The Communication Complexity of Instant-Runoff Voting

Élie de Panafieu[1] , **François Durand**[1] and **Jérôme Lang**[2]
[1]Nokia Bell Labs
[2]CNRS, LAMSADE, Université Paris-Dauphine, PSL, France
{francois.durand, elie.de_panafieu}@nokia-bell-labs.com, lang@lamsade.dauphine.fr

Abstract

The communication complexity of a voting rule is the worst-case number of bits that n voters must transmit to a central authority under the most efficient elicitation protocol in an election with m candidates. We study the communication complexity of Instant-Runoff Voting (IRV). Conitzer and Sandholm [2005] established an upper bound of $O(n(\log m)^2)$, but did not provide a matching lower bound beyond $\Omega(n \log m)$. We resolve this open problem by raising the lower bound to $\Omega(n(\log m)^2)$ using the fooling set technique, thereby showing that the communication complexity of IRV is $\Theta(n(\log m)^2)$. We further show that this complexity drops to $\Theta(n \log m)$ under the single-peakedness restriction, and that both the IRV-Average variant and Single Transferable Vote (STV), the multiwinner extension of IRV, have the same asymptotic communication complexity as IRV. (Extended version with appendices: https://shs.hal.science/hal-05566718. Short video: https://youtu.be/gTXV3R2DS6o.)

1 Introduction

Assessing the suitability of a voting rule requires analyzing its normative properties, its computational issues, and its *communication complexity*, while also taking into account the rule's understandability by laypeople. Among these criteria, communication complexity is especially relevant: if excessive interaction is required, voters may abandon the voting process altogether or, at best, submit their preferences hastily and without due care. In this respect, communication complexity is arguably more directly relevant to practical deployment than computational complexity: the cost of computation falls on a machine, whereas the cost of communication falls on voters.

A communication protocol for a voting rule is an interactive elicitation procedure in which, at each step, some voters are asked to transmit information about their vote, conditional on the responses obtained in previous steps, and that fully determines the outcome of the election. The complexity of a communication protocol is the total number of bits transmitted by voters in the worst-case execution of the protocol. The communication complexity of a voting rule is the minimum complexity over all its communication protocols.

In this paper, our main focus is on the communication complexity of *Instant-Runoff Voting* (IRV), a prominent single-winner voting rule used in many large-scale political elections. IRV proceeds by successive elimination of the candidate with the fewest top-ranked votes. The rule is sometimes also referred to as *Single Transferable Vote* (STV); following standard terminology, we use IRV for the single-winner rule and STV for its multiwinner variant.

The communication complexity of voting rules was first, and mostly, studied by Conitzer and Sandholm [2005]. In this framework, only the bits sent by voters are counted: messages sent by the central authority and received by voters are not taken into account, as the communication (or elicitation) burden falls on voters rather than on the central authority, which is typically a computer. Their upper bounds are established via explicit communication protocols, whereas their lower bounds rely on the classical *fooling set* technique [Kushilevitz and Nisan, 1997], which we recall later. Their results identify the communication complexity of most common voting rules, with a notable exception: IRV, for which they leave a gap between a lower bound $\Omega(n \log m)$ and an upper bound $O(n(\log m)^2)$. To date, this gap has not been filled yet.

Answering this open question is particularly relevant given the prevalence of IRV in real-world elections, both in large-scale political contexts and in lower-stakes settings. If the communication complexity of IRV is $\Theta(n(\log m)^2)$, then the known upper bound is asymptotically optimal, and efforts should focus on refining and implementing the existing protocol on interactive voting platforms. If, on the other hand, the communication complexity of IRV has a smaller order of magnitude than $n(\log m)^2$, this would imply the existence of a more efficient protocol yet to be discovered.

We do not resist the temptation to spoil the reader: we close this gap by proving that the communication complexity of IRV is indeed $\Theta(n(\log m)^2)$.

Section 2 provides background and related work. Section 3 introduces the framework. Section 4 presents our main result: the communication complexity of IRV is bounded below by $\Omega(n(\log m)^2)$. In Section 5 we show that the communication complexity of IRV drops to $\Theta(n \log m)$ under the restriction to single-peaked profiles, and that both the IRV-Average variant and the multiwinner rule STV have the same asymptotic

communication complexity as IRV. Section 6 concludes.

2 Background and Related Work

2.1 Instant-Runoff Voting (IRV)

IRV is used in a wide range of elections across several countries, including Australia, India, Ireland, the UK, and the US. It is one of the few commonly studied single-winner rules that are independent of clones [Tideman, 1987], and arguably the only one used in large-scale real-world elections. It is computationally hard to manipulate, even by a single voter [Bartholdi *et al.*, 1989]; the probability that a preference profile is manipulable is low in most settings [Durand, 2023; Durand, 2025]; and manipulation is cognitively demanding, which tends to encourage sincere voting [Van der Straeten *et al.*, 2010]. From a computational perspective, while IRV with immediate tie-breaking is easy to compute, its *parallel universe* version is NP-hard, though still tractable in practice [Csar *et al.*, 2017; Wang *et al.*, 2019]. IRV has also been the subject of numerous empirical studies: see [Institute for Mathematics and Democracy, 2026] for a continuously updated survey.

2.2 Communication Complexity of Voting Rules

The seminal paper by Conitzer and Sandholm [2005] identifies three broad classes of voting rules. *High-communication* rules, such as Borda, Copeland, and ranked pairs, require voters to transmit their full rankings in the worst case, yielding a communication complexity of $\Theta(nm \log m)$. At the other end of the spectrum, *low-communication* rules, such as plurality and plurality with runoff, require each voter to transmit not much more than a single candidate, leading to a communication complexity of $\Theta(n \log m)$. *Medium-communication* rules lie between these extremes; notable examples include voting trees and Bucklin, both of which have communication complexity $\Theta(nm)$. For IRV, the gap between $\Omega(n \log m)$ and $O(n(\log m)^2)$ prevents classification as a medium-communication rule, although it is widely believed to require substantially more communication than $n \log m$.

Service and Adams [2012] study the communication complexity of approximating voting rules. Mandal *et al.* [2019] investigate the trade-off between communication complexity and distortion. Chevaleyre *et al.* [2009], Xia and Conitzer [2010], and Karia and Lang [2024] focus on the one-round communication complexity of voting rules, known as *compilation complexity*.

2.3 Communication Issues for IRV

The raw definition of IRV assumes full rankings and therefore entails a high communication burden. Two natural escape routes are *simplified ballots* and *interactive protocols*.

The most common form of simplified ballots consists of *truncated ballots*: voters rank only their top k candidates, where k may depend on the voter. Smaller values of k make the input easier to elicit but provide weaker guarantees that the resulting winner coincides with the true IRV winner. Such truncated-ballot approximations of IRV and STV have been studied both theoretically and empirically [Burnett and Kogan, 2015; Ayadi *et al.*, 2019; Kilgour *et al.*, 2020; Hoffman *et al.*, 2021; Tomlinson *et al.*, 2023]. Other works allow voters to express indifference [Delemazure and Peters, 2024], or consider missing and truncated ballots [Han *et al.*, 2024]. All these works assume one-shot elicitation.

By contrast, to guarantee determination of the true IRV winner, one can rely on interactive elicitation protocols. Such protocols ask voters to provide partial preference information to the central authority as needed during execution. As mentioned above, Conitzer and Sandholm [2005] propose an IRV protocol with communication complexity $\mathcal{O}(n(\log m)^2)$. Ayadi *et al.* [2019] improve its practical performance while preserving the same worst-case complexity.

3 Framework

3.1 Rankings and Profiles

We consider a set of n voters $\mathcal{V} = \{0, \dots, n-1\}$ and a set of m candidates $\mathcal{C} = \{0, \dots, m-1\}$. A *ranking* r is a linear order over $\mathcal{C}$, for instance $(2 \succ 0 \succ 1)$. A *profile* P is a collection of n rankings, one for each voter. Given two profiles P and Q, and a voter $v \in \mathcal{V}$, we denote by $P^{-v} + Q^v$ the profile obtained from P by replacing the ranking of voter v in P with their ranking in Q.

3.2 Communication Complexity of Voting Rules

Following [Conitzer and Sandholm, 2005], a *voting rule* f maps a profile to a winning candidate. The *communication complexity* of f is the worst-case number of bits transmitted by the voters in the most efficient protocol implementing f. A *fooling set* $\mathcal{F}$ is a collection of profiles such that:

1. f has the same outcome w for all profiles $P \in \mathcal{F}$; and
2. for any two distinct profiles $P, Q \in \mathcal{F}$, there exists a profile T obtained by selecting, for each voter v, either P_v or Q_v, such that $f(T) \neq w$.

Its existence implies a communication complexity of at least $\log_2 |\mathcal{F}|$ for f [Kushilevitz and Nisan, 1997]. As observed by Conitzer and Sandholm [2005], it suffices to construct a fooling set for a representative set of pairs (n, m), namely for infinitely many values of m and, for each such m, infinitely many values of n. All these notions extend naturally to multiwinner voting rules, where f outputs a subset of candidates rather than a single winner.

3.3 Instant-Runoff Voting (IRV)

Under *Instant-Runoff Voting* (IRV), the outcome is determined through successive rounds of elimination. Initially, all candidates in $\mathcal{C}$ are *active*. In each round, each voter assigns one point to their most-preferred active candidate, and the candidate with the lowest score is eliminated from the active set. The procedure iterates until a single active candidate remains and is declared the winner. In practical implementations, it may terminate as soon as a candidate obtains a strict majority of the votes, that is, a score greater than $n/2$, since this candidate is then guaranteed to win.

Throughout the paper, we assume that ties are broken in every round in favor of candidates with smaller indices: if several candidates are tied for the lowest score, the one with

Protocol 1 PPR Protocol for IRV

$A \leftarrow \mathcal{C}$ ▷ active candidates
for $v \in \mathcal{V}$:
  $t(v) \leftarrow$ **elicit** the top candidate of v
while true:
  for $c \in A$:
    $V(c) \leftarrow \{v \in \mathcal{V} \mid t(v) = c\}$ ▷ voters supporting c
    score$(c) \leftarrow |V(c)|$
  if $\exists\, c \in A$ such that score$(c) > n/2$:
    return c ▷ c is the winner
  else:
    let c be a candidate in A minimizing score(c) (possibly using tie-breaking)
    remove c from A
    for $v \in V(c)$:
      $t(v) \leftarrow$ **elicit** the top candidate of v in A

the *largest* index is eliminated. Alternative tie-breaking rules are discussed in Section 4.5.

4 Communication Complexity of IRV

Section 4.1 recalls the communication protocol for IRV described by Conitzer and Sandholm [2005], yielding an upper bound of $\mathcal{O}(n(\log m)^2)$. To obtain a matching lower bound, Sections 4.2, 4.3, and 4.4 introduce a family of profiles $\mathcal{F}$, show that it forms a fooling set, and derive its asymptotic log-cardinality $\Theta(n(\log m)^2)$. Section 4.5 shows that this bound holds independently of the tie-breaking rule.

4.1 Progressive Preference Revelation Protocol

Conitzer and Sandholm [2005] establish the $\mathcal{O}(n(\log m)^2)$ upper bound on the communication complexity of IRV using Protocol 1, which we call the *Progressive Preference Revelation* (PPR) protocol. To make it more suitable for real-world usage, we add early termination as soon as a candidate obtains a strict majority of the votes. Beyond its interest for IRV, this protocol also serves as a basis for analyzing the alternative settings considered in Section 5.

4.2 Definition of the Fooling Set

To build intuition for our fooling set, we start from perfectly symmetric profiles, where each ranking is represented by the same number of voters. In such profiles, the elimination order is $(m-1, m-2, \ldots, 0)$, since ties are broken in favor of candidates with smaller indices, and the winner is candidate 0.

However, some rankings cannot be distinguished by the IRV counting process in such profiles. For instance, the rankings $(2 \succ 4 \succ 3 \succ 0 \succ 1)$ and $(2 \succ 0 \succ 1 \succ 3 \succ 4)$ both initially count for candidate 2 and are then transferred to candidate 0, where they remain until the end. We say that these two rankings have the same *signature* $(2, 0)$. If two profiles differ only by replacing rankings with others having the same signature, then mixing these profiles cannot alter the IRV outcome and thus cannot prevent candidate 0 from winning, contradicting the definition of a fooling set. Thus, all rankings with the same signature will be replaced by a single ranking, called its *representative*.

Candidates not seen by the IRV counting process, i.e., not appearing in the signature, are ranked so as to favor candidates with higher indices as much as possible. When mixing two distinct profiles from the fooling set, this choice will help a candidate other than 0 to win.

Signature. As mentioned above, for a ranking r, its *signature* $\sigma(r)$ is the list of candidates observed during the IRV counting process, assuming the elimination order $(m-1, m-2, \ldots, 0)$. For example, $\sigma\big((2 \succ 4 \succ 3 \succ 0 \succ 1)\big) = (2, 0)$. By convention, we include a final trivial round in which only candidate 0 remains and all voters vote for that candidate. For example, $\sigma\big((1 \succ 2 \succ 0)\big) = (1, 0)$, even if this ballot supports candidate 1 throughout the entire process in practice.

Formally, the signature $\sigma(r)$ is defined by the following algorithm. Start with the empty list. For each rank $i \in \{0, \ldots, m-1\}$, append the candidate r_i if and only if it has the smallest index encountered so far, that is, if $r_i = \min_{j \le i} r_j$. In combinatorics, this construction is known as the *record* of a permutation, except that records are conventionally defined with respect to a maximum rather than a minimum [Dumont and Kreweras, 1988]. The signature of a ranking is therefore always a decreasing list, starting with its top-ranked candidate and ending with candidate 0. The possible signatures are exactly the decreasing lists of candidates ending with 0.

Representative. Intuitively, for a signature s, we define its *representative* $\rho(s)$ as the ranking with signature s that favors candidates with higher indices as much as possible. For example, $\rho\big((2, 0)\big) = (2 \succ 4 \succ 3 \succ 0 \succ 1)$. Formally, $\rho(s)$ is obtained by scanning the signature s: for each element s_i, we append s_i, followed by all candidates with higher indices that have not yet been added, in decreasing order. By construction, for every signature s, we have $\sigma(\rho(s)) = s$.

Fooling set. Assume $n = \ell m!$, where ℓ is a positive integer. Let $\mathcal{P}^{(\ell)}$ denote the set of all profiles in which each ranking appears exactly ℓ times. We define our fooling set $\mathcal{F}$ as

$$\mathcal{F} = \big\{\, \rho(\sigma(P_{\text{ini}})) \mid P_{\text{ini}} \in \mathcal{P}^{(\ell)} \,\big\},$$

where $\rho(\sigma(P_{\text{ini}}))$ denotes the profile obtained from P_{ini} by replacing each ranking r with $\rho(\sigma(r))$, the representative of its signature.

An equivalent definition is as follows. Let $R(s) = |\sigma^{-1}(s)|$ denote the number of rankings with signature s. The fooling set $\mathcal{F}$ then consists of all profiles that contain, for each signature s, exactly $\ell R(s)$ copies of $\rho(s)$.

For example, for $m = 3$, consider this profile $P_{\text{ini}} \in \mathcal{P}^{(1)}$:

0	**0**	**1**	**1**	**2**	**2**
1	2	**0**	2	**0**	**1**
2	1	2	**0**	1	**0**

Each column represents a voter's ranking; for instance, the leftmost voter has ranking $(0 \succ 1 \succ 2)$. Candidates belonging to the signatures are highlighted in bold. Note that the first two rankings have the same signature, as do the next two.

By taking the representative of each signature, we obtain

the profile $P = \rho(\sigma(P_{\text{ini}}))$, which belongs to $\mathcal{F}$:

0	0	1	1	2	2
2	2	2	2	0	1
1	1	0	0	1	0

The other profiles in $\mathcal{F}$ are the reorderings of P, such as the following profile Q, where the reordered columns are highlighted.

0	0	1	2	1	2
2	2	2	0	2	1
1	1	0	1	0	0

In both profiles P and Q, the IRV counting process is the same as in P_{ini}, and therefore candidate 0 wins. However, consider the profile $P^{-3} + Q^3$, obtained by borrowing the *fourth* voter from Q (of index 3, as indices start at 0):

0	0	1	2	2	2
2	2	2	0	0	1
1	1	0	1	1	0

Candidate 1 is eliminated in the first round; the corresponding ballot then transfers to candidate 2, who is eventually elected. In particular, the winner is not candidate 0. To prove that $\mathcal{F}$ is a fooling set, it remains to generalize this observation.

4.3 Proof of the Fooling Set Property

We now prove that $\mathcal{F}$ is a fooling set for IRV.

In any profile $P_{\text{ini}} \in \mathcal{P}^{(\ell)}$, the elimination order is $(m-1, m-2, \dots, 0)$. Hence, in the profile $P = \rho(\sigma(P_{\text{ini}}))$, the IRV counting process is identical, and candidate 0 is elected.

Let P and Q be two distinct profiles of $\mathcal{F}$. Then there exists a voter v whose ranking is r in P and a different ranking r' in Q.

Different tops. Suppose that r and r' differ in their top-ranked candidate, say c for r and d for r', and assume without loss of generality that $c < d$. Consider the profile $P^{-v} + Q^v$. In this profile, candidate c receives one fewer top vote than average, while d receives one more. In the initial round $t = 0$, candidate c is therefore eliminated. By the definition of the representative ρ, all corresponding vote transfers go to candidate $m-1$. At each subsequent round $t > 0$, candidate $m-1$ holds $\frac{(t+1)n}{m}$ votes, while every other remaining candidate has $\frac{n}{m}$ votes, up to a discrepancy of one vote due to voter v. One of these candidates is then eliminated, and again all transfers go to candidate $m-1$. Eventually, candidate $m-1$ is elected.

General case. The two rankings r and r' necessarily have distinct signatures: by construction of $\mathcal{F}$, if they had the same signature s, they would both be equal to the ranking $\rho(s)$. The signatures $\sigma(r)$ and $\sigma(r')$ share a common prefix and then differ at some position:

$$\sigma(r) = (s_0, \dots, s_j, c, \dots), \qquad \sigma(r') = (s_0, \dots, s_j, d, \dots),$$

where, without loss of generality, we assume $c < d$. Note that $s_j \neq 1$, otherwise we would have $c = d = 0$ by definition of a signature. Consider the profile $P^{-v} + Q^v$. Until the elimination of s_j, the counting process proceeds as usual, eliminating candidates $m-1, m-2, \dots, s_j$ in this order. At that point, we are back in the case where r and r' differ in their top-ranked candidate, and candidate $s_j - 1$ is eventually elected. Since $s_j \neq 1$, candidate 0 is not the winner, which establishes that $\mathcal{F}$ is a fooling set.

4.4 Cardinality of the Fooling Set

Rankings with a given signature. As shown by Wilf [1995, item (IV)], the number of rankings with a given signature s is

$$R(s) = \prod_{c \in \mathcal{C} \setminus s} c, \tag{1}$$

where $\mathcal{C} \setminus s$ denotes the set of candidates that do not appear in the signature. The simplicity of this formula follows directly from our choice to label the candidates as $\{0, 1, \dots, m-1\}$, which actually motivated this convention.

Exact cardinality of $\mathcal{F}$. As noted in Section 4.2, the profiles in $\mathcal{F}$ are exactly the reorderings of any one of them. If all voters had distinct rankings, this would yield $n!$ distinct profiles. However, for each signature s, there are $\ell R(s)$ voters with identical rankings, and we must therefore divide by $(\ell R(s))!$ to avoid multiple counting. Using Equation (1) and the change of variables $S = \mathcal{C} \setminus s$, we obtain the following expression for the cardinality of the fooling set:

$$|\mathcal{F}| = \frac{n!}{\prod_{S \subseteq [1,m-1]} (\ell \prod_{c \in S} c)!}, \tag{2}$$

where the empty product corresponding to $S = \emptyset$ evaluates to 1 by convention.

Asymptotics. We give here the main steps of the proof; full details are provided in the technical appendix [de Panafieu *et al.*, 2026]. To approximate Equation (2), we consider $n = \ell m!$, with ℓ, m, or both tending to infinity. We take the logarithm of the exact expression (2) and apply Stirling's formula $\log(a!) = a(\log a - 1) + \mathcal{O}(\log a)$ twice, which yields

$$\begin{aligned}\log|\mathcal{F}| = {} & n\big(\log n - 1\big) + \mathcal{O}(\log n) \\ & - \sum_{S \subseteq [1,m-1]} \ell\Big(\prod_{c \in S} c\Big)\Big(\log \ell - 1 + \sum_{c \in S} \log c\Big) \\ & + \mathcal{O}\Big(\sum_{S \subseteq [1,m-1]} \log\Big(\ell \prod_{c \in S} c\Big)\Big).\end{aligned}$$

Interchanging the sums over $S \subseteq [1, m-1]$ and $c \in S$, and applying twice the identity $\sum_{T \subseteq \mathcal{B}} \prod_{c \in T} c = \prod_{c \in \mathcal{B}} (1 + c)$, we obtain

$$\begin{aligned}\log|\mathcal{F}| = {} & n\big(\log n - 1\big) - \ell m!\big(\log \ell - 1\big) \\ & - \ell m! \sum_{j=1}^{m-1} \frac{j \log j}{j+1} + \mathcal{O}\Big(\sum_{S \subseteq [1,m-1]} \log(\ell m!)\Big)\end{aligned}$$

which reduces, after substituting $n = \ell m!$, to

$$\log|\mathcal{F}| = n \sum_{j=1}^{m-1} \frac{\log j}{j+1} + \mathcal{O}(2^m \log n).$$

The sum is estimated using a Riemann sum-integral comparison, and the error term is found negligible for $n = \ell m!$, which finally yields

$$\log |\mathcal{F}| \sim \frac{n(\log m)^2}{2}.$$

This establishes the desired lower bound on the worst-case communication complexity of any protocol for IRV. Combining this result with the upper bound $\mathcal{O}(n(\log m)^2)$ provided by Conitzer and Sandholm [2005] and based on the PPR protocol (Protocol 1), we deduce the following theorem.

Theorem 1. *The communication complexity of IRV is* $\Theta\left(n(\log m)^2\right)$.

4.5 Discussion of the Tie-Breaking Rule

The fooling set studied in Sections 4.2, 4.3, and 4.4 relies on the assumption that ties are broken in favor of candidates with lower indices. However, the construction can be adapted to eliminate all ties, showing that the resulting lower bound remains valid under any tie-breaking rule.

Indeed, we can append at the end of each profile a small collection of additional voters that simulates tie-breaking in favor of lower indices: one voter whose top choice is candidate $m-1$, two voters whose top choice is candidate $m-2$, ...; and m voters whose top choice is candidate 0. Their rankings are completed by placing the remaining candidates in ascending order of their indices; for instance, the voter whose top choice is $m-1$ has ranking $(m-1 \succ 0 \succ 1 \succ \cdots \succ m-2)$. Altogether, this yields a number $m(m+1)/2$ of *tie-breaking voters*. A closely related idea already appears implicitly in the proofs of Conitzer and Sandholm [2005], although it is not stated explicitly in these terms.

Now assume $\ell > m(m+1)/2$. The outcome of IRV then depends only on the remaining voters, which we call the *significant voters*, except when a tie arises among them, in which case the tie-breaking voters come into play.

Since the number of voters added in this way is negligible, the previous asymptotic estimate of the size of the fooling set remains unchanged, and consequently so does the lower bound on the communication complexity of any protocol for IRV.

What does change with the tie-breaking rule is the upper bound: if the tie-breaking rule requires information that is not revealed during the PPR protocol, then the worst-case communication complexity may be larger.

5 Related Settings

In this section, we consider three variants of our main setting: IRV under single-peaked preferences (Section 5.1), the IRV-Average voting rule (Section 5.2), and STV, the multiwinner version of IRV (Section 5.3).

5.1 IRV in Single-Peaked Profiles

In this section, we add the classical assumption that voters' preferences are *single-peaked* with respect to the reference axis $(0, \ldots, m-1)$. This means that for any voter and any triple of candidates $c_1 < c_2 < c_3$, candidate c_2 cannot be the least preferred among the three. Equivalently, for any voter and any strict non-empty prefix $(c_1 \succ \ldots \succ c_{j-1})$ of her ranking, the next-ranked candidate c_j must have either the immediately smaller or the immediately larger index among those not yet encountered. We further assume that the reference axis is known in advance and can be exploited by the protocol.

Upper bound. In this setting, the PPR protocol becomes significantly cheaper. The first step is unchanged, since each voter is asked to elicit her top-ranked candidate. At each subsequent step, however, it suffices to ask the relevant voters whether their next preferred candidate, among the non-eliminated ones, lies to the left or to the right of the candidate that has just been eliminated. This requires only a single bit of communication, instead of $\log_2 m$ bits. We give an example of execution in the technical appendix [de Panafieu *et al.*, 2026]. When j candidates remain, the candidate eliminated in that round has at most n/j top supporters, since the minimum score is never larger than the average. Therefore, the worst-case communication complexity of the protocol is

$$n \log_2(m) + \frac{n}{m} + \frac{n}{m-1} + \cdots + \frac{n}{3} = \mathcal{O}(n \log m).$$

Definition of the fooling set. We now construct a fooling set that yields a matching lower bound on the communication complexity of IRV in the single-peaked setting.

We may assume without loss of generality that m is a power of two: it suffices to establish the lower bound for an infinite subset of possible values of m. We also assume that $n = \ell m$ for some integer $\ell > 2$; the rationale for this choice will become clear shortly.

Our fooling set $\mathcal{F}$ consists of all profiles with exactly ℓ voters for each ranking of the form

$$(c \succ c-1 \succ \cdots \succ 0 \succ c+1 \succ c+2 \succ \cdots \succ m-1).$$

In other words, there are ℓ voters for each possible peak position c, and all voters complete their rankings by successively filling the remaining positions with all candidates to the left of the peak, followed by all candidates to the right.

Proof of the fooling set property. We write $c \to d$ as shorthand for the fact that candidate c is eliminated and all votes cast for c are transferred to d. In any profile of $\mathcal{F}$, under our usual tie-breaking rule favoring smaller indices, the elimination process unfolds as follows. First, we have $m-1 \to m-2$, $m-3 \to m-4$, ..., $1 \to 0$, that is, odd-indexed candidates transfer to the even candidate immediately to their left. Next, $m-2 \to m-4, \ldots, 2 \to 0$, meaning that candidates congruent to $2 \bmod 4$ transfer to candidates that are multiples of 4. And so on. Eventually, we obtain $m/2 \to 0$, and candidate 0 is declared the winner.

Now let P and Q be two distinct profiles in $\mathcal{F}$. There exists a voter v whose ranking is r in P and a different ranking r' in Q. Let c and d denote the respective peak candidates of r and r'. Since profiles in $\mathcal{F}$ contain only one type of ranking for each peak position, we necessarily have $c \neq d$. Without loss of generality, assume that $c < d$, that is, c has a smaller index than d.

If c is odd, consider the profile $Q^{-v} + P^v$, illustrated in Table 1. Initially, candidate c receives $\ell + 1$ votes, candidate d receives $\ell - 1$ votes, and every other candidate receives exactly ℓ votes. First, candidate d, and possibly some

$\cdots$	ℓ	ℓ	$\ell+1$	ℓ	$\cdots$	$\ell-1$	$\cdots$
$\cdots$	$c-2$	$c-1$	c	$c+1$	$\cdots$	d	$\cdots$
	$c-3$	$c-2$	$c-1$	c		$d-1$	
	$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\vdots$	

Table 1: Excerpt of the profile $Q^{-v} + P^v$ constructed for the proof of the fooling set property in the single-peaked setting, when the peak candidate c of voter v is odd. In each column, the header gives the number of voters, and the entries below specify their common ranking. After some eliminations of candidates with indices greater than c, we have $c-1 \to c-2$, then $c-3 \to c-4$, and so on, until $2 \to 1$, then candidate 0 is eliminated.

candidates to the right of c with ℓ votes, are eliminated; the corresponding votes are transferred to candidates of index at least c. Once all remaining candidates to the right of c have more than ℓ votes, eliminations then start to occur on the left of c: $c-1 \to c-2, \ldots, 2 \to 1$, and finally candidate 0 is eliminated. The winner is therefore not 0.

If c is even, consider the profile $P^{-v} + Q^v$. In the initial round, candidate c receives $\ell - 1$ votes, candidate d receives $\ell+1$ votes, and every other candidate receives exactly ℓ votes. The first candidate eliminated is therefore c. If $c = 0$, we are done. Otherwise, we have $c \to c-1$. Candidate $c-1$ now has $2\ell - 1$ votes, which is strictly greater than $\ell + 1$ since we assumed $\ell > 2$. From this point on, the remainder of the argument is analogous to the case where c is odd: after some eliminations of candidates to the right of $c-1$, eliminations occur on the left, yielding $c-2 \to c-3, \ldots, 2 \to 1$, and finally candidate 0 is eliminated. This completes the proof that $\mathcal{F}$ is a fooling set.

Cardinality of $\mathcal{F}$. The fooling set consists of all possible permutations of the n voters, where the voters are partitioned into m classes of ℓ voters with identical rankings. Consequently,

$$|\mathcal{F}| = \frac{n!}{(\ell!)^m}.$$

Applying Stirling's approximation, straightforward calculations detailed in the technical appendix [de Panafieu *et al.*, 2026] yield

$$\log |\mathcal{F}| \sim n \log m.$$

This establishes the desired lower bound on the worst-case communication complexity of any protocol for IRV in the single-peaked setting. We can therefore state the following theorem.

Theorem 2. *The communication complexity of IRV in the single-peaked setting is* $\Theta(n \log m)$.

In the technical appendix [de Panafieu *et al.*, 2026], we show that, as in Section 4.5, the lower bound remains valid for any tie-breaking rule, using the same tie-breaking voters technique.

5.2 IRV-Average

We now examine *IRV-Average*, a rule similar to IRV in which, at each round, multiple candidates may be eliminated simultaneously, depending on how their scores compare to the average (see, e.g., Durand [2023]). We consider two variants. In the *strict* version, all candidates whose number of votes is strictly below the average are eliminated. In the *weak* version, candidates with a score lower than or equal to the average are eliminated. In both cases, an exception rule must be specified to handle perfect ties and to prevent the elimination of either none or all candidates.

This rule is similar to the Nanson and Kim–Roush rules, which operate analogously using, respectively, the Borda and Veto scores [Niou, 1987; Kim and Roush, 1996].

A straightforward adaptation of the PPR protocol applies to IRV-Average and achieves the same worst-case communication complexity, provided that the tie-handling exception requires no additional information. For example, perfect ties may be resolved by immediately declaring the smallest-index candidate the winner, or by eliminating the largest-index candidate.

Moreover, the fooling-set lower bound established for IRV extends to IRV-Average. To cover both the strict and weak variants, as well as all possible exception rules, it suffices to use the fooling set with tie-breaking voters introduced in Section 4.5. We thus obtain the following theorem, analogous to Theorem 1 for IRV.

Theorem 3. *The communication complexity of IRV-Average is* $\Theta\left(n(\log m)^2\right)$.

5.3 STV

Single Transferable Vote (STV) is the multiwinner analogue of IRV. It satisfies several desirable normative properties and is used in political elections in a number of jurisdictions [Tideman, 1995; Gallagher and Mitchell, 2005].

Let k denote the number of candidates to be elected. STV is defined sequentially, similarly to IRV, with the difference that at each step either a candidate is elected or a candidate is eliminated. Initially, each voter has a voting weight equal to 1.

A candidate is elected at a given step if the total weight of votes supporting it reaches a threshold called the *quota* Q, which remains constant throughout the process. There are several classical ways to define the quota as a function of n and k; the most common is the *Droop quota*, $Q = \lfloor \frac{n}{k+1} \rfloor + 1$. Once a candidate is elected, the voters supporting it see their voting weights reduced and their votes are transferred to their next preferred candidate. Again, there are various ways to update the weights, either deterministically or stochastically. The choice of the quota rule and of the weight-update mechanism does not affect the communication complexity of the rule.

When no candidate is elected at a given step, the candidate with the least support is eliminated, and all votes supporting it are transferred to their next preferred candidate.

The process continues until either k candidates have been elected or $n-k$ candidates have been eliminated; in the latter case, the remaining active candidates become elected.

Upper bound. The STV-PPR protocol (Protocol 2) generalizes the PPR Protocol 1 to an arbitrary number $k \geq 1$

Protocol 2 STV-PPR Protocol

$A \leftarrow \mathcal{C}$ ▷ active candidates
$E \leftarrow \emptyset$ ▷ elected candidates
for $v \in \mathcal{V}$:
  $t(v) \leftarrow$ **elicit** the top candidate of v
  $w(v) \leftarrow 1$ ▷ weight of voter v
while true:
  for $c \in A$:
    $V(c) \leftarrow \{v \in \mathcal{V} \mid t(v) = c\}$ ▷ voters supporting c
    $\text{score}(c) \leftarrow \sum_{v \in V(c)} w(v)$
  let c be a candidate in A maximizing $\text{score}(c)$
  if $\text{score}(c) \geq Q$:
    add c to E and remove it from A
    if $|E| = k$:
      return E
    update the weights $w(v)$ of voters in $V(c)$
      ▷ several variants exist
    remove $\{v \in V(c) \mid w(v) = 0\}$ from $\mathcal{V}$ and $V(c)$
  else:
    let c be a candidate in A minimizing $\text{score}(c)$
    remove c from A
    if $|E| + |A| = k$:
      return $E \cup A$
  for $v \in V(c)$:
    $t(v) \leftarrow$ **elicit** the top candidate of v in A

of winners. Its worst-case communication complexity is attained on profiles in which each possible ranking appears the same number of times. In such profiles, at each step all remaining candidates are tied, and ties are broken, as usual, in favor of smaller indices. Candidates $m-1, \dots, k$ are therefore eliminated one by one; subsequently, all candidates $k-1, \dots, 0$ are elected. Provided that k is constant, the resulting communication complexity of the protocol is $\mathcal{O}\left(n(\log m)^2\right)$.

STV-signature. We replace the notion of IRV-signature $\sigma(r)$ of a ranking r, previously simply called a *signature*, with the *STV-signature* $\tilde{\sigma}(r)$, defined as follows. For each rank $i \in \{0, \dots, m-1\}$, append the candidate r_i if and only if it is the smallest candidate index encountered so far, that is, if $r_i = \min_{j \leq i} r_j$. Stop as soon as an index belonging to $\{0, \dots, k-1\}$ has been added. For $k = 4$, an example of an STV-signature is

$$\tilde{\sigma}(\mathbf{7} \succ 9 \succ \mathbf{3} \succ 4 \succ 5 \succ 0 \succ 1 \succ 6 \succ 2 \succ 8) = (7, 3).$$

Note that, unlike in the IRV case, candidate 0 does not necessarily appear in the signature. The STV-signature $\tilde{\sigma}(r)$ thus corresponds to the list of candidates encountered in ranking r by the STV-PPR protocol when assuming the elimination order $(m-1, \dots, k)$.

STV-representative. The notion of a representative for a signature is adapted similarly. Among the rankings compatible with a given STV-signature, we select the one that favors candidates with higher indices as much as possible. We denote the STV-representative of an STV-signature s by $\tilde{\rho}(s)$. For example, with $m = 10$ and $k = 4$,

$$\tilde{\rho}\big((7, 3)\big) = \mathbf{7} \succ 9 \succ \mathbf{3} \succ 8 \succ 6 \succ 5 \succ 4 \succ 2 \succ 1 \succ 0.$$

Fooling set. Let $\mathcal{P}^{(\ell)}$ again denote the family of all profiles over m candidates in which each ranking appears exactly ℓ times, so that the total number of voters is $n = \ell m!$. We define

$$\mathcal{F} = \left\{ \tilde{\rho}(\tilde{\sigma}(P_{\text{ini}})) \mid P_{\text{ini}} \in \mathcal{P}^{(\ell)} \right\},$$

where $\tilde{\rho}(\tilde{\sigma}(P_{\text{ini}}))$ is the profile obtained from P_{ini} by replacing each ranking r with $\tilde{\rho}(\tilde{\sigma}(r))$, the STV-representative of its STV-signature. By an argument analogous to the one used for IRV, $\mathcal{F}$ is a fooling set for STV.

Size of the fooling set. An STV-signature is a decreasing sequence of elements from $\{0, \dots, m-1\}$ that contains exactly one element in $\{0, \dots, k-1\}$, namely its last element. For any STV-signature $(s_0, \dots, s_{j-1})$, there is a one-to-one correspondence between rankings having this STV-signature and rankings having IRV-signature $(s_0, \dots, s_{j-2}, 0)$: it suffices to interchange candidates s_{j-1} and 0. Recall that $R(s)$ denotes the number of rankings with IRV-signature s, and let $\tilde{R}(s)$ denote the number of rankings with STV-signature s. Using (1), we obtain

$$\tilde{R}((s_0, \dots, s_{j-1})) = \prod_{c \in [1, m-1] \setminus \{s_0, \dots, s_{j-2}\}} c.$$

The size of the fooling set for STV is therefore

$$|\mathcal{F}| = \frac{n!}{\left(\prod_{S \subseteq [k, m-1]} \left(\ell \prod_{c \in [1, m-1] \setminus S} c\right)!\right)^k} \tag{3}$$

Asymptotic communication complexity of STV. A similar analysis to the one carried out for IRV, detailed in the technical appendix [de Panafieu *et al.*, 2026], shows that, for fixed k,

$$\log |\mathcal{F}| = \Theta\left(n(\log m)^2\right).$$

Since this lower bound matches the upper bound derived from the analysis of the STV-PPR protocol, we obtain the following theorem, which is not affected by the choice of the quota rule and of the weight-update mechanism.

Theorem 4. *The communication complexity of STV is* $\Theta\left(n(\log m)^2\right)$.

6 Discussion

Our main conclusion is that the communication complexity of IRV is $\Theta(n(\log m)^2)$, which implies that the PPR protocol is asymptotically optimal and confirms that STV is, as expected, a medium-communication rule. Since we now know that no significantly better protocol can exist asymptotically, it is timely to implement this protocol in interactive voting platforms. Of course, using the PPR protocol in practice requires voters to remain online long enough, or to reconnect when prompted, *e.g.* via email. Although such assumptions remain unrealistic for large-scale political elections, they may be reasonable in low-stake or small-scale settings.

This conclusion also applies to STV and IRV-Average, but not to IRV under single-peaked preferences, for which the protocol simplifies.

Since the lower bound is only attained asymptotically, it makes sense to study the *exact* communication complexity of IRV and STV for small values of m and/or n, in order to design finely optimized protocols for practical voting platforms.

Acknowledgments

We thank Emma Caizergues for fruitful discussions about the fooling set used for IRV. This work has been supported in part by a grant from the French State managed by the National Research Agency (ANR) through the France 2030 program, with the reference ANR-23-IACL-0008 (PR[AI]RIE-PSAI), and through the CONDORCET project with reference ANR-24-EXMA-0001.

References

[Ayadi *et al.*, 2019] Manel Ayadi, Nahla Ben Amor, Jérôme Lang, and Dominik Peters. Single transferable vote: Incomplete knowledge and communication issues. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '19*, pages 1288–1296, 2019.

[Bartholdi *et al.*, 1989] John J. Bartholdi, Craig A. Tovey, and Michael A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[Burnett and Kogan, 2015] Craig M. Burnett and Vladimir Kogan. Ballot (and voter) 'exhaustion' under instant runoff voting: An examination of four ranked-choice elections. *Electoral Studies*, 37:41–49, 2015.

[Chevaleyre *et al.*, 2009] Yann Chevaleyre, Jérôme Lang, Nicolas Maudet, and Guillaume Ravilly-Abadie. Compiling the votes of a subelectorate. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI '09*, pages 97–102, 2009.

[Conitzer and Sandholm, 2005] Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Electronic Commerce, EC '05*, pages 78–87, 2005.

[Csar *et al.*, 2017] Theresa Csar, Martin Lackner, Reinhard Pichler, and Emanuel Sallinger. Winner determination in huge elections with mapreduce. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence, AAAI '17*, pages 451–458, 2017.

[de Panafieu *et al.*, 2026] Élie de Panafieu, François Durand, and Jérôme Lang. The communication complexity of instant-runoff voting, 2026. Extended version with appendices, https://shs.hal.science/hal-05566718.

[Delemazure and Peters, 2024] Théo Delemazure and Dominik Peters. Generalizing instant runoff voting to allow indifferences. In *Proceedings of the 25th ACM Conference on Economics and Computation, EC '24*, page 50, 2024.

[Dumont and Kreweras, 1988] Dominique Dumont and Germain Kreweras. Sur le développement d'une fraction continue liée à la série hypergéométrique et son interprétation en termes de records et anti-records dans les permutations. *European Journal of Combinatorics*, 9(1):27–32, 1988.

[Durand, 2023] François Durand. Coalitional manipulation of voting rules: Simulations on empirical data. *Constitutional Political Economy*, 34(3):390–409, 2023.

[Durand, 2025] François Durand. Why instant-runoff voting is so resilient to coalitional manipulation: Phase transitions in the perturbed culture. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '25*, pages 658–666, 2025.

[Gallagher and Mitchell, 2005] Michael Gallagher and Paul Mitchell. *The politics of electoral systems*. OUP Oxford, 2005.

[Han *et al.*, 2024] Qishen Han, Amelie Marian, and Lirong Xia. Determining winners in elections with absent votes. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence, IJCAI '24*, pages 2816–2824, 2024.

[Hoffman *et al.*, 2021] Christina Hoffman, Jakini Kauba, Julie Reidy, and Thomas Weighill. Proportionality in multi-winner RCV elections: A simulation study with ballot truncation. Technical report, SSRN: https://ssrn.com/abstract=3942892, 2021. Accessed: 2026-05-06.

[Institute for Mathematics and Democracy, 2026] Institute for Mathematics and Democracy. Empirical analysis of ranked choice methods. https://mathematics-democracy-institute.org/empirical-analysis-of-ranked-choice-voting-methods/, 2026. Accessed May 15, 2026.

[Karia and Lang, 2024] Neel Karia and Jérôme Lang. Compiling the votes of a subelectorate for multi-winner voting rules. In *International Conference on Algorithmic Decision Theory, ADT '24*, pages 18–32, 2024.

[Kilgour *et al.*, 2020] D. Marc Kilgour, Jean-Charles Grégoire, and Angèle M. Foley. The prevalence and consequences of ballot truncation in ranked-choice elections. *Public Choice*, 184(1):197–218, 2020.

[Kim and Roush, 1996] Ki H. Kim and Fred W. Roush. Statistical manipulability of social choice functions. *Group Decision and Negotiation*, 5(3):263–282, 1996.

[Kushilevitz and Nisan, 1997] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.

[Mandal *et al.*, 2019] Debmalya Mandal, Ariel D. Procaccia, Nisarg Shah, and David P. Woodruff. Efficient and thrifty voting by any means necessary. In *Annual Conference on Neural Information Processing Systems, NeurIPS '19*, pages 7178–7189, 2019.

[Niou, 1987] Emerson Niou. A note on Nanson's rule. *Public Choice*, 54(2):191–193, 1987.

[Service and Adams, 2012] Travis C. Service and Julie A. Adams. Communication complexity of approximating voting rules. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS '12*, pages 593–602, 2012.

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

[Tideman, 1995] Nicolaus Tideman. The single transferable vote. *Journal of Economic Perspectives*, 9(1):27–38, 1995.

[Tomlinson *et al.*, 2023] Kiran Tomlinson, Johan Ugander, and Jon M. Kleinberg. Ballot length in instant runoff voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence, AAAI '23*, pages 5841–5849, 2023.

[Van der Straeten *et al.*, 2010] Karine Van der Straeten, Jean-François Laslier, Nicolas Sauger, and André Blais. Strategic, sincere, and heuristic voting under four election rules: an experimental study. *Social Choice and Welfare*, 35(3):435–472, 2010.

[Wang *et al.*, 2019] Jun Wang, Sujoy Sikdar, Tyler Shepherd, Zhibing Zhao, Chunheng Jiang, and Lirong Xia. Practical algorithms for multi-stage voting rules with parallel universes tiebreaking. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence, AAAI '19*, pages 2189–2196, 2019.

[Wilf, 1995] Herbert S. Wilf. On the outstanding elements of permutations. *preprint*, 1995.

[Xia and Conitzer, 2010] Lirong Xia and Vincent Conitzer. Compilation complexity of common voting rules. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI '10*, pages 915–920, 2010.