655–1671.

Goel, A.; Krishnaswamy, A. K.; and Munagala, K. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proc. 18th ACM Conf. on Economics and Computation*, 287–304.

Kempe, D. 2019. An analysis framework for metric voting based on LP duality. arXiv: 1911.07162.

Kempe, D. 2020. Communication, distortion, and randomness in metric voting. In *Proc. 34th AAAI Conf. on Artificial Intelligence*.

Moulin, H. 1986. Choosing from a tournament. *Social Choice and Welfare* 3(4):271–291.

Munagala, K., and Wang, K. 2019. Improved metric distortion for deterministic social choice rules. In *Proc. 20th ACM Conf. on Economics and Computation*, 245–262.

Procaccia, A. D., and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *Proc. 10th Intl. Workshop on Cooperative Inform. Agents X*, 317–331.

Procaccia, A. D. 2010. Can approximation circumvent Gibbard-Satterthwaite? In *Proc. 24th AAAI Conf. on Artificial Intelligence*, 836–841.

Schulze, M. 2011. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social Choice and Welfare* 36(2):267–303.

Tideman, T. N. 1987. Independence of clones as a criterion for voting rules. *Social Choice and Welfare* 4(3):185–206.