# Byzantine Preferential Voting 

Darya Melnyk Yuyi Wang Roger Wattenhofer<br>dmelnyk@ethz.ch<br>yuwang@ethz.ch<br>wattenhofer@ethz.ch

ETH Zürich, Switzerland


#### Abstract

In the Byzantine agreement problem, $n$ nodes with possibly different input values aim to reach agreement on a common value in the presence of $t<n / 3$ Byzantine nodes which represent arbitrary failures in the system. This paper introduces a generalization of Byzantine agreement, where the input values of the nodes are preference rankings of three or more candidates. We show that consensus on preferences, which is an important question in social choice theory, complements already known results from Byzantine agreement. In addition preferential voting raises new questions about how to approximate consensus vectors. We propose a deterministic algorithm to solve Byzantine agreement on rankings under a generalized validity condition, which we call Pareto-Validity. These results are then extended by considering a special voting rule which chooses the Kemeny median as the consensus vector. For this rule, we derive a lower bound on the approximation ratio of the Kemeny median that can be guaranteed by any deterministic algorithm. We then provide an algorithm matching this lower bound. To our knowledge, this is the first non-trivial multi-dimensional approach which can tolerate a constant fraction of Byzantine nodes.


Keywords: Social Choice, Byzantine Agreement, Pareto -Validity, Relative Rankings, Arrow Impossibility, Distributed Voting, Multivalued

## 1 Introduction

In the first-past-the-post (FPTP) voting mechanism, each voter indicates his/her candidate of choice on a ballot, and the candidate that received the most votes wins. This is a popular voting mechanism that for instance is used in the United States. FPTP voting is related to consensus or Byzantine agreement in the sense that each voter/node has an input, and the voters need to decide on a single output. If Byzantine voters are present, it is natural to try to agree on a candidate with many votes to be robust against Byzantine behavior, e.g., 7, 8, 10]. FPTP does however face a lot of criticism since it has several issues, in particular the problem of wasted votes to minority parties, but also tactical voting or gerrymandering [30]. Moreover, one may argue that FPTP will eventually lead to a two-party system, e.g., 34].

Preferential voting is a powerful alternative: Each voter ranks the candidates first, second, third, etc. The collection of all rankings forms a preference profile from which a winner (or even a ranking) is determined. Preferential voting is more expressive, harder to manipulate, and solves many of FPTP's problems. In particular there is no problem with wasted votes to minority candidates. If inputs are not just binary, preferential voting will lead to much better decisions. It is therefore remarkable that byzantine agreement research has not given preferential voting any attention.

In this paper we want to investigate how robust preferential voting is in a Byzantine environment. In Section2, we first focus on some basic properties for voting rules, and see that not all of them can be satisfied if the nodes should reach agreement. This is because Byzantine voters are manipulators that modify the result to make it more favorable to themselves. In the main part of the paper (Section (5) we study how well the voting result intended by the correct (non-Byzantine) voters can be approximated. For this purpose we introduce the Kemeny rule which picks the most central ranking as the voting result. We will provide an algorithm that approximates the solution of the Kemeny rule in the presence of Byzantine voters and prove that this algorithm computes the best possible approximation. We believe that our paper will help to get a deeper understanding of both fault-tolerant distributed systems as well as social choice theory.

## 2 Background and Motivation

In search of a fair rule to elect candidates, philosophers and mathematicians started developing various voting mechanisms and rules already in the beginning of the 18th century. In the middle of the 20th century, Kenneth Arrow [2, 3] was one of the first to formalize existing rules and analyze possibility and impossibility results in an axiomatic fashion, thereby introducing the field of Computational Social Choice. In this section we use this formalism in order to show how well Byzantine agreement connects to voting theory.

We start by considering the special case of $n$ voters voting on only two candidates $c_{1}$ and $c_{2}$. In this setting, each voter (node) ranks the two candidates such that its preferred candidate (input value) is ranked first. A vote for a candidate $c_{1}$ means that the voter strictly prefers $c_{1}$ to $c_{2}$, here denoted $c_{1} \succ c_{2}$. A central authority then applies a social choice function (SCF) to a given preference profile in order to determine the winner (decision value), or set of winners in case of a tie. An SCF $f$ can be qualified based on the following properties:

- $f$ is anonymous if interchanging two voters (swapping their names) does not change the result
- $f$ is neutral if renaming the candidates (changing their names) does not change the result
- $f$ is positively responsive if in a case where the decision is a tie ( $c_{1}$ is among the winners) and a voter changes its ranking from $c_{2} \succ c_{1}$ to $c_{1} \succ c_{2}$, candidate $c_{1}$ becomes the unique winner

One example of an SCF is the majority rule. It chooses the candidate that wins most pairwise comparisons against every other candidate. Note that such a winner always exists in elections with
two candidates, but not necessarily in the general case with any number of candidates. Social choice theory shows that the majority rule satisfies all desirable properties for the special case of voting on two candidates:

Theorem 1 (May's Theorem [25]). For two candidates and any number of voters, the majority rule is the unique SCF that satisfies anonymity, neutrality and positive responsiveness.

Interestingly, most known algorithms for binary Byzantine agreement indirectly exploit the properties of May's theorem. Some of them make use of leaders who suggest their decision value to all nodes $[8,9]$. The leader in these algorithms temporarily plays what is known as a dictator in voting theory. Another type of algorithm, e.g., the shared coin algorithm in [10], is biased towards one of the outcomes and thus violates neutrality. In general we can say that most of the proposed algorithms try to use the majority value as the decision value if a majority exists, or an arbitrary input value otherwise, see for example [7, 10]. Such settings may satisfy anonymity and neutrality, but in cases where the correct nodes are undecided, i.e. there is a tie between the two input values, Byzantine nodes have a large influence on the majority value. Thus, if a correct node decides to swap two candidates in its ranking in order to make one of the candidates win, a Byzantine node can perform an opposite swap in its own ranking and return the profile to the previous state. This shows that positive responsiveness cannot be satisfied for these algorithms in the presence of Byzantine nodes.

May's theorem does not apply to the general case with more than two candidates. In fact, the majority rule gives surprisingly bad results for three or more candidates. To illustrate this, let $m$ denote the number of candidates. Assume that $n / 2+1$ voters rank the candidates as $c_{1} \succ c_{2} \ldots \succ c_{m}$, and all other voters rank the candidates as $c_{2} \succ c_{3} \succ \ldots \succ c_{1}$. In this case candidate $c_{1}$ wins every pairwise comparison according to the majority rule, even though $c_{2}$ seems to be the candidate that is approved by more voters.

Moreover, a lot of information is lost when a single winner is sought. When it comes to preferential voting, social choice theory therefore often wants not only the input to be rankings but also the output. More formally:

Definition 1 (Social Welfare Function). A Social Welfare Function (SWF) is a map from a preference profile to a set of consensus rankings.

For an SWF $g$, the following three properties are usually considered:

- $g$ is dictatorial if there is one distinguished voter whose input ranking is chosen as the single consensus ranking
- $g$ is independent of irrelevant alternatives (IIA) if the consensus ranking of two candidates $c_{i}$ and $c_{j}$ only depends on the relative preference of these candidates in each voter's ranking, and not on the ranking of some third candidate $c_{k}$
- $g$ is weakly Paretian if it satisfies the weak Pareto condition [31]: for two candidates $c_{i}$ and $c_{j}$ which are ranked $c_{i} \succ c_{j}$ by all voters, consensus ranking has to rank $c_{i} \succ c_{j}$ as well

In contrast to IIA and weak Pareto, dictatorship is a highly undesirable property in voting theory. Unfortunately, a corresponding result to May's theorem for SWF's on three or more candidates is the famous impossibility result by Arrow:

Theorem 2 (Arrow's Impossibility Theorem [2]). If there are at least three candidates which the members of the society are free to order in any way, then every SWF that is weakly Paretian and IIA must be dictatorial.

From the viewpoint of Byzantine agreement, an SWF should not be dictatorial since one does not want a dictator to be a Byzantine node and choosing more than one dictator may also result in different decision values. Consequently, any reasonable Byzantine agreement protocol must either violate IIA
or weak Pareto. The IIA condition implies that the consensus ranking should remain the same if the input of every correct node does not change, no matter what the Byzantine nodes do. However, a Byzantine node can pretend to be a correct node but change its ranking in two executions in which the correct nodes have the same inputs. This change may lead to a different consensus ranking which would violate IIA. For the weak Pareto condition consider the case with two candidates: if every nonByzantine voter ranks $c_{1} \succ c_{2}$, the consensus ranking should also rank $c_{1} \succ c_{2}$. This corresponds to a well-known validity condition in Byzantine agreement - the All-Same-Validity: If all correct nodes have the same input value, all correct nodes have to decide on this value. We use the weak Pareto condition to impose a validity rule on Byzantine Agreement with rankings:

Pareto-Validity for any pair of candidates $c_{i}$ and $c_{j}$ : if all correct nodes rank $c_{i} \succ c_{j}$, then the consensus ranking should rank $c_{i} \succ c_{j}$ as well.

Given $m$ candidates, Pareto-Validity can be viewed as All-Same-Validity applied on each of the $\binom{m}{2}$ pairs of candidates in a ranking. Note that Byzantine agreement on a ranking is at least as hard as binary Byzantine agreement: Consider a case where the nodes agree on the ranking of the candidates $c_{3}, \ldots c_{m}$ which they rank last, but not on the two first candidates $c_{1}$ and $c_{2}$. The Pareto condition is then satisfied for every binary relation which contains at least one of the candidates $c_{3}, \ldots c_{m}$. Agreement in this case is reduced to binary Byzantine agreement on the two candidates $c_{1}$ and $c_{2}$, under the All-Same-Validity condition.

Unfortunately, there is no straightforward way to apply a binary Byzantine agreement protocol to solve Byzantine agreement on rankings. Other than binary relations on two candidates, preference profiles can form cycles, e.g., they can contain all three relations $c_{i} \succ c_{j}, c_{j} \succ c_{k}$ and $c_{k} \succ c_{i}$ which are each preferred by a majority of nodes. The smallest preference profile which produces such a cycle of binary relations is called a Condorcet cycle. It contains three rankings $c_{i} \succ c_{j} \succ c_{k}, c_{j} \succ c_{k} \succ c_{i}$ and $c_{k} \succ c_{i} \succ c_{j}$ which induce the three relations from above. Simply agreeing on each pair of candidates can thus lead to a circular decision which does not form a ranking. In order to get rid of cycles one could think of applying the quicksort algorithm on the candidates sorted with respect to the majority. This procedure however violates Pareto-Validity: Consider a candidate $c_{i}$ that Pareto dominates candidate $c_{j}$. Assume that the quicksort algorithm compares both candidates to some third candidate $c_{k}$ first. Then $c_{j}$ might win against $c_{k}$ and $c_{i}$ might lose, thus swapping $c_{i}$ and $c_{j}$ in the consensus ranking. This consideration makes the problem of finding a consensus ranking in the presence of Byzantine nodes rather an instance of multi-valued agreement, as we discuss in Section 4, which makes the problem both interesting and challenging.

## 3 Related Work

Byzantine agreement was first proposed as the Byzantine Generals problem by Pease, Shostak and Lamport [24, 32]. In these papers the authors showed that three nodes cannot establish agreement in the presence of one Byzantine node even if the communication system is synchronous. Given $n$ nodes, it was shown for the synchronous model that at least $t+1$ rounds are required to establish agreement [18], where $t<n / 3$ is the number of Byzantine nodes in the system; the corresponding upper bound was provided in [8, 9]. For the asynchronous model, the FLP impossibility result [19] states that there is no deterministic agreement protocol which can tolerate even one Byzantine node. The first randomized algorithm for solving Byzantine agreement proposed in [7] had expected exponential running time for a constant fraction of Byzantine nodes. This result was recently improved in [23], where the authors showed that it is possible to establish agreement within expected polynomial running time using spectral methods.

Byzantine agreement with more than two input values has mostly been considered in approximate agreement [15, 17], where the input values of the nodes converge towards some value over rounds.

More recent results seek to establish agreement on a value that makes sense for applications. In (14], the values converge towards a value at most $\sqrt{n \log n}$ positions away from the median. In [26, 36] an exact algorithm to establish agreement on a value that is at most $t / 2$ positions away from the median or $t$ positions away from a minimum or a maximum was proposed. In [27, 28, 37], Byzantine agreement was further generalized to several dimensions and the nodes converge to a vector inside the convex hull of all correct input vectors. While the one-dimensional case has been investigated in depth, all previous approaches for multiple dimensions struggle to derive an algorithm which either can tolerate a constant fraction of Byzantine nodes independent on the number of dimensions, or find a solution that is not trivial.

In social choice theory, Byzantine behavior can be interpreted as manipulation of a ballot in an election, in which the manipulating party has full knowledge about all votes. Bartholdi et al. [5] defined manipulation as a preference profile where one single voter can change its ranking such that this voter's most preferred candidate wins the election. Groups of voters have also been considered in this context, but mostly from the perspective of how hard it is for a group of nodes to manipulate the voting result given a certain voting rule [12, 29]. Other types of Byzantine behavior have been considered with respect to robustness of proposed voting rules. In [6], the authors investigate robustness of Borda's mean and median in the presence of outlier ballots. In [33], robustness of scoring rules is considered under arbitrary noise which is described in terms of pairwise swaps of candidates in the ranking of one voter.

In this paper we will consider the Kemeny rule which was first proposed in 20, 21]. The corresponding Kemeny median satisfies additional properties to those presented in Section 2 but it was shown to be NP-hard to compute for an increasing number of candidates and already for four voters in [4, 16]. At least three different 2-approximation algorithms for the Kemeny median have been proposed in [1] and [13] respectively. In [1], the approximation ratio was improved to $4 / 3$ using randomization, and later derandomized in [38]. A good overview over the Kemeny rule and an extended introduction into social choice theory can be found in [11].

## 4 A Deterministic Algorithm for Pareto-Validity

This section focuses on Byzantine agreement protocols for rankings that satisfy Pareto-Validity. By using single transferable voting and a multi-valued Byzantine agreement algorithm, a ranking satisfying Pareto-Validity can be obtained in $(m-1) \cdot(t+1)$ rounds: In the first $t+1$ rounds, we let the voters apply the King algorithm in order to agree on the top candidate. Then every node removes this candidate from its ranking. In the next step, they will agree on the top candidate from the reduced rankings, and so on. While this procedure is simple, the number of rounds depends not only on the number of nodes, but also on the number of candidates.

In the following we present a deterministic algorithm which solves this problem in only $t+1$ rounds using the same number of messages. We do this by modifying the King algorithm to broadcast rankings instead of single candidates. In the proposed algorithm we select $t+1$ different nodes and assign each of them to one of the $t+1$ rounds of the algorithm. Such a node is called a dictator of the corresponding round. This dictator then suggest its own, possibly adjusted, ranking to all nodes, which will always be accepted if the dictator is a correct node. This way, dictators decide on the ranking of all pairs of candidates which do not satisfy the Pareto-Validity. Algorithm $\mathbb{1}$ presents this procedure in pseudocode.

Since we are dealing with rankings, it is not trivial to see that the nodes will always be able to agree on a proper ranking at the end of the algorithm. In the following lemmas we will prove that the nodes can adjust their rankings in Step 8 of Algorithm 1 in order to guarantee Pareto-Validity and that the outcome of the algorithm will be a proper ranking. It is easy to see that the algorithm is correct for $t<n / 4$ Byzantine nodes, since the correct nodes will not be able to propose binary relations which form a Condorcet cycle in this case. In order to show that the algorithm can tolerate

```
Algorithm 1 Byzantine agreement protocol on rankings (for \(t<n / 3\) )
    Every node \(v\) executes the following algorithm
    for round 1 to \(t+1\) do
        Communication Phase:
            Broadcast own input ranking \(r_{v}\)
            for all pairs of candidates \(c_{i}\) and \(c_{j}\) do
                        if \(c_{i}\) is ranked above \(c_{j}\) in at least \(n-t\) rankings then
                        Broadcast "propose \(c_{i} \succ c_{j}\) "
                end if
            end for
            if some "propose \(c_{k} \succ c_{l}\) " received at least \(t+1\) times fix \(\mathbf{c}_{\mathbf{k}} \succ \mathbf{c}_{\mathbf{1}}\)
                Adjust own ranking \(r_{v}\) such that \(c_{k}\) appears before \(c_{l}\)
            end if
        Dictator Phase:
            Let node \(w\) be the predefined dictator of the current round
            The dictator broadcasts its ranking \(r_{\text {dictator }}:=r_{w}\)
        Decision Phase:
            if \(r_{\text {dictator }}\) agrees with \(r_{v}\) in all fixed pairs \(c_{i} \succ c_{j}\) from step 8 then
                \(r_{v}:=r_{\text {dictator }}\)
            end if
    end for
    Return \(r_{v}\)
```

$t<n / 3$ Byzantine nodes, we need to exploit the fact that no Byzantine node can propose relations that form a Condorcet cycle at any point of the algorithm.

Lemma 1. There is no Condorcet cycle that can be proposed by the correct nodes if $t<n / 3$.
Proof. Assume by means of contradiction that the three relations $c_{i} \succ c_{j}, c_{j} \succ c_{k}$ and $c_{k} \succ c_{i}$ were each proposed by at least $t+1$ nodes in Step 4 of Algorithm 1 . Each binary relation was proposed by at least one correct node who must have seen $n-t$ nodes having a ranking with such a pair.

Let $t_{1}$ be the number of all Byzantine nodes who proposed $c_{i} \succ c_{j}$, $t_{2}$ the number of those who proposed $c_{j} \succ c_{k}$ and $t_{3}$ those nodes who proposed $c_{k} \succ c_{i}$. Further, let $t_{1 \cap 2}$ denote the number of Byzantine nodes who proposed $c_{i} \succ c_{j} \succ c_{k}$. The following inequality then holds: $t_{1}+t_{2}-t_{1 \cap 2} \leq t$.

The number of correct nodes who proposed $c_{i} \succ c_{j} \succ c_{k}$ is then $\left(n-t-t_{1}\right)+\left(n-t-t_{2}\right)+t_{1 \cap 2}-$ $(n-t)=n-t-t_{1}-t_{2}+t_{1 \cap 2}$. The number of correct nodes who proposed $c_{k} \succ c_{i}$ is $n-t-t_{3} \geq n-2 t$. However, the two sets must have a nonempty intersection, since

$$
n-t-t_{1}-t_{2}+t_{1 \cap 2}+n-2 t-(n-t)=n-2 t-t_{1}-t_{2}+t_{1 \cap 2} \geq n-3 t>1 .
$$

Therefore, there must be at least one correct node who proposed $c_{i} \succ c_{j} \succ c_{k}$ and $c_{k} \succ c_{i}$ simultaneously. This is a contradiction.

Note that by the properties of the King algorithm, no two opposite binary relations can be proposed in Step 4 simultaneously. Lemma 1 additionally shows that a Condorcet cycle cannot be proposed in Step 4 and thus all proposed pairs can form a ranking. It remains to show that the nodes will always be able to adjust their rankings to incorporate the proposed pairs.

Lemma 2. In Step 8 a correct node will always be able to incorporate the proposed pairs into its own ranking.

Proof. This can be achieved by the following strategy: Divide the candidates into two sets. The first set contains all candidates which are in at least one of the pairs proposed by the $t+1$ nodes in Step 8 . This set of nodes will be ranked first. The second set will contain all candidates for which the node has not received any propose message. The candidates will be ranked second and will be dominated by all candidates from the first set. Next, we can rank all candidates in the first set according to the proposed relations, possibly leaving some pairs of the candidates not ranked. In the last step, all candidates which have not been ranked in each of the sets can be ranked by choosing binary relations from the local ranking of the node. This strategy outputs a ranking of candidates in which all proposed binary relations are satisfied.

The next theorem summarizes the correctness results of Algorithm 1 and states that the consensus ranking will be valid, which can be derived with help of previous lemmas. The corresponding proof can be found in Appendix A.

Theorem 3. At the end of Algorithm $\mathbb{1}$ all nodes will have agreed on the same ranking which additionally satisfies Pareto -Validity.

## 5 Kemeny Median with Byzantine Nodes

Weakly Paretian voting rules are often not sufficient to pick the most fair ranking from a set of individual preference rankings. In search of the best possible consensus ranking we have to add restrictions on the voting rules without violating the known impossibility results of [2]. This leads us to majoritarian SWFs, one of which is the Kemeny rule. In the following we will introduce the Kemeny rule and use it to derive a better consensus ranking in the presence of Byzantine nodes. Since Byzantine nodes have influence on the final ranking, the corresponding solutions can be qualified with respect to their approximation ratio which we define in Section 5.1. In Section 5.2, we will derive a lower bound on the approximation ratio of the Kemeny median in the presence of Byzantine nodes and further provide a matching upper bound in Section 5.3.

Definition 2 (Kendall's $\tau$ distance [22]). The Kendall's $\tau$ distance measures the distance between two rankings $r$ and $p$ on candidates $c_{1}, \ldots, c_{m}$ by counting pairs of candidates on which they disagree:

$$
\tau(r, p) \triangleq \mid\left\{\left(c_{i}, c_{j}\right) \mid c_{i} \succ_{r} c_{j} \text { and } c_{j} \succ_{p} c_{i}\right\} \mid .
$$

This metric $\tau$ on ballots can be extended to a distance function between a ranking $r$ and a profile $\mathcal{P}$ :

$$
\tau(r, \mathcal{P}) \triangleq \sum_{p \in \mathcal{P}} \tau(r, p) .
$$

Definition 3 (Kemeny median). For a given profile $\mathcal{P}$, the Kemeny median is the ranking $r$ which minimizes $\tau(r, \mathcal{P})$.

The Kemeny median satisfies many nice properties and to some extent guarantees that the chosen ranking is "fair". The most prominent quality is probably monotonicity: if voters increase a candidate's preference level, the ranking result either does not change or the promoted choice increases in overall popularity. This quality makes the median solution more robust to Byzantine behavior. The Kemeny rule is also a so called Condorcet method because if there exists a Condorcet winner, i.e., a candidate that wins all pairwise majority comparisons, it will always be ranked as the most popular choice. Besides, it only depends on the number of voters who prefer one alternative over the other and it is reinforcing, meaning that winners which were chosen independently by two different sets of voters will also become winners if the two groups are joined.

### 5.1 Byzantine Setting

The Kemeny median cannot be computed exactly in the presence of Byzantine nodes since they might suggest rankings which have a large distance to the Kemeny median of the correct nodes thus moving the median preference away from the actual median. A notion for approximate median rankings is therefore introduced as follows:

Definition 4 ( $\alpha$-approximation of Kemeny median). Let $m$ be a Kemeny median of a preference profile $\mathcal{P}$. An $\alpha$-approximation of $m$ is a preference ranking $m_{\alpha}$ satisfying

$$
\tau\left(m_{\alpha}, \mathcal{P}\right) \leq \alpha \cdot \tau(m, \mathcal{P})
$$

As an example consider binary agreement $(m=2)$ : Here $\tau$ counts the number of correct nodes who disagree with the consensus value. Any binary Byzantine agreement algorithm that satisfies All-Same-Validity will also satisfy $\alpha<n-t-1$.

Unlike binary agreement, it is not straightforward to see what a Byzantine node would choose as its ranking when the Kemeny rule determines the consensus ranking. Since the input vectors of nodes are rankings, each voter has to propose a strict order between candidates and the corresponding preference relation is transitive. A possible strategy for the Byzantine nodes would then be to choose exactly the opposite ranking of the Kemeny median of all correct nodes. We show in Appendix B that this strategy works, but such a solution is not unique for most preference profiles. It is therefore difficult for the correct nodes to find out which of the rankings might have been Byzantine.

### 5.2 Lower Bounds on the Approximation Ratio

In this section we discuss preference profiles that are vulnerable to Byzantine nodes. The first case is based on the reduction of rankings to binary agreement and gives the highest approximation ratio for $t<n / 3$. Binary agreement does however assume that there are two groups of voters who completely disagree in their preferences. This is somewhat unlikely in practice when $m$ is sufficiently large. In the second case we therefore exclude such binary instances and provide a lower bound based on Condorcet cycles within a preference profile which converges to the same value for large $m$. The approximation ratio usually depends on the ratio $n / t$, which will be denoted $k$ for the sake of simplicity.

For our analysis, we represent the preference profile $\mathcal{P}$ as a weighted tournament graph, i.e., a graph where the nodes represent the candidates and weighted edges represent how many voters prefer one candidate to the other. The sum of the forward and the backward edge should be equal to the total number of rankings in the corresponding preference profile. The ranking of a node is a directed Hamiltonian path following the order of the ranking, and all other edges are derived from the transitivity. For any two candidates we call the edge between these candidates a majority edge if its backward edge has a smaller weight. The backward edge we then call a minority edge. A Kemeny median of a weighted tournament graph is the ranking that minimizes the sum of the weights of all backward edges of the graph. Note that rankings restrict the power of Byzantine nodes in the sense that Byzantine nodes can only send transitive tournament graphs where every edge has weight 1 . Using tournament graphs, we can derive a lower bound for the binary case:

Theorem 4. There is a tournament graph corresponding to a preference profile for which the Byzantine nodes may change the edge weights such that the median of the resulting preference profile is $a \frac{k}{k-2}$-approximation of the optimal median, where $k=n / t$. For $t$ close to $n / 3$, this gives a 3approximation.

Proof. This tournament graph is equivalent to binary agreement. Consider therefore one pair of candidates: $t$ Byzantine nodes are only able to change the median, i.e., the majority edge, between two candidates if they can swap the majority and minority edge by supporting the minority edge with


Figure 1: Two indistinguishable views on $m$ candidates for binary relations
Note that the labels of the edges correspond to all edges in the same color. The left tournament graph is reached if $n / 2$ nodes choose the order $c_{1} \succ c_{2} \succ \ldots \succ c_{m}$ and $n / 2-t$ nodes choose $c_{m} \succ c_{m-1} \succ \ldots \succ c_{1}$. The right profile can be reached from a profile where $n / 2$ nodes choose the order $c_{m} \succ c_{m-1} \succ \ldots \succ c_{1}$ and $n / 2-t$ nodes choose $c_{1} \succ c_{2} \succ \ldots \succ c_{m}$. The Byzantine nodes can make all correct nodes see the tournament graph in the center by adding $t$ preference vectors $c_{m} \succ c_{m-1} \succ \ldots \succ c_{1}$ or $c_{1} \succ c_{2} \succ \ldots \succ c_{m}$ respectively.
their ranking. Assume the worst case, where the forward and the backward edge both have the same weight $n / 2$ after the Byzantine nodes have added their preferences. In the worst case the tournament graph of correct nodes had the weight $n / 2$ for the majority edge. Since the correct nodes will not be able to determine the actual majority edge, they might agree on a minority edge with weight $n / 2-t$ instead. The corresponding approximation ratio is then $\frac{n / 2}{n / 2-t}=\frac{k}{k-2}$.

Figure 5.2 shows a simple generalization of this argument to $m$ candidates and proves that the lower bound of $\frac{k}{k-2}$ holds, for all $m$.

Now we present another lower bound using Condorcet cycles which can result in ambiguous views as well. Here we assume that every majority edge has a weight of more than $n / 2$, thus discarding the possibility to reduce any pair of forward and backward edges in the tournament graph to binary agreement. The main difficulty in finding a good example comes from the fact that not every tournament graph has an underlying preference profile. In Appendix $C$ we discuss the necessary properties that preference profiles induce on tournament graphs and show how the best lower bound can be derived. The next theorem present the best lower bound for cycles and its generalization to $m$ candidates.

Theorem 5. By modifying directed majority cycles in the tournament graph, Byzantine nodes can increase the approximation ratio by a factor of at most 5/4.

Proof. We start by considering a tournament graph formed by one directed cycle of candidates $c_{1}, c_{2}, c_{3}$, i.e., one directed cycle formed by majority edges. Assume all correct nodes receive a view where $n-2 t$ nodes prefer $c_{1}$ to $c_{2}$, i.e., $\left(c_{1}, c_{2}\right)$ is a majority edge; $n / 2+t$ nodes prefer $c_{2}$ to $c_{3}$ and $n / 2+t$ nodes prefer $c_{3}$ to $c_{1}$. For $n>8 f$ or equivalently $k>8$, the edge $\left(c_{1}, c_{2}\right)$ is in the median ranking of all nodes. Since the edges $\left(c_{2}, c_{3}\right)$ and ( $c_{3}, c_{1}$ ) cannot be both in the median ranking, the nodes have to decide for one of the rankings. In the worst case, one of these two edges was supported by all $t$ Byzantine nodes while the other edge was not supported by any Byzantine node. This leads to two views which are not distinguishable for the correct nodes, as shown in Figure 5.2. The approximation ratio for these views is $\frac{2 t+n / 2-t+n / 2+t}{2 t+n / 2-2 t+n / 2}=\frac{n+2 t}{n}=\frac{k+2}{k}<\frac{5}{4}$.

An extension to $m$ candidates gives an approximation ratio of $\frac{2 t+(m-2) \cdot(n / 2-t)+(m-2) \cdot(n / 2+t)}{2 t+(m-2) \cdot n / 2+(m-2) \cdot(n / 2-2 t)}=$ $\frac{2 t+(m-2) \cdot n}{2 t+(m-2) \cdot(n-2 t)}=\frac{2+(m-2) k}{2+(m-2) \cdot(k-2)} \approx \frac{k}{k-2}$ for large $m$.

The received approximation ratio converges to the same approximation ratio as in the binary case for large $m$.


Figure 2: Two indistinguishable views on $m$ candidates for directed cycles
We have two views which show the profiles of correct nodes only. The left tournament graph results from a profile where $n / 2-t$ nodes choose $c_{1} \succ c_{2} \succ c_{m} \succ \ldots \succ c_{3}, n / 2-2 t$ nodes choose $c_{m} \succ \ldots \succ c_{3} \succ c_{1} \succ c_{2}$ and $2 t$ nodes choose $c_{2} \succ c_{m} \succ \ldots \succ c_{3} \succ c_{1}$. The right tournament graph results from $n / 2-2 t$ nodes choosing $c_{1} \succ c_{2} \succ c_{m} \succ \ldots \succ c_{3}, n / 2-t$ nodes choosing $c_{m} \succ \ldots \succ c_{3} \succ c_{1} \succ c_{2}$ and $2 t$ nodes choosing $c_{2} \succ c_{m} \succ \ldots \succ c_{3} \succ c_{1}$. If the Byzantine nodes add $t$ profiles $c_{m} \succ \ldots \succ c_{3} \succ c_{1} \succ c_{2}$ to the left view, and $t$ profiles $c_{1} \succ c_{2} \succ c_{m} \succ \ldots \succ c_{3}$ to the right view, the resulting profiles become indistinguishable to the correct nodes.

### 5.3 Algorithm for Kemeny Median Approximation

In this section we present a synchronous algorithm for computing a consensus median which matches the lower bound on the approximation ratio presented in the previous section. A simple idea is to use interactive consistency [10, 35$]$ : For $t+1$ rounds, the nodes exchange all information they have received until that round and after the $(t+1)$-st round they compute the Kemeny median from a set of rankings which they have received often enough. This algorithm guarantees that the set of rankings will be the same for each node and therefore all nodes will decide on the same ranking. The main drawback of interactive consistency is that it has a large message complexity. Since each of the correct rankings will be forwarded by each of the correct nodes, this message complexity is in $\Theta\left(m n^{t}\right)$.

Instead of exchanging large amounts of information, we can directly exploit the fact that the Byzantine nodes cannot change a Kemeny median of the preference profile of the correct nodes by more than a transitive tournament graph with edge weights $t$. The corresponding strategy is presented in Algorithm 2,

```
Algorithm 2 Byzantine agreement protocol for the Kemeny median (for \(t<n / 3\) )
    Every node \(v\) executes the following algorithm
    broadcast own ranking \(r_{v}\)
    compute the Kemeny median of the received preference profile, call it \(m_{v}\)
    apply Algorithm 1 with \(m_{v}\) as an input value
```

The presented algorithm has the same order of round and message complexity as Algorithm 1 .
Theorem 6. Algorithm 圆 terminates within $t+3$ rounds exchanging $O\left(t n^{2} m \log m\right)$ messages. The computed consensus ranking satisfies the lower bounds from Section 5.2 and Pareto-Validity.

In the following we give two lemmas proving the correctness of this algorithm. A full proof of Theorem 6 is provided in Appendix D.

Lemma 3．In Step 图 of Algorithm 图，every correct node chooses a median ranking that matches the bounds from Section 5．2．
Proof．Instead of all nodes in the previous section，we can consider that the Byzantine nodes change just one node＇s view．Since the number of Byzantine nodes remains the same and the rankings of all correct nodes are received by every node in the synchronous communication model，the Byzantine nodes can in the worst case only reach the lower bound for any correct node，but not exceed it．

Lemma 4．The computed median ranking by Step 园 of the algorithm satisfies the approximation ratios from Section 5．2．
Proof．Observe that the consensus ranking is derived from a preference profile formed by the medians $m_{v}$ ．Unless some correct node disagrees on an edge in this profile，the edge will be inside the consensus median since Algorithm 1 satisfies Pareto－Validity．This edge may not be inside the median ranking of all correct nodes，but the approximation ratio still satisfies the bounds due to Lemma 3．

If the correct nodes disagree on an edge in the preference profile of medians，there was at least one correct node who either chose the opposite edge（binary case）for its median ranking or a different edge in a directed cycle（non－binary case）．Consider the binary case first．There，the forward and the backward edge chosen as the Kemeny median $m_{v}$ will both satisfy the lower bound，since there is a correct node choosing either of the cases．For the non－binary case we need to consider directed cycles formed by the median rankings．Every directed cycle in a tournament graph implies that there is a directed sub－cycle formed by three candidates $c_{i}, c_{j}, c_{k}$ ．The corresponding preference profile of correct median rankings must contain the three opposite rankings $c_{i} \succ c_{j} \succ c_{k}, c_{j} \succ c_{k} \succ c_{i}$ and $c_{k} \succ c_{i} \succ c_{j}$ ．This，however，implies that all three median rankings could have been derived from the preference profile of the rankings $r_{v}$ by modifying edge weights by $t$ ．The only case from which such a situation can result is when the forward and backward edge weights of each of the three pairs of candidates differed by at most $t$ in the preference profile of rankings $r_{v}$ ．This case，again，is equivalent to the binary case and satisfies the lower bound．

Note that the computed median can have a larger Kendall＇s $\tau$ distance to the preference ranking of all correct nodes than any $m_{v}$ has in the algorithm，since the Byzantine nodes can propose their own rankings as dictators in Algorithm 1．Such a ranking would still satisfy the lower bound．

## 6 Discussion and Future Work

In this paper we introduced a new Byzantine agreement problem which extends binary Byzantine agreement to rankings．We showed that rules for choosing a consensus ranking in voting theory fit well with requirements from Byzantine agreement．We further considered a special voting rule，the Kemeny median，for which we provided an optimal Byzantine agreement protocol that can tolerate up to $t<n / 3$ Byzantine nodes．We do not claim to have chosen the best voting rule at this point， since such a rule simply does not exist due to impossibility results in voting theory．Instead，we think of our results as an inspiration to consider a larger pool of voting rules，such as approval voting，the Godgson＇s rule，and many others．

Byzantine agreement on rankings can also be of interest in several applications．Consider dis－ tributed machine learning as an example．Training data is usually collected by different（potentially Byzantine）parties and is stored in different places．When the amount of data is very large or the data is too sensitive to share，it is impossible to transmit the data and then train a global model．Instead， these parties could train their own models locally and then vote to predict the labels for new data points．A data point could be a picture，and every trained model outputs a ranking on the possible labels（panda $\succ$ gibbon $\succ \ldots \succ$ cat）．The different parties could then use our Byzantine preferential voting protocols to find the best ranking．

## References

[1] Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating Inconsistent Information: Ranking and Clustering. Journal of the ACM, 55(5):23:1-23:27, November 2008.
[2] Kenneth J. Arrow. Social Choice and Individual Values. CT: Cowles Foundation, New Haven, 1st edition, 1951.
[3] Kenneth J. Arrow. Social Choice and Individual Values. John Wiley, New York, 2nd edition, 1963.
[4] J. Bartholdi, C. A. Tovey, and M. A. Trick. Voting Schemes for which It Can Be Difficult to Tell Who Won the Election. Social Choice and Welfare, 6(2):157-165, April 1989.
[5] J. J. Bartholdi, C. A. Tovey, and M. A. Trick. The Computational Difficulty of Manipulating an Election. Social Choice and Welfare, 6(3):227-241, July 1989.
[6] Gilbert W. Bassett and Joseph Persky. Robust voting. Public Choice, 990(3):299-310, June 1999.
[7] Michael Ben-Or. Another Advantage of Free Choice (Extended Abstract): Completely Asynchronous Agreement Protocols. In Proceedings of the Second Annual ACM Symposium on Principles of Distributed Computing, PODC '83, pages 27-30. ACM, 1983.
[8] Piotr Berman and Juan A. Garay. Asymptotically optimal distributed consensus. In Automata, Languages and Programming: 16th International Colloquium, ICALP, July 1989.
[9] Piotr Berman, Juan A. Garay, and Kenneth J. Perry. Towards Optimal Distributed Consensus. In 30th Annual Symposium on Foundations of Computer Science, FOCS, October 1989.
[10] Gabriel Bracha. Asynchronous Byzantine Agreement Protocols. Information and Computation, 75(2):130143, 1987.
[11] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. Handbook of Computational Social Choice. Cambridge University Press, New York, NY, USA, 1st edition, 2016.
[12] Jessica Davies, George Katsirelos, Nina Narodytska, and Toby Walsh. Complexity of and Algorithms for Borda Manipulation. volume abs/1105.5667, 2011.
[13] Persi Diaconis and R. L. Graham. Spearman's Footrule as a Measure of Disarray. Journal of the Royal Statistical Society. Series B (Methodological), 39:262-268, 1977.
[14] Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. Stabilizing Consensus with the Power of Two Choices. In Proceedings of the Twenty-third Annual ACM Symposium on Parallelism in Algorithms and Architectures, SPAA, June 2011.
[15] Danny Dolev, Nancy A. Lynch, Shlomit S. Pinter, Eugene W. Stark, and William E. Weihl. Reaching Approximate Agreement in the Presence of Faults. Journal of the ACM, 33(3):499-516, July 1986.
[16] Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. Rank Aggregation Methods for the Web. In Proceedings of the 10th International Conference on World Wide Web, WWW '01, pages 613-622, New York, NY, USA, 2001. ACM.
[17] Alan D. Fekete. Asymptotically optimal algorithms for approximate agreement. Distributed Computing, 4(1):9-29, 1990.
[18] Michael J. Fischer and Nancy A. Lynch. A Lower Bound for the Time to Assure Interactive Consistency. Information Processing Letters, 14(4):183-186, 1982.
[19] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of Distributed Consensus with One Faulty Process. Journal of the ACM, 32(2):374-382, April 1985.
[20] John G. Kemeny. Mathematics without Numbers. Daedalus, 88(4):577-591, 1959.
[21] John G. Kemeny and J. Laurie Snell. Mathematical models in the social sciences. Introductions to higher mathematics. Blaisdell, Waltham (Mass.), 1962.
[22] M. G. Kendall. A New Measure of Rank Correlation. Biometrika, 30(1/2):81-93, 1938.
[23] Valerie King and Jared Saia. Byzantine Agreement in Expected Polynomial Time. Journal of the ACM, 63(2):13:1-13:21, 2016.
[24] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine Generals Problem. ACM Transactions on Programming Languages and Systems, 4(3):382-401, July 1982.
[25] Kenneth O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision. Econometrica, 20(4):680-684, 1952.
[26] Darya Melnyk and Roger Wattenhofer. Byzantine Agreement with Interval Validity. 2017. Unpublished manuscript.
[27] Hammurabi Mendes and Maurice Herlihy. Multidimensional Approximate Agreement in Byzantine Asynchronous Systems. In Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC, June 2013.
[28] Hammurabi Mendes, Maurice Herlihy, Nitin Vaidya, and Vijay K. Garg. Multidimensional agreement in Byzantine systems. Distributed Computing, 28(6):423-441, January 2015.
[29] Nadja Betzler and Rolf Niedermeier and Gerhard J. Woeginger. Unweighted Coalitional Manipulation Under the Borda Rule Is NP-hard. IJCAI'11, pages 55-60. AAAI Press, 2011.
[30] Stephen Orvis and Carol Ann Drogus. Introducing comparative politics: Concepts and cases in context. CQ Press, 2013.
[31] V. Pareto. Manuale di Economia Politica con una Introduzione alla Scienza Sociale. Società Editrice Libraria, 1919.
[32] Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching Agreement in the Presence of Faults. Journal of the ACM, 27(2):228-234, April 1980.
[33] Ariel D. Procaccia and and Gal A. Kaminka Jeffrey S. Rosenschein. On the Robustness of Preference Aggregation in Noisy Environments. AAMAS '07, pages 66:1-66:7, New York, NY, USA, 2007. ACM.
[34] Jeffrey D Sachs. The price of civilization: Reawakening American virtue and prosperity. Random House, 2011.
[35] T.K. Srikanth and Sam Toueg. Simulating Authenticated Broadcasts to Derive Simple Fault-Tolerant Algorithms. Distributed Computing, 2(2):80-94, June 1987.
[36] David Stolz and Roger Wattenhofer. Byzantine Agreement with Median Validity. In 19th International Conference on Priniciples of Distributed Systems, OPODIS, December 2015.
[37] Nitin H. Vaidya and Vijay K. Garg. Byzantine Vector Consensus in Complete Graphs. In Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing, PODC, July 2013.
[38] Anke van Zuylen and David P. Williamson. Deterministic Algorithms for Rank Aggregation and Other Ranking and Clustering Problems, pages 260-273. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

## A Proof of Theorem 3

We will divide the proof of Theorem 3 into two parts:
Lemma 5. Algorithm $\mathbb{1}$ satisfies Pareto-Validity.
Proof. Assume that all correct nodes agree on a binary relation between two candidates, e.g. $c_{i} \succ c_{j}$. Then all $n-t$ correct nodes will propose $c_{i} \succ c_{j}$ in Step 4 and every correct node will receive at least $n-t$ propose messages for the same pair. Thus, no correct node will change its ranking until a round in which the current dictator agrees on this binary relation $c_{i} \succ c_{j}$ in Step 13. This holds for any pair of binary relations.

Lemma 6. At the end of the algorithm all nodes will have agreed on the same ranking.
Proof. We will show that all correct nodes will always adjust their ranking to the ranking of a correct dictator. For this we need to verify that every correct node will enter the if-clause in Step 13 if the dictator is a correct node. Note that any propose message that a correct node has received $n-t$ times in Step 8, will be received at least $n-2 t>t+1$ times by every other correct node who will incorporate this relation into its own ranking. This way every correct dictator will agree on all binary relations that any correct node has received in Step 8 of the algorithm.

At latest in a round with a correct dictator all nodes will adjust their ranking to the ranking of the dictator and, by Lemma 5 , will receive the same ranking at the end of the algorithm.

## B Opposite Ranking to the Kemeny Median

First we show that the opposite ranking of the Kemeny median always gives the worst possible solution which a Byzantine nodes can pick. Note that the sum of the weights of all edges in a tournament graph produced by the rankings of correct nodes only is $(n-t) \cdot\binom{m}{2}$. The tournament graph formed by a ranking $r$ and the tournament graph formed by the opposite ranking $\bar{r}$ are complementary graphs with respect to the tournament graph of all correct rankings. Consider the weights of all possible rankings. Each ranking that minimizes the Kendall's $\tau$ distance must have an opposite ranking that maximizes the same distance, i.e., $\tau(r, \mathcal{P})+\tau(\bar{r}, \mathcal{P})=(n-t) \cdot\binom{m}{2}$. This shows that the opposite ranking of the Kemeny median is also the ranking furthest away from it and can always be chosen as a Byzantine value.

The opposite ranking is however not the only ranking that the Byzantine nodes can choose. Assume all correct nodes agree on the preference $c_{i} \succ c_{j}$ such that this pair will always belong to the Kemeny median of the correct rankings. Then, the Byzantine nodes can pick either $c_{i} \succ c_{j}$ or $c_{j} \succ c_{i}$ for their ranking, since this strategy does not have any influence on the Kemeny median of all rankings. One example for such a case is depicted in Figure 5.2, where the Byzantine rankings are deliberately chosen not to be the opposite rankings of the Kemeny median of all correct nodes, although such a strategy would give equivalent results. The number of possible solutions which the Byzantine nodes can choose increases with the number of binary relations on which all correct nodes agree.

## C Proof of the Lower Bound for Cycles

Consider a cycle on the three candidates $c_{1}, c_{2}, c_{3}$ with majority edges $\left(c_{1}, c_{2}\right),\left(c_{2}, c_{3}\right)$ and $\left(c_{3}, c_{1}\right)$. Let the weights on majority edges $\left(c_{1}, c_{2}\right),\left(c_{2}, c_{3}\right)$ and $\left(c_{3}, c_{1}\right)$ be $x, y, z$ and the weights on corresponding minority edges $n-t-x, n-t-y, n-t-z$ respectively. In order to show that a cycle with such weights is the worst case we need to consider additional properties of a tournament graph:

- The weights of the forward and the backward edge between any two candidates sum up to the number of rankings which is $n-t$ without counting Byzantine rankings and $n$ if we count all rankings.
- Consider the tournament graph of correct nodes. Since we exclude cases which reduce to binary agreement and due to the definition of majority edges, we can assume $x \geq y \geq z \geq n / 2$.
- All weights satisfy the triangle inequality on directed edges, which states that for any triplet $i, j, k$ holds

$$
w_{i, j}+w_{j, k} \geq w_{i, k}
$$

where $w_{i, j}$ denotes the weight on the directed edge $\left(c_{i}, c_{j}\right)$. According to the triangle inequality, it holds that

$$
n-t \leq x+y+z \leq 2(n-t)
$$

- Byzantine nodes can change the tournament graph by only adding new edges which also satisfy the triangle inequality.

Based on the above conditions, the correct ranking should be $c_{1} \succ c_{2} \succ c_{3}$ and the total Kendall's distance is $2(n-t)-x-y+z$. In order to see which case leads us to the largest approximation ratio, we examine all other five possible rankings after considering Byzantine rankings. Since we exclude the binary case, we are able to reduce the five cases to only two which do not reduce to binary agreement: $c_{2} \succ c_{3} \succ c_{1}$ and $c_{3} \succ c_{1} \succ c_{2}$.

For the case $c_{2} \succ c_{3} \succ c_{1}$, the total Kendall's distance is $2(n-t)+x-y-z$. In order for this ranking to be chosen as the Kemeny median after adding Byzantine rankings, the inequality $x-z \leq t$ has to be satisfied. Finding the worst case for the correct nodes is equivalent to solving the optimization problem

$$
\begin{array}{cl}
\max & \frac{2(n-t)+x-y-z}{2(n-t)-x-y+z} \\
\text { s.t. } & x-z \leq t \\
& n-t \leq x+y+z \leq 2(n-t) \\
& n / 2 \leq z \leq y \leq x
\end{array}
$$

This optimization problem has a solution when $k=n / t \geq 4$.

- When $k \in[4,6]$, the maximum ratio is obtained by letting $x=n-2 t$ and $y=z=n / 2$. In this case, the maximum ratio is $2-4 / k$.
- When $k \in(6,8]$, the maximum ratio is obtained by letting $x=n / 2+t, y=n-3 t$ and $z=n / 2$. In this case, the maximum ratio is $1+2 / k$.
- When $k \in(8, \infty)$, the maximum ratio is obtained by letting $x=y=(2 n-t) / 3$ and $z=$ $(2 n-4 t) / 3$. In this case, the maximum ratio is $(2 k-1) /(2 k-4)$.

For the second case $c_{3} \succ c_{1} \succ c_{2}$, the total Kendall's distance is $2(n-t)-x+y-z$. The condition for this ranking to be the median ranking after adding Byzantine rankings is $y-z \leq t$. Finding the worst case is then equivalent to solving

$$
\begin{array}{cl}
\max & \frac{2(n-t)-x+y-z}{2(n-t)-x-y+z} \\
\text { s.t. } & y-z \leq t \\
& n-t \leq x+y+z \leq 2(n-t) \\
& n / 2 \leq z \leq y \leq x
\end{array}
$$

This optimization problem also only has a solution when $k=n / t \geq 4$.

- When $k \in[4,8]$, the maximum ratio is obtained by letting $x=y=3 n / 4-t, z=n / 2$. In this case, the maximum ratio is $3 / 2-2 / k$.
- When $k \in(8, \infty)$, the maximum ratio is obtained by letting $x=n-3 t, y=n / 2+t$ and $z=n / 2$. In this case, the maximum ratio is $1+2 / k$.


## D Proof of Theorem 6

In Lemma 3 and 4 we proved that the computed solution in Step 3 of Algorithm 2 satisfies the bounds from Section 5.2. In the following lemma we will show that the computed median also satisfies Pareto -Validity.

Lemma 7. Algorithm 圆 satisfies Pareto-Validity.
Proof. Assume, all nodes agree on a preference profile $c_{i} \succ c_{j}$. Then, there is a directed edge between $c_{i}$ and $c_{j}$ of weight $n-t$. Such an edge always belongs to a Kemeny median, since there cannot be any directed cycle formed by correct nodes only with weights $n-2 t>(n-t) / 2$ on all three majority edges due to the triangle inequality. This way, all correct nodes will have the preference $c_{i} \succ c_{j}$ in their median ranking $m_{v}$. Since Algorithm 1 satisfies Pareto-Validity, the consensus median ranking will also satisfy $c_{i} \succ c_{j}$.

It remains to show that Algorithm 2 terminates and has the same order of message complexity as Algorithm 1

Lemma 8. Algorithm terminates within $t+3$ rounds exchanging $O\left(n^{2} m \log m\right)$ messages.
Proof. Note that Algorithm 1 terminates after $t+1$ rounds. Algorithm 2 has two additional rounds in which the messages are exchanged. This way, our algorithm terminates within $t+3$ rounds. The message complexity of Algorithm $\square$ is $O\left(n^{2} m \log m\right)$, since in every round, each of the $n$ nodes sends a messages of size $m \log m$ to $n$ other nodes. In the additional two steps of Algorithm 2 $2 n^{2} m \log m$ messages are exchanged which gives the same message complexity $O\left(t n^{2} m \log m\right)$ for the second algorithm.

