

Determining Majority in Networks with Local Interactions and very Small Local Memory

George B. Mertzios · Sotiris E. Nikolettseas ·
Christoforos L. Raptopoulos · Paul G. Spirakis

Received: date / Accepted: date

Abstract We study the problem of determining the majority type in an arbitrary connected network, each vertex of which has initially two possible types. The vertices may later change into other types, out of a set of a few additional possible types, and can interact in pairs only if they share an edge. Any (population) protocol is required to stabilize in the initial majority. First

A preliminary conference version of this work appeared in the *41st International Colloquium on Automata, Languages, and Programming (ICALP)*, Copenhagen, Denmark, 2014 [18]. George B. Mertzios was partially supported by the EPSRC Grant EP/K022660/1. Sotiris E. Nikolettseas was partially supported by the MULTIPLEX project – 317532. Christoforos L. Raptopoulos was partially supported by the SHARPEN project – PE6 (1081). Paul G. Spirakis was partially supported by the MULTIPLEX project – 317532, the EU ERC Project ALGAME and the EEE/CS School of the University of Liverpool.

George B. Mertzios
School of Engineering and Computing Sciences, Durham University, UK
E-mail: george.mertzios@durham.ac.uk

Sotiris E. Nikolettseas
Computer Technology Institute and Press “Diophantus”, Greece and Computer Science Department, University of Patras, Greece
E-mail: nikole@cti.gr

Christoforos L. Raptopoulos (corresponding author)
Computer Technology Institute and Press “Diophantus”, Greece and Computer Science Department, University of Geneva, Switzerland
Tel.: +30 2610 996966
Fax.: +30 2610 960490
E-mail: raptopox@ceid.upatras.gr

Paul G. Spirakis
Computer Technology Institute and Press “Diophantus”, Greece, Computer Science Department, University of Patras, Greece and Department of Computer Science, University of Liverpool, UK
E-mail: p.spirakis@liverpool.ac.uk

we prove that there does not exist any population protocol that always computes majority in any interaction graph by using at most 3 types per vertex. However this does not rule out the existence of a protocol with 3 types per vertex that is correct with high probability (whp). To this end, we examine an elegant and very natural majority protocol with 3 types per vertex, introduced in [4], whose performance has been analyzed for the clique graph. In particular, we study the performance of this protocol in arbitrary networks, under the probabilistic scheduler. We prove that, if the initial assignment of types to vertices is random, the protocol of [4] converges to the initial majority with probability higher than the probability of converging to the initial minority. In contrast, we show that the resistance of the protocol to failure when the underlying graph is a clique causes the failure of the protocol in general graphs.

Keywords majority in networks · population protocol · probabilistic scheduler · coupling

1 Introduction

One of the most natural computational problems in many physical systems is to compute the *majority*, i.e. to determine accurately which type of an element of the system appears more frequently. For instance, the majority problem is encountered in various settings such as in voting [12, 14], in epidemiology and interacting particles systems [17], in diagnosis of multiprocessor systems [25], in social networks [19, 21] etc. In distributed computing, the majority problem is an important and natural special case of the central problem of reaching *consensus* within a system [9, 16], where a number of processes have to agree on any single data value (e.g. leader election [11]). In all these physical systems, some

pairs of elements may interact with each other while other pairs may not be able to interact directly. This structure of the possible pairwise interactions between elements of the system can be modeled by a network (i.e. graph), where elements and possible interactions are represented by vertices and edges, respectively.

In order to solve the majority computation problem in a network, we first need to make some assumptions on the underlying model of computation. Much research has been done under the assumption that there exists a central authority, as well as unlimited available memory and full information about the whole network (see e.g. [8, 27]). However, in many real systems we do not have (or we do not wish to have) such a powerful computational model. The weaker the considered model of computation is (e.g. lack of central authority, partial or no information about the system, lack of memory etc.), the more challenging the majority computation becomes.

One of the ways to study distributed systems where agents may interact in pairs and each individual agent is extremely limited (in fact, being able to change into only a finite number of possible types) is by using *population protocols* [2, 5]. Then the complex behavior of the system emerges from the rules governing the possible pairwise interactions of the agents. Population protocols have been defined by analogy to population processes [15] in probability theory and have already been used in various fields, such as in statistical physics, genetics, epidemiology, chemistry and biology [7].

In particular, population protocols are *scalable*, i.e. they work independently of the size n of the underlying network (called the *interaction graph*) and the value of n is not even known to the protocol. Furthermore they are *anonymous*, i.e. there is only one transition function which is common to all entities/agents: the result of an interaction of an agent u of type q_u with an agent v of type q_v is the same regardless of the identity of u and v . The transition function of a population protocol only specifies the result of every possible interaction, without specifying *which* pairs of agents interact or *when* they are chosen to interact. Usually it is assumed that interactions between agents happen under some kind of a *fairness* condition. For a survey about population protocols we refer to [5].

In this direction, a very natural and simple population protocol for the majority problem on the clique (i.e. the complete graph), where initially every vertex has one of two possible types, has been introduced and analyzed in [4]. In particular, the protocol of [4] assigns only 3 possible types to every agent (i.e. there is a 3×3 transition table capturing all possible interactions) and the interactions between agents are dictated

by a *probabilistic scheduler* (i.e. all pairs have the same probability to interact at any step). Every vertex has an identity v , but it is unaware of the identity of any other vertex, as well as of its own identity. Although the underlying interaction graph in [4] is assumed to be a clique, the authors distinguish in their protocol the agents u and v participating in an interaction into an “initiator” and a “responder” of the interaction (when agents u and v interact, each of them becomes initiator or responder with equal probability). Their main result is that, if initially the difference between the initial majority from the initial minority in the complete graph with n vertices is $\omega(\sqrt{n} \log n)$, their protocol converges to the correct initial majority value after $O(n \log n)$ interactions with high probability.

Most works on population and majority dynamics so far considered only two entity types (e.g. the voter model [12], the Moran process [20]). The analysis of population dynamics with more than two types is challenging. As an example we refer to the model of [4], in which, although agents can have initially one of only two types (red and green), the protocol itself allows every agent to be of one among three different types (red, green and blank) at every subsequent time point. Even though this model is quite simple, it is very hard to be analyzed. Computing the majority with as few types as possible in the more general case, where the interaction graph has an arbitrary structure (as opposed to the complete graph that has been mainly considered so far) remained an open problem.

1.1 Our contribution

In this paper we study the majority problem in an *arbitrary* underlying interaction graph G , where initially every vertex can be of one out of two possible types (red and green). We consider here the weakest and simplest possible model of computation. In particular, we assume the existence of no central authority and we allow every vertex of G to have only a (small) constant number of available types it can change into. Although every vertex of G has a unique identity, no vertex is aware of its own identity or the identity of any other vertex. Furthermore, although only two adjacent vertices can interact, vertices of G do not even know to which other vertices they are adjacent.

First, we focus on the problem of *always* computing the correct majority value in an *arbitrary* (directed or undirected) interaction graph G , regardless of how large the initial difference between the majority and the minority is. In particular, assuming that the interacting pairs of vertices are chosen by an arbitrary fair scheduler, we prove that there does not exist any population

protocol that always computes majority in any interaction graph by using at most 3 types per vertex. Our result implies that the 4-type interval consensus protocol proposed by [6] (thoroughly analyzed in [10] under the probabilistic scheduler and also rediscovered in the conference version of our paper [18] as the *ambassador protocol*) is optimal in this sense.

Second, we provide a detailed analysis of the 3-type protocol of [4] on an arbitrary interaction graph G . Our first result in this direction is that, when the two initial types (red and green) are distributed on the vertices of an arbitrary graph G uniformly at random, the protocol of [4] will converge to the initial majority with higher probability than to the initial minority. The proof of this relies on a well known result in extremal combinatorics (in particular, on Hall's marriage Theorem). Furthermore we present an infinite family of graphs $\{G_n\}_{n \in \mathbb{N}}$, together with an initial assignment of types to vertices of each graph in this family, for which the protocol of [4] can fail (i.e. it can converge to the initial minority) with high probability, even when the difference between the initial majority and the initial minority is as large as $n - \Theta(\ln n)$. Then we present another infinite family of graphs $\{G'_n\}_{n \in \mathbb{N}}$, together with an initial assignment of types to vertices of each graph in this family, for which the protocol of [4] can take an exponential expected number of steps to converge. In particular, this rules out the possibility to use a Markov chain Monte-Carlo approach to approximate the probability that the protocol of [4] converges to the correct majority value.

In order to prove our results on the classes $\{G_n\}_{n \in \mathbb{N}}$ and $\{G'_n\}_{n \in \mathbb{N}}$, we first proved the intermediate result that for any $\varepsilon > 0$, if the minority has size at most $(\frac{1}{7} - \varepsilon)n$ in the complete graph with n vertices, then the protocol of [4] converges to the initial minority with exponentially small probability. In particular, this is an improvement upon the failure probability implied by the result of [4], where it is shown that the probability of failure is at most n^{-c} , for some constant c . Our result shows that, although the performance of the protocol of [4] can drop significantly when the interaction graph G is not the complete graph, it is quite robust when G is the complete graph. After submission of our original work in [18], we became aware of the paper [24], the results of which can also be used to prove a tighter version of our result concerning the robustness of the protocol of [4] on the clique. However, the techniques used there cannot be applied to graph structures other than the clique. On the other hand, our proof technique can also be applied in cases where the underlying interaction graph is different than (but close to being) a

clique, by using a non-trivial coupling argument which might be of independent interest.

1.2 Structure of the paper

The rest of the paper is organized as follows: In Section 2, we formally define population protocols and we present the 4-type interval consensus protocol of [6] and the 3-type majority protocol of [4]. Furthermore, we discuss how to abstract computation under the probabilistic scheduler by using Markov chains and we provide some useful preliminary results on birth-death processes that we use in our proofs. In Section 3, we present our impossibility result, that at least 4 types are necessary in any population protocol that always computes the correct majority value. In Section 4 we provide a more detailed analysis of the 3-type majority protocol of [4]. In particular, in Subsection 4.1, we present an easy result on the probability of success of the protocol in the case where the initial types are distributed uniformly at random on the vertices of the graph. In Subsection 4.2, we present our coupling method for analyzing the robustness of the majority protocol of [4], which is the main contribution of our paper. Finally, in Subsections 4.3 and Subsection 4.4 we show how our coupling method can be used to prove that there are certain graphs in which the majority protocol of [4] can fail to compute the correct majority or take an exponential number of steps to converge.

2 The model and notation

A *population protocol* consists of a finite set Q of types¹, a finite set of input symbols X , an input function $\iota : X \rightarrow Q$, a finite set of output symbols Y , an output function $\gamma : Q \rightarrow Y$, and a joint transition function $\delta : Q \times Q \rightarrow Q \times Q$. If, for any pair of types $q_a, q_b \in Q$, $\delta(q_a, q_b) = (q'_a, q'_b)$ implies that $\delta(q_b, q_a) = (q'_b, q'_a)$, then the population protocol is called *symmetric*. A population protocol is executed by a fixed finite *population* of agents with types in Q . We assume that each agent has an identity $v \in V$, but agents are oblivious to their own identity and to identities of agents they interact with.

Initially, each agent is assigned a type according to an *input* $x : V \rightarrow X$ that maps agent identities to input symbols. In the general population protocol model, agents are identified with the vertices of an *interaction graph*, whose edges indicate the possible

¹ In the original formulation of population protocols these are called *states*, but we chose to use the term *type* in order to avoid confusion with the states in a Markov chain.

agent interactions that may take place. Here, an interaction graph is a simple connected graph G (i.e. without loops or multiple edges), which can be directed or undirected. A function $C : V \rightarrow Q$ is called a *configuration*. For two configurations C, C' , we will say that C' is *directly reachable from C* (or reachable in one step from C) if there is $(v, u) \in E$, such that $(C'(v), C'(u)) = \delta(C(v), C(u))$ and $C'(w) = C(w)$, for all $w \in V \setminus \{v, u\}$. Unless specified otherwise, an interaction will correspond to a single *time step* and we will use the term *time* to mean number of interactions². In particular, we will say that the population protocol reaches configuration C at time t if, for every agent v in the population, the type of v after t interactions is $C(v)$.

In the original model, agents do not send messages or share memory; instead, an interaction between two agents updates both of their types according to a joint transition function (which can be also represented by a table). Interactions between agents are planned by a scheduler under a general “fairness” condition; the actual mechanism for choosing which agents interact is abstracted away. More precisely, following the notation in [2], a *schedule* is a finite or infinite sequence of population configurations C_0, C_1, C_2, \dots , such that C_{i+1} is directly reachable from C_i , for every $i \geq 0$. An infinite execution is *fair* if the following holds: for every two configurations C, C' , if C occurs infinitely often and C' is directly reachable from C , then C' also occurs infinitely often. We will say that a scheduler is fair if it produces fair schedules.

In this paper, we consider a special case of scheduler, namely the *probabilistic scheduler*, which was first considered in [3] and was also used in the time analysis for the protocol of [4]. The probabilistic scheduler is fair with probability 1, and is defined on directed graphs as follows. During each execution step, a directed edge $(v, u) \in E$ is chosen uniformly at random from E , where v (i.e. the tail of (v, u)) is called the *initiator* and u (i.e. the head of (v, u)) is called the *responder* of the interaction. Then, agents v and u update their types jointly according to δ . In particular, if v is of type q_v and u is of type q_u , the type of v (respectively u) becomes q'_v (respectively q'_u), where $(q'_v, q'_u) = \delta(q_v, q_u)$. The types of all other agents remain unchanged. The probabilistic scheduler is defined on undirected graphs similarly, by replacing every undirected edge $\{v, u\}$ by the two directed edges (v, u) and (u, v) . That is, in an

undirected graph G , the probabilistic scheduler selects first an undirected edge uniformly at random and then it selects equiprobably one of its endpoint as the initiator. Note that a symmetric protocol does not distinguish between initiators and responders. Therefore, if the protocol is symmetric, the probabilistic scheduler on undirected graphs just chooses at each execution step one undirected edge uniformly at random and lets its endpoints interact according to the transition function.

Given the probabilistic scheduler, a population protocol *computes* a (possibly partial) function $g : X^V \rightarrow Y$ with error probability at most ϵ , if for all $x \in g^{-1}(Y)$, the population eventually reaches a configuration C that satisfies the following properties with probability at least $1 - \epsilon$: (a) all agents agree on the correct output, i.e. $g(x) = \gamma(C(v))$ for all $v \in V$ and (b) this is also true for every configuration reachable from C . Furthermore, a population protocol *stably computes* a (possibly partial) function $g : X^V \rightarrow Y$ if, for every fair scheduler, the population eventually reaches a configuration C that satisfies both the above properties (a) and (b).

Observation 1 *If a symmetric population protocol stably computes a function on an undirected interaction graph G , then it also stably computes the same function on a directed interaction graph G' that comes from G by assigning to every edge of G one or two directions.*

2.1 Computing majority in arbitrary interaction graphs

Bénézit et al. [6] proposed a symmetric population protocol with 4 types which, given an *arbitrary* undirected graph $G = (V, E)$ as the underlying interaction graph of the population, *stably computes* the initial majority of the types of the vertices of G (even if the initial majority differs only by one from the initial minority). This protocol is known as the *interval or quantized consensus protocol* and was thoroughly analyzed under the probabilistic scheduler in [10]. It was also rediscovered in the conference version of our paper [18] as the *ambassador protocol*.

Assuming that the input symbols are g (for green) and r (for red), the set of types in the protocol is $Q = \{g0, g1, r0, r1\}$. The input function ι is such that $\iota(g) = g1$ and $\iota(r) = r1$. The output function γ is such that $\gamma(gi) = g$ and $\gamma(ri) = r$, where $i \in \{0, 1\}$. Finally, we present the transition function δ in the form of a symmetric table in Figure 1.

The proof of correctness of the interval consensus protocol can be found in [6]. A detailed analysis under the probabilistic scheduler can be found in [10]. It is also

² In the literature of population protocols, the term time sometimes refers to *parallel time*, which is equal to the number of interactions divided by the size of the population. However, we do not make this consideration in our paper; here, the terms *time* and *number of interactions* are interchangeable.

$u \setminus v$	$g0$	$g1$	$r0$	$r1$
$g0$	–	$(g1, g0)$	–	$(r1, r0)$
$g1$	$(g0, g1)$	–	$(g0, g1)$	$(g0, r0)$
$r0$	–	$(g1, g0)$	–	$(r1, r0)$
$r1$	$(r0, r1)$	$(r0, g0)$	$(r0, r1)$	–

Fig. 1 The transition matrix of the interval consensus protocol.

worth noting that population protocols that converge faster to the correct majority value were proposed and analyzed in the recent work [1] at the cost of using more than 4 types per vertex.

2.2 Computing majority with high probability on the clique

Angluin et al. [4] proposed a 3-type population protocol for computing majority with high probability (whp) in the case where the interaction graph is a clique. Their protocol uses just 3 types $Q = \{b, g, r\}$. For convenience, we will sometimes refer to these types as the *blank*, *green* and *red* type respectively. The joint transition function δ is given by:

$$\delta(x, y) = \begin{cases} (x, x), & \text{if } x = y \text{ or } y = b \\ (x, b), & \text{if } (x, y) \in \{(g, r), (r, g)\}. \end{cases} \quad (1)$$

One of the main results in [4] is that if the underlying interaction graph is a clique K_n (i.e. the complete graph on n vertices) and interactions are planned according to the probabilistic scheduler, then with high probability $1 - o(1)$, the above 3-type majority protocol converges to the initial majority value if the difference between the initial majority and initial minority populations is $\omega(\sqrt{n} \log n)$.

2.3 Representation

Using the probabilistic scheduler to plan agent interactions has the advantage that we can describe evolution by using a discrete time Markov chain \mathcal{M} . For general interaction graphs, the state space \mathcal{S} of \mathcal{M} can have up to $|V|^{|Q|}$ states³, namely one for each configuration $C : V \rightarrow Q$.

³ Here the term “state” is used in a different way than the term “configuration”: a “state” denotes a state of the Markov chain and a “configuration” is an assignment of different types of Q to the vertices V of the graph. We note that, when the underlying interaction graph has a high degree of symmetry, the number of different states of the corresponding Markov chain can be significantly smaller. Therefore, in general, we distinguish between Markov states and configurations, since states are encodings of sets of configurations.

Specifically for the model of [4], we denote by W_t (respectively R_t and G_t) the set of agents of type b (respectively r and g) at time t . Note that if the interaction graph has a high degree of symmetry, then \mathcal{S} can be reduced significantly. One such example is the clique K_n , in which case we can describe a state of \mathcal{M} by the tuple $(|R_t|, |G_t|)$ (where we have also used the fact that $|W_t| = n - |G_t| - |R_t|$).

2.4 Preliminaries on Birth-Death Processes

Consider a (*discrete time*) *birth-death process* \mathcal{B} with state space $\mathcal{S}_{\mathcal{B}} = \{S_0, S_1, \dots, S_m\}$, for some integer $m \in \mathbb{N}$ and transition probability matrix P given by

$$P(S_i, S_j) = \begin{cases} p_i, & \text{if } j = i + 1, 0 \leq i \leq m - 1 \\ q_i, & \text{if } j = i - 1, 1 \leq i \leq m \\ 1 - p_i - q_i, & \text{if } i = j \notin \{0, m\} \\ 1 - p_i, & \text{if } i = j = 0 \\ 1 - q_i, & \text{if } i = j = m. \end{cases}$$

In particular, when $p_0 = 1$ (resp. $p_0 = 0$), we say that \mathcal{B} has a *reflecting barrier* at state S_0 (resp. we say that S_0 is an *absorbing state*). Similarly, when $q_m = 1$ (resp. $q_m = 0$), we say that \mathcal{B} has a reflecting barrier at state S_m (resp. we say that S_m is an absorbing state). The fraction $\frac{p_i}{q_i}$ (resp. $\frac{q_i}{p_i}$) is called the *forward bias* (resp. *backward bias*) at state S_i . If $\frac{p_i}{q_i} = \frac{p_j}{q_j}$, for all $i, j \notin \{0, m\}$, we refer to $\frac{p_i}{q_i}$ (resp. $\frac{q_i}{p_i}$) as the *forward bias* (resp. *backward bias*) of \mathcal{B} .

We now present some useful preliminary results on discrete time birth-death processes. Other similar results concerning birth-death processes can be found for example in [22, 23, 26].

Lemma 1 (*Absorption probability*) *Consider a discrete time birth-death processes \mathcal{B} satisfying $p_i = p$ and $q_i = q$, with $p \neq q$, for all $i = 1, \dots, m - 1$. Then, given that the process starts at state S_i , the probability of reaching S_0 before reaching S_m is equal to*

$$\Pr\{\mathcal{B} \text{ reaches } S_m \text{ before } S_0 | \mathcal{B}(0) = S_i\} = \frac{\left(\frac{q}{p}\right)^i - 1}{\left(\frac{q}{p}\right)^m - 1}.$$

In particular, $\Pr\{\mathcal{B} \text{ reaches } S_m \text{ before } S_0 | \mathcal{B}(0) = S_1\} = \frac{\frac{q}{p} - 1}{\left(\frac{q}{p}\right)^m - 1}$.

Proof Let $h_i = \Pr\{\mathcal{B} \text{ reaches } S_m \text{ before } S_0 | \mathcal{B}(0) = S_i\}$, where $i = 0, \dots, m$. Then $h_0 = 0, h_m = 1$ and $h_i = ph_{i+1} + qh_{i-1} + (1 - p - q)h_i$, for $i = 1, \dots, m - 1$. Setting $u_i = h_i - h_{i-1}$, we then have that $u_{i+1} = \frac{q}{p}u_i$,

hence $u_{i+1} = \left(\frac{q}{p}\right)^i u_1$, for $i = 1, \dots, m-1$. Furthermore, $1 = u_1 + \dots + u_m = \left(1 + \frac{q}{p} + \dots + \left(\frac{q}{p}\right)^{m-1}\right) u_1$. Since $u_1 = h_1$, we have that $h_1 = \frac{1}{1 + \frac{q}{p} + \dots + \left(\frac{q}{p}\right)^{m-1}} = \frac{\frac{q}{p}-1}{\left(\frac{q}{p}\right)^m - 1}$. Finally, we have that $h_i = \sum_{j=0}^{i-1} u_{j+1} = \sum_{j=0}^{i-1} \left(\frac{q}{p}\right)^j u_1 = \frac{1 + \frac{q}{p} + \dots + \left(\frac{q}{p}\right)^{i-1}}{1 + \frac{q}{p} + \dots + \left(\frac{q}{p}\right)^{m-1}}$, which completes the proof. \square

Lemma 2 (*Expected time to absorption*) Consider a discrete time birth-death processes \mathcal{B} satisfying $p_i = p$ and $q_i = q$, with $p \neq q$, for all $i = 1, \dots, m-1$. If \mathcal{B} has a reflecting barrier at S_0 (i.e. $p_0 = 1$), then the expected time that \mathcal{B} takes to reach S_m is

$$\begin{aligned} & \mathbb{E}[\text{time for } \mathcal{B} \text{ to reach } S_m | \mathcal{B}(0) = S_i] \\ &= \left(1 + \frac{1}{p\left(\frac{q}{p}-1\right)}\right) \frac{\left(\frac{q}{p}\right)^m - \left(\frac{q}{p}\right)^i}{\frac{q}{p}-1} - \frac{m-i}{p\left(\frac{q}{p}-1\right)}. \end{aligned}$$

In particular, for large enough m , we have the following:

(a) If $\frac{q}{p} < 1$ and $i < m$, then

$$\mathbb{E}[\text{time for } \mathcal{B} \text{ to reach } S_m | \mathcal{B}(0) = S_i] = O\left(\frac{m-i}{p\left(1-\frac{q}{p}\right)}\right).$$

(b) If $\frac{q}{p} > 1$ and $i = 1$, then

$$\begin{aligned} & \mathbb{E}[\text{time for } \mathcal{B} \text{ to reach } S_m | \mathcal{B}(0) = S_1] \\ &= \Omega\left(\frac{1}{\frac{q}{p}-1} \left(\frac{q}{p}\right)^m\right). \end{aligned}$$

Proof Let $\mu_i = \mathbb{E}[\text{time for } \mathcal{B} \text{ to reach } S_m | \mathcal{B}(0) = S_i]$, where $i = 0, \dots, m$. Then $\mu_0 = 1 + \mu_1$, $\mu_m = 0$ and $\mu_i = 1 + p\mu_{i+1} + q\mu_{i-1} + (1-p-q)\mu_i$, for $i = 1, \dots, m-1$. Denote $y_{i+1} = \mu_{i+1} - \mu_i$, and thus $y_{i+1} = -\frac{1}{p} + \frac{q}{p}y_i = -\left(\frac{q}{p}\right)^i - \frac{1}{p}\left(1 + \frac{q}{p} + \dots + \left(\frac{q}{p}\right)^{i-1}\right) = -\left(\frac{q}{p}\right)^i - \frac{1}{p}\frac{\left(\frac{q}{p}\right)^i - 1}{\frac{q}{p}-1}$, for all $i = 1, \dots, m-1$. Note that $\sum_{i=1}^{m-1} y_{i+1} = -\mu_1$ and thus $\mu_1 = \sum_{i=1}^{m-1} \left[\left(1 + \frac{1}{p\left(\frac{q}{p}-1\right)}\right) \left(\frac{q}{p}\right)^i - \frac{1}{p\left(\frac{q}{p}-1\right)}\right] = \left(1 + \frac{1}{p\left(\frac{q}{p}-1\right)}\right) \frac{\left(\frac{q}{p}\right)^m - \frac{q}{p}}{\frac{q}{p}-1} - \frac{m-1}{p\left(\frac{q}{p}-1\right)}$. Furthermore, we have that

$$\begin{aligned} \mu_i - \mu_1 &= \sum_{j=1}^{i-1} y_{j+1} \\ &= \sum_{j=1}^{i-1} \left[-\left(1 + \frac{1}{p\left(\frac{q}{p}-1\right)}\right) \left(\frac{q}{p}\right)^j + \frac{1}{p\left(\frac{q}{p}-1\right)} \right] \\ &= -\left(1 + \frac{1}{p\left(\frac{q}{p}-1\right)}\right) \frac{\left(\frac{q}{p}\right)^i - \frac{q}{p}}{\frac{q}{p}-1} + \frac{i-1}{p\left(\frac{q}{p}-1\right)}, \end{aligned}$$

and thus $\mu_i = \left(1 + \frac{1}{q-p}\right) \frac{\left(\frac{q}{p}\right)^m - \left(\frac{q}{p}\right)^i}{\frac{q}{p}-1} - \frac{m-i}{p\left(\frac{q}{p}-1\right)}$ as in the statement of the Lemma. Notice then that, if $\frac{q}{p} < 1$ (or equivalently $q < p$) and m is large enough, then μ_i is dominated by the term $\frac{m-i}{p\left(1-\frac{q}{p}\right)}$, which implies part (a) of the Lemma. On the other hand, setting $i = 1$, if $\frac{q}{p} > 1$ and m is large enough, then μ_1 is dominated by the term $\left(1 + \frac{1}{q-p}\right) \frac{\left(\frac{q}{p}\right)^m - \frac{q}{p}}{\frac{q}{p}-1}$, which implies part (b) of the Lemma. \square

3 At least 4 types are needed for majority

We begin by defining the *rank* of a family of population protocols.

Definition 1 (*rank*) For any population protocol P , denote by $Q(P)$ the set of types used by P . Let \mathcal{P}_g be a class of population protocols that stably compute the function g . The *rank* of \mathcal{P}_g is $R(\mathcal{P}_g) \stackrel{\text{def}}{=} \min_{P \in \mathcal{P}_g} |Q(P)|$.

In this section, we prove that 3 types are not sufficient to stably compute the majority function in a population of agents that have initially one of two possible types and for any interaction graph. Therefore, the protocol of [6] is optimal with respect to the number of types that it uses.

Theorem 1 Let $\mathcal{P}_{\text{majority}}$ be the class of population protocols such that every protocol in $\mathcal{P}_{\text{majority}}$ stably computes the majority function in a population of agents that have initially one of two possible types and in any interaction graph. Then $R(\mathcal{P}_{\text{majority}}) > 3$.

Proof Assume for the sake of contradiction that there is a population protocol $P \in \mathcal{P}_{\text{majority}}$ that uses only 3 types, namely $Q(P) = \{r, b, g\}$, among which r and g are input symbols. Notice also, that since we initially have a 2-type population, the final answer given by P can have only two possible values (i.e. indicating the type of the initial majority). Therefore, if the output

function of P assigns different values to each distinct type in $Q(P)$, then one of the types must never occur in a final configuration. Thus, we can set the output value of this type arbitrarily equal to the output value of one of the other two types. Consequently, we may assume that the output function of P is $\gamma : Q(P) \rightarrow \{0, 1\}$.

Since, by assumption $P \in \mathcal{P}_{majority}$, we cannot have $\gamma(r) = \gamma(b) = \gamma(g)$, otherwise P would not be able to tell the difference between any two different input populations. Therefore, we may assume without loss of generality that there is a type $q_1 \in Q(P)$ such that $\gamma(q_1) = 1$, while for both the other two types q_2 and q_3 in $Q(P) \setminus \{q_1\}$, we have $\gamma(q_2) = \gamma(q_3) = 0$ (the case where q_1 is the only types having $\gamma(q_1) = 0$, while $\gamma(q_2) = \gamma(q_3) = 1$ is treated similarly, by symmetry). We can also assume without loss of generality, that when the initial majority is of type r , the protocol eventually reaches a configuration where all agents are of type q_1 , and thus P gives output 1.

Assume now that we have a population V of agents, among which $R \subseteq V$ are initially of type r and $G \subseteq V$ are initially of type g . We will denote this configuration by (R, G) . By the above discussion, if $|V| = 2k + 1$, for some integer $k \geq 2$, then P will output 1 in both the following two configurations: (a) $C_1 = (R_1, G_1)$, where $|R_1| = k + 1$, $|G_1| = k$ and (b) $C_2 = (R_2, G_2)$, where $|R_2| = k + 2$, $|G_2| = k - 1$. In particular, running P on input either C_1 or C_2 , we will eventually reach a configuration where all vertices are of type q_1 (in fact we will *never* leave this configuration once we have reached it). In other words, there is a sequence of transitions T_1 (equivalently, there is a sequence of pairs of agents picked by the scheduler) that transforms configuration C_1 to the configuration where all agents are of type q_1 . Similarly, there is a sequence of transitions T_2 that transforms configuration C_2 to the configuration where all agents are of type q_1 .

Suppose now that we have a population $V' = V \cup \{v, u\}$, i.e. V' consists of V together with two extra agents v and u (where $u, v \notin V$). Consider also the following two initial configurations: (a) $C'_1 = (R_1, G_1 \cup \{v, u\})$ and (b) $C'_2 = (R_2, G_2 \cup \{v, u\})$. In particular, in configuration C'_1 , agents in R_1 are of type r and agents in $G_1 \cup \{v, u\}$ are of type g . Furthermore, in configuration C'_2 , agents in R_2 are of type r and agents in $G_2 \cup \{v, u\}$ are of type g . Note that the majority in C'_1 is of type g , while the majority in C'_2 is of type r . Therefore, since $P \in \mathcal{P}_{majority}$, it follows that P must output 0 when the starting configuration is C'_1 and 1 when the starting configuration is C'_2 . But starting at C'_1 it is possible to follow the sequence of transitions T_1 (i.e. ignoring agents v and u during these steps), thus reaching a configuration C' where all agents in $V' \setminus \{v, u\}$ are of

type q_1 and agents v and u remain of type r . Similarly, starting at C'_2 it is possible to follow the sequence of transitions T_2 , thus reaching to the same configuration C' . This is a contradiction, since P will not be able to tell the difference between the starting configurations C'_1 and C'_2 . In particular, the output of P after reaching C' will be wrong for exactly one of the two initial configurations C'_1 or C'_2 , contradicting the assumption that $P \in \mathcal{P}_{majority}$. \square

4 The model of Angluin et al. in arbitrary interaction graphs

In this section, we provide a detailed analysis of the 3-type protocol of [4] on arbitrary interaction graphs G . In particular, in Subsection 4.1, we present our result concerning the random initial placement of individuals on the vertices of the interaction graph. In Subsection 4.2, we prove our auxiliary result that, when the minority is sufficiently small, the probability that the protocol of Angluin et al. fails in computing the majority value is exponentially small. In particular, this is an improvement upon the failure probability implied by the result of [3], where it was shown that the probability of failure is upper bounded by the inverse of a constant power of the number of vertices. Although this result shows the robustness of the protocol of [4] in the clique, we use it as an intermediate step in proving in Subsection 4.3 that there exists a family of graphs in which the protocol can fail with high probability. Finally, in Subsection 4.4, we prove the existence of a family of graphs in which the protocol of [4] can take an exponential expected number of steps to reach consensus.

4.1 Random initial placement

We prove in this subsection a preliminary result concerning the model of [4] when the interaction graph is an arbitrary, strongly connected, directed graph G . In particular, we prove that if the initial assignment of individuals to the vertices of G is random, then the majority protocol described in [4] correctly identifies the initial majority with probability at least $\frac{1}{2}$.

For the proof of Theorem 2 below, we will need a result from extremal combinatorics concerning systems of distinct representatives. A *system of distinct representatives* for a sequence of (not necessarily distinct) sets $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_x$ is a sequence of distinct elements e_1, e_2, \dots, e_x such that $e_i \in \mathcal{T}_i$ for all $i = 1, 2, \dots, x$. The following is a consequence of Hall's marriage Theorem (for a proof see chapter 5 in [13]):

Corollary 1 ([13]) *Let $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_x$ be r -element subsets of a universe of y elements, such that each element belongs to the same number $d \geq 1$ of these sets. If $x \leq y$, then the sets $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_x$ have a system of distinct representatives.*

We are now ready to prove the following Theorem.

Theorem 2 *For any strongly connected directed graph G , if the initial assignment of individuals to the vertices of G is random, then the majority protocol described in [4] correctly identifies the initial majority with probability at least $\frac{1}{2}$.*

Proof Assume without loss of generality that the initial majority is of type g (green), i.e. $|G_0| \geq |R_0|$. The state space of the Markov chain \mathcal{M} describing the evolution of the protocol will be the set $\mathcal{S} = \{(A, B) : A, B \subseteq V, A \cap B = \emptyset\}$. In particular, when \mathcal{M} is in state (A, B) at time t , we will write that $(R_t, G_t) = (A, B)$, and this will mean that the subset of vertices A (respectively B) is of type r (respectively g). It is also evident that the set of vertices $V - (A \cup B)$ is of type b . It then follows that the majority protocol in [4] fails to correctly identify type g as the initial majority if \mathcal{M} eventually reaches state (V, \emptyset) ; we will then say that the *initial minority wins* (whereas in the opposite case we will say that the *initial majority wins*). Since the initial placement of individuals on the vertices of the graph is random and the initial minority is of type r , we also have that

$$\begin{aligned} & \Pr\{\text{initial minority wins}\} \\ &= \sum_{S \subseteq V: |S|=|G_0|} \left[\frac{1}{\binom{n}{|G_0|}} \cdot \Pr\{(V, \emptyset) \text{ reached} | (R_0, G_0) = (V - S, S)\} \right]. \quad (2) \end{aligned}$$

Consider now a fixed choice of S in the above sum and an arbitrary set $A \subseteq S$ of size $|A| = |R_0| = |V - S|$ (notice that it is always possible to find such a set A , since $|S| \geq |V - S|$, by assumption). Then we have

$$\begin{aligned} & \Pr\{(V, \emptyset) \text{ reached} | (R_0, G_0) = (V - S, S)\} \\ & \leq \Pr\{(V, \emptyset) \text{ reached} | (R_0, G_0) = (V - A, A)\} \\ & = \Pr\{(\emptyset, V) \text{ reached} | (R_0, G_0) = (A, V - A)\}. \quad (3) \end{aligned}$$

For the inequality we used the fact that $V - S \subseteq V - A$ and that augmenting the initial set of agents (inclusion-wise) of type r increases the probability that agents of type r win. The basic proof idea behind this fact is that, simulating the total possible executions that lead to either a total red or total green configuration, the extra agents of type r can only work in favor of the r type.

For the equality we used the fact that the protocol is symmetric for types g and r . In words, (3) states that the probability that the initial minority wins, starting from $(V - S, S)$, is at most the probability that the initial majority wins if we exchange agents of type r initially placed in $V - S$ with agents of type g initially placed in A .

Notice now that if for any choice of S we assign a unique choice of $A = A_S \subseteq S$, i.e. such that for any $S \neq S'$ we have $A_S \neq A_{S'}$, then we are done, since by (2) and (3) we have

$$\begin{aligned} & \Pr\{\text{initial minority wins}\} \\ &= \sum_{S \subseteq V: |S|=|G_0|} \left[\frac{1}{\binom{n}{|G_0|}} \cdot \Pr\{(V, \emptyset) \text{ reached} | (R_0, G_0) = (V - S, S)\} \right] \\ & \leq \sum_{S \subseteq V: |S|=|G_0|} \left[\frac{1}{\binom{n}{|G_0|}} \cdot \Pr\{(\emptyset, V) \text{ reached} | (R_0, G_0) = (A_S, V - A_S)\} \right] \\ & \leq \Pr\{\text{initial majority wins}\}. \end{aligned}$$

We now show that such a one to one correspondence of sets A_S to sets S is possible: For $k = |R_0|$, let $T_1, T_2, \dots, T_{\binom{n}{k}}$ be an arbitrary enumeration of subsets of V of size exactly k . For any $1 \leq i \leq \binom{n}{k}$, let also \mathcal{T}_i be the set of subsets of $V - T_i$ of size k . Notice that the set \mathcal{T}_i is non-empty (because $n - k \geq k$ by assumption) and is also uniquely determined by T_i . Furthermore, each set inside \mathcal{T}_i can be formed by selecting exactly k vertices in $V \setminus T_i$, therefore, the sets \mathcal{T}_i have the same size $r = \binom{n-k}{k}$. Additionally, by definition, a set T_j belongs to some \mathcal{T}_i if and only if T_i (which uniquely determines \mathcal{T}_i) does not have any vertices in common with T_j , and so there are exactly $d = \binom{n-k}{k}$ distinct \mathcal{T}_i 's that contain T_j . Therefore, we can apply Corollary 1 to the sets \mathcal{T}_i , with $x = y = \binom{n}{k}$. In particular, this implies that the sets \mathcal{T}_i , $1 \leq i \leq \binom{n}{k}$, have a system of distinct representatives, which concludes the proof. \square

In the above theorem we provided a sufficient condition under which the majority protocol described in [4] correctly identifies the initial majority with probability at least $\frac{1}{2}$. This result is in wide contrast to the negative result of Subsection 4.3 (cf. Theorem 4), in which we highlight a case where the majority protocol of [4] fails with high probability.

It is worth noting that Theorem 2 does not imply any bounds on the expected convergence time of the protocol of [4] on arbitrary interaction graphs under the

probabilistic scheduler; finding such bounds remains an open problem.

4.2 Clique

In this Subsection we provide an upper bound on the probability that all agents eventually become of type r , given that we start with ϵn agents of type r and $(1 - \epsilon)n$ agents of type g on the n -vertex clique, where $\frac{1}{n} \leq \epsilon < \frac{1}{7}$ (the upper bound on ϵ is used to facilitate the exposition of our arguments, while the lower bound guarantees that there is at least one vertex of type r). We assume without loss of generality that ϵn is an integer. By the discussion in Subsection 2.3, the state space of the Markov chain \mathcal{M} describing the evolution of the protocol at time t contains tuples of the form $(|R_t|, |G_t|)$, where R_t (resp. G_t) is the set of vertices of type r (resp. g) at time t . In particular, we are interested in upper bounding

$$\Pr\{\text{absorption at } (n, 0) \mid \text{initially at } (\epsilon n, n - \epsilon n)\}.$$

The core of our proof lies in the definition of two discrete time processes \mathcal{W} and \mathcal{C} that “filter” the information from the original Markov chain \mathcal{M} .

Definition 2 (The Blank Process \mathcal{W}) This process keeps track of the number of blank vertices over time, i.e. $\mathcal{W}(t) \stackrel{\text{def}}{=} \langle \# \text{ vertices of type } b \text{ at time } t \rangle$.

For convenience, we will use the following notation to describe transitions of \mathcal{M} : We will write $g \rightarrow r$ to describe a transition of the form $(x, y) \rightarrow (x - 1, y)$, for some $x, y \in \{1, 2, \dots, n\}$. More specifically, $g \rightarrow r$ is used to describe a transition where a directed edge (v, u) is chosen by the scheduler, v is of type g , and u is of type r . Similarly, we will use $r \rightarrow g$ for transitions of the form $(x, y) \rightarrow (x, y - 1)$, $g \rightarrow b$ for transitions of the form $(x, y) \rightarrow (x, y + 1)$ and $r \rightarrow b$ for transitions of the form $(x, y) \rightarrow (x + 1, y)$. We note that the state of \mathcal{M} at any time t can be fully described by the initial state and by a sequence of transitions among $\{g \rightarrow r, r \rightarrow g, g \rightarrow b, r \rightarrow b\}$.

Definition 3 (The Contest Process \mathcal{C}) Transitions of \mathcal{M} are paired recursively, starting from time 0 as follows: Every transition that increases the number of blanks is paired with the earliest subsequent transition that decreases the number of blanks and is not paired yet⁴. For an arbitrary time t , we denote by $\tau(t)$ (or

just τ for short) the number of pairs until time t . The Contest Process \mathcal{C} is defined over the time scale τ , where $\mathcal{C}(0) = |R_0|$ and for $\tau = 1, 2, \dots$,

$$\mathcal{C}(\tau) = \begin{cases} \mathcal{C}(\tau - 1) + 1, & \text{if the } \tau\text{-th pair is } (r \rightarrow g, r \rightarrow b), \\ \mathcal{C}(\tau - 1) - 1, & \text{if the } \tau\text{-th pair is } (g \rightarrow r, g \rightarrow b), \\ \mathcal{C}(\tau - 1), & \text{otherwise.} \end{cases}$$

For example, suppose that the underlying clique has $n = 4$ vertices and we have the sequence of states $S_0 = (1, 3), S_1 = (1, 2), S_2 = (1, 3), S_3 = (1, 2), S_4 = (1, 1), S_5 = (2, 1), S_6 = (2, 0), S_7 = (3, 0), S_8 = (4, 0)$. In the new notation, we start at $S_0 = (1, 3)$ and then we have the transitions (1) $r \rightarrow g$, (2) $g \rightarrow b$, (3) $r \rightarrow g$, (4) $r \rightarrow g$, (5) $r \rightarrow b$, (6) $r \rightarrow g$, (7) $r \rightarrow b$, (8) $r \rightarrow b$. The value of \mathcal{W}_t , for any $t = 0, \dots, 8$ is then easy to find (e.g. $\mathcal{W}_5 = 1$). For the Contest process \mathcal{C} , we pair the transitions as follows: (a) $S_0 \rightarrow S_1$ (which is $r \rightarrow g$) is paired to $S_1 \rightarrow S_2$ (which is $g \rightarrow b$), (b) $S_2 \rightarrow S_3$ (which is $r \rightarrow g$) is paired to $S_4 \rightarrow S_5$ (which is $r \rightarrow b$), (c) $S_3 \rightarrow S_4$ (which is $r \rightarrow g$) is paired to $S_6 \rightarrow S_7$ (which is $r \rightarrow b$) and (d) $S_5 \rightarrow S_6$ (which is $r \rightarrow g$) is paired to $S_7 \rightarrow S_8$ (which is $r \rightarrow b$). In particular, there are only 4 transitions for \mathcal{C} and in particular, we have $\mathcal{C}(0) = 1, \mathcal{C}(1) = 1, \mathcal{C}(2) = 2, \mathcal{C}(3) = 3, \mathcal{C}(4) = 4$.

Notice that the processes \mathcal{W} and \mathcal{C} are dependent. As a matter of fact, \mathcal{C} is not even defined using the same time scale as \mathcal{W} and \mathcal{M} (to indicate this, we have used the convention that t is the time variable for processes \mathcal{M}, \mathcal{W} , while τ is the time variable for process \mathcal{C}). However, observe that if we initially begin with no blanks (i.e. $|R_0| + |G_0| = n$ hence $\mathcal{W}(0) = 0$), then whenever \mathcal{W} decreases its value, we have a transition step of \mathcal{C} . Additionally, we can prove the following:

Lemma 3 (Relating \mathcal{C} and \mathcal{M}) For any $T \in \mathbb{N}$, denote by $\mathcal{C}_{|T}$ the value of \mathcal{C} given only states $\mathcal{M}(t), t = 0, 1, \dots, T$ (i.e. given the history of \mathcal{M} up to time T). Then, $\mathcal{C}_{|T} \geq |R_T|$ for any $T \in \mathbb{N}$. Furthermore, if $\mathcal{C}_{|T} = 0$, then all vertices are of type g .

Proof Notice that if all transitions up to time T were paired according to the pairing in the definition of \mathcal{C} , then we would have exactly $\mathcal{C}_{|T} = |R_T|$. Indeed, each pair of transitions of the form $(r \rightarrow g, r \rightarrow b)$ increases the number of vertices of type r by 1, each pair of transitions of the form $(g \rightarrow r, g \rightarrow b)$ decreases the number of vertices of type r by 1 and any pair of transitions of the form $(r \rightarrow g, g \rightarrow b)$ or $(g \rightarrow r, r \rightarrow b)$ does not change the number of vertices of type r . Finally, notice that transitions that are not paired are either of the form $r \rightarrow g$, or $g \rightarrow r$, which can only decrease the number of vertices of type r . This completes the proof of the first part.

⁴ We assume that the pairing concerns only transitions that change the state of \mathcal{M} . In particular, transitions of the form $b \rightarrow r, b \rightarrow g, b \rightarrow b, g \rightarrow g$ and $r \rightarrow r$ are ignored in this pairing as irrelevant.

For the second part of the Lemma, let T_0 be a time where $\mathcal{C}_{|T_0} = 0$ and note that because of the first part, we already have that $|R_t| = 0$, so we only need to prove that also $\mathcal{W}(T_0) = 0$. Assume for the sake of contradiction that there is some vertex w that is of type b at time T_0 . By definition, this implies that there is an unpaired transition of the form $r \rightarrow g$ or $g \rightarrow r$. Indeed, this follows from the observation that paired transitions do not change the number of vertices of type b , therefore if all transitions are paired there must be no vertices of type b remaining. If now the unpaired transition is of the form $r \rightarrow g$, then the existence of a vertex of type r at the time of the transition (say time $t < T_0$) together with the fact that $|R_{T_0}| = 0$ would imply that there is also an unpaired subsequent transition (i.e. a transition that happened at some time t' , with $T_0 > t' > t$) of the form $g \rightarrow r$. But if there is an unpaired transition of the form $g \rightarrow r$, then pairing it with a subsequent transition of the form $g \rightarrow b$ after time T_0 would decrease the value of \mathcal{C}' by 1, making it negative, which leads to a contradiction because of the first part of the Lemma. This completes the proof of the second part. \square

We will use the following domination statements in Lemmas 4 and 5, which concern the domination of processes \mathcal{W} and \mathcal{C} by appropriate birth-death processes:

Lemma 4 (Domination of \mathcal{W}) *Let $\alpha, \beta, \kappa \in \{1, \dots, n-1\}$, with $\alpha < \beta$. Let also $\mathcal{B}_{\mathcal{W}}$ be a birth-death process, which has state space $\mathcal{S}_{\mathcal{B}_{\mathcal{W}}} = \{S_0, \dots, S_n\}$, with S_n an absorbing state and transition probability matrix P , with $P(S_i, S_{i+1}) = 1$ for all $i \in \{0, \dots, \alpha\} \cup \{\beta, \dots, n-1\}$, $\frac{P(S_i, S_{i+1})}{P(S_i, S_{i-1})} = \frac{2\kappa}{\alpha}$ for all $i \in \{\alpha+1, \dots, \beta-1\}$ and $P(S_i, S_i) = \Pr(\mathcal{W}(t) = i | \mathcal{W}(t-1) = i)$, for all $t \geq 1$ and for all $i \in \{\alpha+1, \dots, \beta-1\}$. Then, given that the vertices of type r are at most κ , the process \mathcal{W} is stochastically dominated by $\mathcal{B}_{\mathcal{W}}$ in the following sense: $\Pr(\mathcal{W}(t) > x | \mathcal{W}(0) = 0) \leq \Pr(\mathcal{B}_{\mathcal{W}}(t) \in \cup_{y>x} S_y | \mathcal{B}_{\mathcal{W}}(0) = S_0)$, for any time t and $x \in \{0, \dots, n\}$.*

Proof For the proof, it suffices to show that, for any $t \geq 1$ and for any $i \in \{0, \dots, n-1\}$,

$$\begin{aligned} \Pr(\mathcal{W}(t) = i + 1 | \mathcal{W}(t-1) = i) \\ \leq \Pr(\mathcal{B}_{\mathcal{W}}(t) = S_{i+1} | \mathcal{B}_{\mathcal{W}}(t-1) = S_i). \end{aligned}$$

This is trivially true for all $i \in \{0, \dots, \alpha\} \cup \{\beta, \dots, n-1\}$, because the right hand side of the above inequality is 1. For $i \in \{\alpha+1, \dots, \beta-1\}$, given that $\mathcal{W}(t-1) = i$ (i.e. there are exactly i blanks at time t), the probability that \mathcal{W} increases by 1 in the next time step is equal to the probability that either a transition $g \rightarrow r$ or a transition $r \rightarrow g$ occurs, which is equal to $\frac{2|R_{t-1}| |G_{t-1}|}{n(n-1)}$. On the other hand, given that $\mathcal{W}(t-1) = i$, the probability

that \mathcal{W} decreases by 1 in the next time step is equal to the probability that either a transition $g \rightarrow b$ or a transition $r \rightarrow b$ occurs, which is equal to $\frac{i(|R_{t-1}| + |G_{t-1}|)}{n(n-1)}$. But then

$$\begin{aligned} \frac{\Pr(\mathcal{W}(t) = i + 1 | \mathcal{W}(t-1) = i)}{\Pr(\mathcal{W}(t) = i - 1 | \mathcal{W}(t-1) = i)} &= \frac{2|R_{t-1}| |G_{t-1}|}{i(|R_{t-1}| + |G_{t-1}|)} \\ &\leq \frac{2|R_{t-1}|}{i}. \end{aligned}$$

By assumption, this is at most $\frac{2\kappa}{\alpha}$, which combined with the fact that $P(S_i, S_i) = \Pr(\mathcal{W}(t) = i | \mathcal{W}(t-1) = i)$, for all $i \in \{\alpha+1, \dots, \beta-1\}$, concludes the proof. \square

Lemma 5 (Domination of \mathcal{C}) *Let β, κ be positive integers, with $\beta + \kappa < n$. Let also $\mathcal{B}_{\mathcal{C}}$ be a birth-death process, which has state space $\mathcal{S}_{\mathcal{B}_{\mathcal{C}}} = \{T_0, \dots, T_n\}$, with T_0, T_n absorbing states and transition probability matrix Q , with $Q(T_i, T_{i+1}) = 1$ for all $i \in \{\kappa, \dots, n-1\}$, $\frac{Q(T_i, T_{i+1})}{Q(T_i, T_{i-1})} = \frac{\kappa}{n-\beta-\kappa}$ for all $i \in \{1, \dots, \kappa-1\}$ and $Q(T_i, T_i) = \Pr(\mathcal{C}(\tau) = i | \mathcal{C}(\tau-1) = i)$, for all $\tau \geq 1$ and for all $i \in \{1, \dots, \kappa-1\}$. Then, given that the vertices of type b are at most β , the process \mathcal{C} is stochastically dominated by $\mathcal{B}_{\mathcal{C}}$ in the following sense: $\Pr(\mathcal{C}(\tau) > x | \mathcal{C}(0) = |R_0|) \leq \Pr(\mathcal{B}_{\mathcal{C}}(\tau) \in \cup_{y>x} T_y | \mathcal{B}_{\mathcal{C}}(0) = |R_0|)$, for any τ and $x \in \{0, \dots, n\}$.*

Proof It suffices to show that, for any $\tau \geq 1$ and for any $i \in \{1, \dots, n-1\}$,

$$\begin{aligned} \Pr(\mathcal{C}(\tau) = i + 1 | \mathcal{C}(\tau-1) = i) \\ \leq \Pr(\mathcal{B}_{\mathcal{C}}(\tau) = T_{i+1} | \mathcal{B}_{\mathcal{C}}(\tau-1) = T_i). \end{aligned}$$

This is trivially true for all $i \in \{\kappa, \dots, n-1\}$, because the right hand side of the above inequality is 1. For $i \in \{1, \dots, \kappa-1\}$, we apply the *principle of deferred decisions*. In particular, let t_1 (resp. t_2) be the time in the time scale of \mathcal{M} that corresponds to the first (resp. second) transition of the τ -th transition pair in the definition of \mathcal{C} (notice that both t_1 and t_2 are random variables). Given that $\mathcal{C}(\tau-1) = i$, the probability that \mathcal{C} increases by 1 in the next time step in the time scale of \mathcal{C} is equal to the probability that the τ -th transition pair is $(r \rightarrow g, r \rightarrow b)$, which is equal to $\Pr(\mathcal{C}(\tau) = i + 1 | \mathcal{C}(\tau-1) = i, t_1, t_2) = \frac{|R_{t_1}| |G_{t_1}|}{2|R_{t_1}| |G_{t_1}|} \frac{|R_{t_2}| (n - |R_{t_2}| - |G_{t_2}|)}{(|R_{t_2}| + |G_{t_2}|)(n - |R_{t_2}| - |G_{t_2}|)}$. On the other hand, the probability that the τ -th transition pair in the definition of \mathcal{C} is $(g \rightarrow r, g \rightarrow b)$ is $\Pr(\mathcal{C}(\tau) = i - 1 | \mathcal{C}(\tau-1) = i, t_1, t_2) = \frac{|G_{t_1}| |R_{t_1}|}{2|R_{t_1}| |G_{t_1}|} \frac{|G_{t_2}| (n - |R_{t_2}| - |G_{t_2}|)}{(|R_{t_2}| + |G_{t_2}|)(n - |R_{t_2}| - |G_{t_2}|)}$. But then

$$\frac{\Pr(\mathcal{C}(\tau) = i + 1 | \mathcal{C}(\tau-1) = i, t_1, t_2)}{\Pr(\mathcal{C}(\tau) = i - 1 | \mathcal{C}(\tau-1) = i, t_1, t_2)} = \frac{|R_{t_2}|}{|G_{t_2}|}.$$

Since $i \leq \kappa - 1$, by assumption, it follows by Lemma 3 that $|R_{t_2}| < \kappa$. Therefore, we have that

$$\frac{\Pr(\mathcal{C}(\tau) = i + 1 | \mathcal{C}(\tau-1) = i)}{\Pr(\mathcal{C}(\tau) = i - 1 | \mathcal{C}(\tau-1) = i)} \leq \frac{\kappa}{n - \beta - \kappa},$$

which, combined with the fact that $Q(T_i, T_i) = \Pr(\mathcal{C}(\tau) = i | \mathcal{C}(\tau - 1) = i)$, for all $i \in \{1, \dots, \kappa - 1\}$, concludes the proof. \square

We are now ready to prove our main Theorem, which is stated below.

Theorem 3 *Let $\epsilon < \frac{1}{7}$. For large enough n , starting from ϵn agents of type r and $(1 - \epsilon)n$ agents of type g on the clique K_n , the probability that the clique eventually contains only agents of type r is at most $\left(\frac{5}{2}\right)^{n(\epsilon - \frac{1}{7})} + 2\left(\frac{7}{6}\right)^{-\frac{n}{13}} = e^{-\Theta(n)}$.*

Proof For the proof, we will upper bound the probability that the initial minority wins by providing upper bounds to the probability that the following events occur, in which κ, β are predefined integers (which we fix later to facilitate exposition):

- (i) A_1 is the event that \mathcal{C} reaches κ before the number of vertices of type b reaches β .
- (ii) A_2 is the event that the number of vertices of type b reaches β before \mathcal{C} reaches κ .

In particular, we note that

$$\Pr\{\text{initial minority wins} | \mathcal{M}(0) = (\epsilon n, n - \epsilon n)\} \leq \Pr(A_1 \cup A_2 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n)) \quad (4)$$

so we need to provide upper bounds for $\Pr(A_1 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n))$ and $\Pr(A_2 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n))$.

We now set $\alpha = \frac{n}{3}$, $\beta = \frac{n}{2}$ and $\kappa = \frac{n}{7}$. In that case, the forward bias of the birth-death process \mathcal{B}_W at states $\{S_{\alpha+1}, S_{\alpha+2}, \dots, S_{\beta-1}\}$ is at most $r_W \stackrel{\text{def}}{=} \frac{6}{7} < 1$ and the forward bias of the birth-death process \mathcal{B}_C at states $\{T_1, T_2, \dots, T_{\kappa-1}\}$ is at most $r_C \stackrel{\text{def}}{=} \frac{2}{5} < 1$.

In particular, by Lemma 5 and Lemma 1, given that the number of vertices of type b is less than β , the probability that the process \mathcal{C} reaches value κ before reaching 0, given that it starts at ϵn , is at most $\frac{\left(\frac{1}{r_C}\right)^{\epsilon n} - 1}{\left(\frac{1}{r_C}\right)^{\frac{n}{7}} - 1}$.

Therefore,

$$\begin{aligned} \Pr(A_1 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n)) &\leq \frac{\left(\frac{1}{r_C}\right)^{\epsilon n} - 1}{\left(\frac{1}{r_C}\right)^{\frac{n}{7}} - 1} \\ &= \frac{\left(\frac{5}{2}\right)^{\epsilon n} - 1}{\left(\frac{5}{2}\right)^{\frac{n}{7}} - 1}. \end{aligned} \quad (5)$$

For the bound on $\Pr(A_2 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n))$, notice that, given $\mathcal{M}(0) = (\epsilon n, n - \epsilon n)$, the event A_2 can only happen if \mathcal{W} reaches β before \mathcal{C} reaches either 0 or κ . For integers k_0 and k_1 we will define the following events: (i) B_1 is the event that \mathcal{W} reaches β and backtracks less than k_0 times until then and (ii) B_2 is the

event that \mathcal{W} reaches β and \mathcal{C} takes more than k_1 steps (in the time scale of \mathcal{C}) to reach either value 0 or κ .

We will set $k_0 = k_1 = \left(\frac{1}{r_W}\right)^{\frac{n}{12}} = \left(\frac{7}{6}\right)^{\frac{n}{12}}$ (however, note here that equation (6) remains valid for any $k_1 \leq k_0$). By observing that, every time \mathcal{W} backtracks, we have a step in \mathcal{C} , we have that

$$\begin{aligned} \Pr(A_2 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n)) &\leq \Pr(B_1 \cup B_2 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n)). \end{aligned} \quad (6)$$

Indeed, given $\mathcal{M}(0) = (\epsilon n, n - \epsilon n)$, the probability that \mathcal{W} reaches β before \mathcal{C} reaches either 0 or κ is at most the probability of the event that \mathcal{W} reaches β and backtracks less than k_0 times until then (i.e. B_1) or that \mathcal{W} reaches β and backtracks at least k_0 times but it takes longer for \mathcal{C} to reach either 0 or κ (i.e. $B_2 \cap \overline{B_1}$).

Notice now that, by Lemma 4 and Lemma 1, the probability that the process \mathcal{W} reaches β before going back to α , given that it starts with $\alpha + 1$ blank vertices is at most $\frac{\frac{1}{r_W} - 1}{\left(\frac{1}{r_W}\right)^{\frac{n}{6}} - 1}$. Therefore, by coupling \mathcal{W} and \mathcal{B}_W (defined in Lemma 4), we have that

$$\begin{aligned} \Pr(B_1 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n)) &\leq \Pr\left\{\mathcal{B}_W \text{ reaches } S_\beta \text{ and backtracks } < k_0 \text{ times} \mid \mathcal{B}_W(0) = \alpha + 1\right\} \end{aligned} \quad (7)$$

$$\leq \Pr\left\{\mathcal{B}_W \text{ reaches } S_\beta \text{ and visits } S_\alpha < k_0 \text{ times} \mid \mathcal{B}_W(0) = \alpha + 1\right\} \quad (8)$$

$$\leq 1 - \left(1 - \frac{\frac{1}{r_W} - 1}{\left(\frac{1}{r_W}\right)^{\frac{n}{6}} - 1}\right)^{k_0} \quad (9)$$

$$\leq k_0 \left(\frac{1}{r_W}\right)^{-\frac{n}{6}} = \left(\frac{7}{6}\right)^{-\frac{n}{12}}. \quad (10)$$

Inequality (7) follows by domination in Lemma 4. Inequality (9) follows from the memoryless property of Markov chains and the observation that every time \mathcal{B}_W takes the value α it immediately takes the value $\alpha + 1$ on the next step, since it has a reflecting barrier on S_α . Finally, to get inequality (10) we used the fact that $(1 - x)^y \geq 1 - 2xy$ whenever $0 < xy < \frac{1}{2}$.

Let now X_C be the number of steps (in the time scale of \mathcal{C}) needed for process \mathcal{C} to reach either value 0 or κ . Notice also that we can use Lemma 2(a) in order to upper bound the expected time needed for \mathcal{B}_C to reach either T_0 or T_κ . In particular, this is at most the time needed for \mathcal{B}_C to reach T_0 if it had a reflecting barrier at T_κ , which is a mirror birth-death process to that of Lemma 2(a) if we ignore all states other than T_1, \dots, T_κ (i.e. with $m = \kappa$ and $S_i = T_{\kappa-i}, i = 0, \dots, \kappa$). Therefore, the expected time needed for \mathcal{B}_C to reach either T_0

or T_κ , given that it starts at $T_{\epsilon n}$ is $O\left(\frac{\epsilon n |E|^2}{\frac{1}{rc} - 1}\right) = O(n^5)$, where we also used the fact that the probability that \mathcal{C} (and thus also $\mathcal{B}_\mathcal{C}$) does not stay on the same state, given that it starts at one of the states $T_1, \dots, T_{\kappa-1}$, is at least $\Omega\left(\frac{1}{|E|^2}\right)$ (in particular, this is a lower bound on the value of p in Lemma 2(a)). Therefore, by Lemma 5 and taking into consideration that $\mathcal{C}(0) = \epsilon n$, the expected value of $X_\mathcal{C}$ satisfies

$$\mathbb{E}[X_\mathcal{C} | \mathcal{C}(0) = \epsilon n] = O(n^5).$$

By Markov's inequality, for large enough n , we then have that

$$\Pr(X_\mathcal{C} > k_1 | \mathcal{C}(0) = \epsilon n) < \frac{\mathbb{E}[X_\mathcal{C} | \mathcal{C}(0) = \epsilon n]}{k_1} = \frac{O(n^5)}{k_1}.$$

In particular, we have that

$$\begin{aligned} \Pr(B_2 \cap \overline{B_1} | \mathcal{M}(0) = (\epsilon n, n - \epsilon n)) \\ \leq \Pr(X_\mathcal{C} > k_1 | \mathcal{C}(0) = \epsilon n) \\ = \frac{O(n^5)}{k_1} \leq \left(\frac{7}{6}\right)^{-\frac{n}{13}}. \end{aligned} \quad (11)$$

By using now the bounds we have for $\Pr(A_1 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n))$ in (5), $\Pr(B_1 | \mathcal{M}(0) = (\epsilon n, n - \epsilon n))$ in (10), and $\Pr(B_2 \cap \overline{B_1} | \mathcal{M}(0) = (\epsilon n, n - \epsilon n))$ in (11), together with inequalities (4) and (6), we get that

$$\begin{aligned} \Pr\{\text{initial minority wins} | \mathcal{M}(0) = (\epsilon n, n - \epsilon n)\} \\ \leq \frac{\left(\frac{1}{rc}\right)^{\epsilon n} - 1}{\left(\frac{1}{rc}\right)^{\frac{n}{7}} - 1} + k_0 \left(\frac{1}{r\mathcal{W}}\right)^{-\frac{n}{6}} + \frac{O(n^5)}{k_1} \\ \leq \frac{\left(\frac{5}{2}\right)^{\epsilon n} - 1}{\left(\frac{5}{2}\right)^{\frac{n}{7}} - 1} + \left(\frac{7}{6}\right)^{-\frac{n}{13}} + \left(\frac{7}{6}\right)^{-\frac{n}{13}}. \end{aligned} \quad (12)$$

This concludes the proof of the Theorem. \square

We note here that the upper bound on ϵ in the statement of Theorem 3 is only used to facilitate exposition of our arguments in the proof. We claim that this upper bound can be increased further by using the same proof ideas, but that we cannot reach arbitrarily close to $\frac{1}{2}$. However, we conjecture that the constant ϵ in Theorem 3 can be as close to $\frac{1}{2}$ as desired, as long as it remains bounded away from it.

Conjecture 1 Let ϵ be a positive constant strictly less than $\frac{1}{2}$. For large enough n , starting from ϵn agents of type r and $(1 - \epsilon)n$ agents of type g on the clique K_n , the probability that the clique eventually contains only agents of type r is at most $e^{-\Theta(n)}$.




Fig. 2 A lollipop graph consisting of a K_{n_1} clique and a L_{n_2} line.

4.3 Minority domination

Consider the lollipop graph (see Figure 2), which consists of a complete graph K_{n_1} on n_1 vertices, among which vertex v is connected to the leftmost vertex u of a line graph L_{n_2} on n_2 vertices, with $n_1 + n_2 = n$. Suppose that initially vertex v and all vertices in L_{n_2} are of type r , while all vertices in $K_{n_1} \setminus \{v\}$ are of type g . Without loss of generality, we also assume that $n_1 < n_2$, so that the color green is the initial minority in the graph. In this subsection we will provide a lower bound on the probability that the initial minority eventually wins, given that we start with this configuration.

For the analysis we will use the analysis of Subsection 4.2, together with a suitable domination argument and the following fact:

Lemma 6 Consider a line graph L_m , in which the leftmost vertex is of type g and all other $m - 1$ vertices are of type r . Starting from this initial configuration, the probability that eventually all vertices become of type g is $\frac{1}{2(m-1)}$.

Proof Let \mathcal{M} denote the Markov chain that describes the evolution of the protocol. Without loss of generality, we will assume that only transitions that result in some vertex changing its type are allowed. In particular, discarding any transitions of \mathcal{M} resulting from interactions of the form $b \rightarrow r, b \rightarrow g, b \rightarrow b, g \rightarrow g$ or $r \rightarrow r$ does not change the probability of absorption of \mathcal{M} to one of its two absorbing states (namely the one with all vertices of type r and the one with all vertices of type g).

Let us now construct a stochastic process \mathcal{B} by taking snapshots of \mathcal{M} every two transitions that change the state of \mathcal{M} . Then we can describe the set of states reachable by \mathcal{M} (given the initial configuration) by a single number between 0 and m . Indeed, starting from a single vertex of type g on the leftmost vertex of the line and all other vertices of type r , and taking snapshots of \mathcal{M} every two transitions, we can only reach configurations S_k , $k = 0, \dots, m$, in which a number of k consecutive vertices to the left are of type g and all others are of type r . By the Markov property we can easily verify that \mathcal{B} is a Markov chain with state

space $\{S_0, S_1, \dots, S_m\}$ (where S_0, S_m are absorbing) and transition probability matrix P given by

$$P(S_i, S_j) = \begin{cases} \frac{1}{4}, & \text{if } |j - i| = 1, 2 \leq i \leq m - 2, \\ \frac{1}{4}, & \text{if } j = i + 1 = 2 \text{ or } j = i - 1 = m - 2, \\ \frac{1}{2}, & \text{if } j = i + 1 = m \text{ or } j = i - 1 = 0, \\ 1, & \text{if } i = j = 0 \text{ or } i = j = m, \\ 1 - P(S_i, S_{i+1}) - P(S_i, S_{i-1}), & \text{if } i = j, 1 \leq i \leq m - 1 \end{cases}$$

In particular, we have that

$$\frac{P(S_i, S_{i-1})}{P(S_i, S_{i+1})} = 1,$$

for all $i = 2, 3, \dots, m - 2$ and $\frac{P(S_1, S_0)}{P(S_1, S_2)} = \frac{P(S_{m-1}, S_m)}{P(S_{m-1}, S_{m-2})} = 2$. Furthermore, we have that

$$\begin{aligned} & \Pr\{\mathcal{M} \text{ reaches } S_m | \mathcal{M}(0) = S_1\} \\ &= \Pr\{\mathcal{B} \text{ reaches } S_m | \mathcal{B}(0) = S_1\}. \end{aligned}$$

In order to compute the above probability, we proceed as in the proof of Lemma 1. In particular, let $h_i = \Pr\{\text{absorption at } S_m | \mathcal{B}(0) = S_i\}$. We then have that $h_0 = 1 - h_m = 0$ and $h_{i+1} - h_i = \frac{P(S_i, S_{i-1})}{P(S_i, S_{i+1})}(h_i - h_{i-1})$, for $i = 1, \dots, m - 1$. In particular, $h_{i+1} - h_i = 2h_1$, for all $i = 1, \dots, m - 2$ and $h_m - h_{m-1} = h_1$. Finally, we have that $1 = \sum_{i=0}^{m-1} (h_{i+1} - h_i) = 2(m - 1)h_1$, which completes the proof of the Lemma. \square

For the proof of the main Theorem in this subsection, we define stochastic processes \mathcal{W}' and \mathcal{C}' just as the processes \mathcal{W} and \mathcal{C} from Subsection 4.2, but concerning the clique K_{n_1} . In particular, these processes take into account only transitions involving either edges that belong to the clique or the directed edge (u, v) . We note that Lemma 3 continues to hold for the modified process \mathcal{C}' with only one modification (because of the existence of the edge $\{v, u\}$) concerning its second part: T must satisfy $\mathcal{C}'|_T = 0$ and $\mathcal{C}'|_{T-1} > 0$ in order to be able to deduce that K_{n_1} has only vertices of type g .⁵ Furthermore, because of the existence of the directed edge (u, v) , the domination Lemmas 4 and 5 become as follows:

Lemma 7 (Domination of \mathcal{W}') *Let α', β', κ' be integers from the set $\{1, \dots, n_1 - 1\}$, with $\alpha' < \beta'$. Let also $\mathcal{B}_{\mathcal{W}'}$ be a birth-death process, which has state space $\mathcal{S}_{\mathcal{B}_{\mathcal{W}'}} = \{S'_0, \dots, S'_{n_1}\}$, with S'_{n_1} an absorbing state and transition probability matrix P' , with $P'(S'_i, S'_{i+1}) = 1$*

⁵ If this is not the case, i.e. if only $\mathcal{C}'|_T = 0$ is given, then we can still claim that all vertices except v are of type g , whereas v can be of type g or b . However, this is not needed in our analysis.

for all $i \in \{0, \dots, \alpha'\} \cup \{\beta', \dots, n_1 - 1\}$, $\frac{P'(S'_i, S'_{i+1})}{P'(S'_i, S'_{i-1})} = \frac{2\kappa'+1}{\alpha'}$ for all $i \in \{\alpha' + 1, \dots, \beta' - 1\}$ and $P'(S'_i, S'_i) = \Pr(\mathcal{W}'(t) = i | \mathcal{W}'(t-1) = i)$, for all $t \geq 1$ and for all $i \in \{\alpha' + 1, \dots, \beta' - 1\}$. Then, given that the vertices of type r in K_{n_1} are at most κ' , the process \mathcal{W}' is stochastically dominated by $\mathcal{B}_{\mathcal{W}'}$ in the following sense: $\Pr(\mathcal{W}'(t) > x | \mathcal{W}'(0) = 0) \leq \Pr(\mathcal{B}_{\mathcal{W}'}(t) \in \cup_{y>x} S'_y | \mathcal{B}_{\mathcal{W}'}(0) = S'_0)$, for any time t and $x \in \{0, \dots, n_1\}$.

Proof It suffices to show that, for any $t \geq 1$ and for any $i \in \{0, \dots, n_1 - 1\}$,

$$\begin{aligned} & \Pr(\mathcal{W}'(t) = i + 1 | \mathcal{W}'(t-1) = i) \\ & \leq \Pr(\mathcal{B}_{\mathcal{W}'}(t) = S'_{i+1} | \mathcal{B}_{\mathcal{W}'}(t-1) = S'_i). \end{aligned}$$

This is trivially true for all $i \in \{0, \dots, \alpha'\} \cup \{\beta', \dots, n_1 - 1\}$, because the right hand side of the above inequality is 1. For $i \in \{\alpha' + 1, \dots, \beta' - 1\}$, given that $\mathcal{W}'(t-1) = i$ (i.e. there are exactly i blanks at time t), the probability that \mathcal{W}' increases by 1 in the next time step is equal to the probability that either a transition $g \rightarrow r$ or a transition $r \rightarrow g$ occurs either involving vertices that are both inside the clique, or involving u and v . If we denote by R'_{t-1} (resp. G'_{t-1}) the set of vertices of type r (resp. of type g) in the clique K_{n_1} at time $t-1$, then the above probability is at most $\frac{2|R'_{t-1}| + |G'_{t-1}| + 1}{n_1(n_1-1)+1}$. On the other hand, the probability that \mathcal{W}' decreases by 1 in the next time step is equal to the probability that either a transition $g \rightarrow b$ or a transition $r \rightarrow b$ occurs either involving vertices that are both inside the clique, or involving the edge (u, v) . Given that $\mathcal{W}'(t-1) = i$, this probability is at least $\frac{i(|R'_{t-1}| + |G'_{t-1}|)}{n_1(n_1-1)+1}$. But then

$$\begin{aligned} \frac{\Pr(\mathcal{W}'(t) = i + 1 | \mathcal{W}'(t-1) = i)}{\Pr(\mathcal{W}'(t) = i - 1 | \mathcal{W}'(t-1) = i)} & \leq \frac{2|R'_{t-1}| + |G'_{t-1}| + 1}{i(|R'_{t-1}| + |G'_{t-1}|)} \\ & \leq \frac{2|R'_{t-1}| + 1}{i}. \end{aligned}$$

By assumption, this is at most $\frac{2\kappa'+1}{\alpha'}$, which combined with the fact that $P'(S'_i, S'_i) = \Pr(\mathcal{W}'(t) = i | \mathcal{W}'(t-1) = i)$, for all $i \in \{\alpha' + 1, \dots, \beta' - 1\}$, concludes the proof of Lemma 7. \square

Lemma 8 (Domination of \mathcal{C}') *Let β', κ' be positive integers, with $\beta' + \kappa' < n_1$. Let also $\mathcal{B}_{\mathcal{C}'}$ be a birth-death process, which has state space $\mathcal{S}_{\mathcal{B}_{\mathcal{C}'}} = \{T'_0, \dots, T'_{n_1}\}$, with a reflecting barrier at T'_0 , an absorbing state T'_{n_1} and transition probability matrix Q' , with $Q'(T'_i, T'_{i+1}) = 1$ for all $i \in \{0\} \cup \{\kappa', \dots, n_1 - 1\}$, and $\frac{Q'(T'_i, T'_{i+1})}{Q'(T'_i, T'_{i-1})} = \left(1 + \frac{1}{n_1 - \beta' - \kappa'}\right) \frac{\kappa'+1}{n_1 - \beta' - \kappa'}$, for all $i \in \{1, \dots, \kappa' - 1\}$, and $Q'(T'_i, T'_i) = \Pr(\mathcal{C}'(\tau) = i | \mathcal{C}'(\tau-1) = i)$, for all $\tau \geq 1$ and for all $i \in \{1, \dots, \kappa' - 1\}$. Then, given that the vertices of type b are at most β' , the process \mathcal{C}' is stochastically dominated by $\mathcal{B}_{\mathcal{C}'}$*

in the following sense: $\Pr(\mathcal{C}'(\tau) > x | \mathcal{C}'(0) = |R'_0|) \leq \Pr(\mathcal{B}_{\mathcal{C}'}(\tau) \in \cup_{y>x} T'_x | \mathcal{B}_{\mathcal{C}'}(0) = |R'_0|)$, for any time τ and $x \in \{0, \dots, n_1\}$, where $|R'_0|$ is the number of vertices of type r in K_{n_1} at time 0.

Proof It suffices to show that, for any $\tau \geq 1$ and for any $i \in \{1, \dots, n_1 - 1\}$,

$$\begin{aligned} & \Pr(\mathcal{C}'(\tau) = i + 1 | \mathcal{C}'(\tau - 1) = i) \\ & \leq \Pr(\mathcal{B}_{\mathcal{C}'}(\tau) = T'_{i+1} | \mathcal{B}_{\mathcal{C}'}(\tau - 1) = T'_i). \end{aligned}$$

This is trivially true for all $i \in \{0\} \cup \{\kappa', \dots, n_1 - 1\}$, because the right hand side of the above inequality is 1. For $i \in \{1, \dots, \kappa' - 1\}$, we apply the *principle of deferred decisions*. In particular, let t_1 (resp. t_2) be the time in the time scale of \mathcal{M} (i.e. the Markov chain describing the evolution of the protocol on H) that corresponds to the first (resp. second) transition of the τ -th transition pair in the definition of \mathcal{C}' . Given that $\mathcal{C}'(\tau - 1) = i$, the probability that \mathcal{C}' increases by 1 in the next time step in the time scale of \mathcal{C}' is equal to the probability that the τ -th transition pair is $(r \rightarrow g, r \rightarrow b)$. If we denote by R'_t (resp. G'_t) the set of vertices of type r (resp. of type g) in the clique K_{n_1} at time t , then the above probability is at most $\frac{|R'_{t_1}| |G'_{t_1}| + 1}{2|R'_{t_1}| |G'_{t_1}| + 1} \frac{|R'_{t_2}| (n_1 - |R'_{t_2}| - |G'_{t_2}|) + 1}{(|R'_{t_2}| + |G'_{t_2}|)(n_1 - |R'_{t_2}| - |G'_{t_2}|) + 1}$. On the other hand, given that $\mathcal{C}'(\tau - 1) = i$, the probability that \mathcal{C}' decreases by 1 in the next time step is equal to the probability that the τ -th transition pair in the definition of \mathcal{C}' is $(g \rightarrow r, g \rightarrow b)$, which is at least $\frac{|G'_{t_1}| |R'_{t_1}|}{2|R'_{t_1}| |G'_{t_1}| + 1} \frac{|G'_{t_2}| (n_1 - |R'_{t_2}| - |G'_{t_2}|)}{(|R'_{t_2}| + |G'_{t_2}|)(n_1 - |R'_{t_2}| - |G'_{t_2}|) + 1}$. But then

$$\begin{aligned} & \frac{\Pr(\mathcal{C}'(\tau) = i + 1 | \mathcal{C}'(\tau - 1) = i, t_1, t_2)}{\Pr(\mathcal{C}'(\tau) = i - 1 | \mathcal{C}'(\tau - 1) = i, t_1, t_2)} \\ & \leq \frac{|R'_{t_1}| |G'_{t_1}| + 1}{|R'_{t_1}| |G'_{t_1}|} \frac{|R'_{t_2}| (n_1 - |R'_{t_2}| - |G'_{t_2}|) + 1}{|G'_{t_2}| (n_1 - |R'_{t_2}| - |G'_{t_2}|)} \quad (13) \\ & \leq \left(1 + \frac{1}{|G'_{t_1}|}\right) \left(\frac{|R'_{t_2}|}{|G'_{t_2}|} + \frac{1}{|G'_{t_2}|}\right). \end{aligned}$$

We note here that we do not need to worry about division with 0 in the above fractions. In particular, by the assumption $\beta' + \kappa' < n_1$, we have that $|G'_{t_1}| |G'_{t_2}| \neq 0$. Furthermore, by the definition of the time t_2 , there must exist at least one vertex of type b at time t_2 , which implies that $n_1 - |R'_{t_2}| - |G'_{t_2}| > 0$. Finally, we do not need to worry about $|R'_{t_1}| = 0$ for the following reason: by the definition of \mathcal{C}' , having no vertices of type r and $\mathcal{C}'(\tau - 1) \geq 1$ at the same time would mean that there is a transition $g \rightarrow r$ that needs to be paired before any other transition of the form $r \rightarrow g$ and thus the fraction in the left hand side of (13) would be 0 (which is smaller than the bound given in the statement of the Lemma).

Since $i \leq \kappa' - 1$ by assumption, it follows by the first part of Lemma 3 that $|R'_{t_2}| < \kappa'$. Therefore, we have that $\frac{\Pr(\mathcal{C}'(\tau) = i + 1 | \mathcal{C}'(\tau - 1) = i)}{\Pr(\mathcal{C}'(\tau) = i - 1 | \mathcal{C}'(\tau - 1) = i)} \leq \left(1 + \frac{1}{n_1 - \beta' - \kappa'}\right) \frac{\kappa' + 1}{n_1 - \beta' - \kappa'}$, which, in combination with the fact that $Q'(T'_i, T'_i) = \Pr(\mathcal{C}'(\tau) = i | \mathcal{C}'(\tau - 1) = i)$, for all $\tau \geq 1$ and for all $i \in \{1, \dots, \kappa' - 1\}$, concludes the proof of Lemma 8. \square

We are now ready to prove the main result of this subsection.

Theorem 4 *Consider a lollipop graph on n vertices, which consists of a complete graph K_{n_1} on $n_1 \geq 100 \ln n$ vertices, among which vertex v is connected to the left-most vertex u of a line graph L_{n_2} on $n_2 = n - n_1$ vertices. Suppose that initially vertex v and all vertices in L_{n_2} are of type r , while all vertices in $K_{n_1} \setminus \{v\}$ are of type g . Then with probability at least $1 - 16n \left(\frac{7}{6.1}\right)^{-\frac{n_1}{13}} = 1 - o(1)$, eventually only vertices of type g will remain in the graph.*

Proof For the proof, we modify accordingly the proof of the main Theorem in Subsection 4.2. We begin by setting $\alpha' = \frac{n_1}{3}$, $\beta' = \frac{n_2}{2}$ and $\kappa' = \frac{n_1}{7}$. In that case, the forward bias of the birth-death process $\mathcal{B}_{\mathcal{W}'}$ at states $\{S'_{\alpha'+1}, S'_{\alpha'+2}, \dots, S'_{\beta'-1}\}$ is at most $r_{\mathcal{W}'} \stackrel{def}{=} \frac{6.1}{7}$ (taking also into account that $n_1 \rightarrow \infty$) and the forward bias of the birth-death process $\mathcal{B}_{\mathcal{C}'}$ at states $\{T'_1, T'_2, \dots, T'_{\kappa'-1}\}$ is at most $r_{\mathcal{C}'} \stackrel{def}{=} \frac{2.1}{5}$.

We will also denote by R'_t (resp. G'_t) the set of vertices of type r (resp. of type g) in the clique K_{n_1} at time t . For simplicity, the initial configuration will be denoted by \mathcal{E} .

We first provide a lower bound on the probability that the clique K_{n_1} reaches a configuration where all its vertices (i.e. including v) are of type g . To this end, we define the following events, which are similar to the events A_1 and A_2 used in the proof of the main Theorem in Subsection 4.2 (only now they are defined for the clique K_{n_1}):

- (i) Γ_1 is the event that \mathcal{C}' reaches κ before the number of vertices of type b reaches β' .
- (ii) Γ_2 is the event that the number of vertices of type b reaches β' before \mathcal{C}' reaches either 0 or κ' .

Notice now that, by the second part of Lemma 3 (see also the discussion before Lemma 7 in this subsection), whenever we have $\mathcal{C}'_{|t} = 0$ and $\mathcal{C}'_{|t-1} > 0$, all vertices in K_{n_1} must be of type g . In view of this, we are interested in lower bounding the probability that \mathcal{C}' reaches 0 before either Γ_1 or Γ_2 happens, given that we initially start with a configuration where all vertices in $K_{n_1} \setminus \{v\}$ are of type g and

v is of type r . Equivalently, we are interested in upper bounding the probability that $\Gamma_1 \cup \Gamma_2$ happen before \mathcal{C}' reaches 0. But for the latter, we can follow the same steps as in the proof of Theorem 3. In particular, using $n_1, \alpha', \beta', \kappa', r_{\mathcal{W}'}, r_{\mathcal{C}'}$, instead of $n, \alpha, \beta, \kappa, r_{\mathcal{W}}, r_{\mathcal{C}}$ respectively, setting $k'_0 = k'_1 = \left(\frac{1}{r_{\mathcal{W}'}}\right)^{\frac{n_1}{12}}$, instead of k_0, k_1 , and setting $\epsilon = \frac{1}{n_1}$ (i.e. initially we start with a single vertex of type r in K_{n_1}) we get the following equivalent to equation (12):

$$\begin{aligned} & \Pr\{\Gamma_1 \cup \Gamma_2 \text{ happens before } \mathcal{C}' \text{ reaches } 0 \mid \mathcal{E}\} \\ & \leq \frac{\frac{1}{r_{\mathcal{C}'}} - 1}{\left(\frac{1}{r_{\mathcal{C}'}}\right)^{\frac{n_1}{7}} - 1} + k'_0 \left(\frac{1}{r_{\mathcal{W}'}}\right)^{-\frac{n_1}{6}} + \frac{O(n_1^5)}{k'_1} \\ & \leq \left(\frac{7}{6.1}\right)^{-\frac{n_1}{13}} + \left(\frac{7}{6.1}\right)^{-\frac{n_1}{12}} + \frac{2.9}{2.1} \frac{1}{\left(\frac{5}{2.1}\right)^{\frac{n_1}{7}} - 1} \\ & \leq 8 \left(\frac{7}{6.1}\right)^{-\frac{n_1}{13}} \stackrel{\text{def}}{=} \phi_K. \end{aligned}$$

By Lemma 6, we also have that, given that there are no vertices of type r in the clique K_{n_1} and that v is of type g , the probability that eventually only vertices of type g remain on the line L_{n_2} is at least $\frac{1}{2n} \stackrel{\text{def}}{=} \phi_L$. Indeed, this follows from the observation that $n \geq n_2$ and that the existence of any other vertex of type g or b inside the line can only increase the probability that the type of the leftmost vertex wins without v ever changing its type from g to b .

Using all the above we then have that, given \mathcal{E} , the probability that eventually only vertices of type r remain in the lollipop graph at most the probability that at some point $\Gamma_1 \cup \Gamma_2$ happens before \mathcal{C}' reaches 0, which is at most

$$\sum_{i=0}^{\infty} \phi_K (1 - \phi_L)^i = \frac{\phi_K}{\phi_L} = 16n \left(\frac{7}{6.1}\right)^{-\frac{n_1}{13}},$$

which is $o(1)$ (i.e. it goes to 0 as n goes to infinity) provided $n_1 \geq 100 \ln n$. This completes the proof of the Theorem. \square

4.4 Expected Exponential time to absorption

In this subsection we show that there are graphs in which the majority protocol of [4] needs an expected exponential time to reach consensus. We prove this by analyzing the expected time to reach consensus in the case where the interaction graph consists of 2 disjoint cliques on n_1 and n_2 vertices each (i.e. $n = n_1 + n_2$) and a single edge between them, as shown in Figure 3. Initially, all vertices in the K_{n_1} clique are of type g , while all vertices in the K_{n_2} clique are of type r .




Fig. 3 An interaction graph G with 2 disjoint K_{n_1} and K_{n_2} , connected with a single edge u, v .




Fig. 4 The graph H consisting of a clique K_{n_1} with a directed edge (u, v) , where only u does not belong to the clique.

We will assume w.l.o.g. that $n_1 \leq n_2$ and that n_1 is an increasing function of the total number of vertices (i.e. $n_1 \rightarrow \infty$ as n goes to infinity). In order to provide a lower bound on the time after which only one type of vertices remains (e.g. either r or g), we will use a directed graph H consisting of a single clique K_{n_1} and a vertex v outside the clique, which is connected to a vertex u of the clique with a directed edge (v, u) .

We now prove the following:

Lemma 9 *Let T_H be the number of steps that the majority protocol of [4] needs to reach consensus when the underlying graph is H and initially all vertices are of type g except for u which is of type r . Then the mean value of T_H is at least exponential in n_1 .*

Proof It is evident from the definition of the protocol of [4] that eventually only vertices of type r will remain. Indeed, vertex u can never change its type (notice that the protocol of Angluin et al. is not symmetric), but it can affect the type of v , and through v the type of every vertex in the clique. In order to analyze the time needed to reach (forced) consensus, we define stochastic processes \mathcal{W}'' and \mathcal{C}'' just as the processes \mathcal{W} and \mathcal{C} from Subsection 4.2, but concerning the clique K_{n_1} . In particular, these processes take into account only transitions involving either edges that belong to the clique or the directed edge (u, v) . We note that the first part of Lemma 3 continues to hold for process \mathcal{C}'' . However, because of the existence of the directed edge (u, v) , the domination Lemmas 4 and 5 become as the next two Lemmas (their proofs are identical to the proofs of Lemmas 7 and 8 respectively).

Lemma 10 (Domination of \mathcal{W}'') *Let $\alpha'', \beta'', \kappa'' \in \{1, \dots, n_1 - 1\}$, with $\alpha'' < \beta''$. Let also $\mathcal{B}_{\mathcal{W}''}$ be a*

birth-death process, which has state space $\mathcal{S}_{\mathcal{B}_{\mathcal{W}''}} = \{S_0'', \dots, S_{n_1}''\}$, with S_{n_1}'' an absorbing state and transition probability matrix P'' , with $P''(S_i'', S_{i+1}'') = 1$ for all $i \in \{0, \dots, \alpha''\} \cup \{\beta'', \dots, n_1 - 1\}$, $\frac{P''(S_i'', S_{i+1}'')}{P''(S_i'', S_{i-1}'')} = \frac{2\kappa''+1}{\alpha''}$ for all $i \in \{\alpha''+1, \dots, \beta''-1\}$ and $P''(S_i'', S_i'') = \Pr(\mathcal{W}''(t) = i | \mathcal{W}''(t-1) = i)$, for all $t \geq 1$ and for all $i \in \{\alpha''+1, \dots, \beta''-1\}$. Then, given that the vertices of type r in K_{n_1} are at most κ'' , the process \mathcal{W}'' is stochastically dominated by $\mathcal{B}_{\mathcal{W}''}$ in the following sense: $\Pr(\mathcal{W}''(t) > x | \mathcal{W}''(0) = 0) \leq \Pr(\mathcal{B}_{\mathcal{W}''}(t) \in \cup_{y>x} S_y'' | \mathcal{B}_{\mathcal{W}''}(0) = S_0'')$, for any time t and $x \in \{0, \dots, n_1\}$.

Lemma 11 (Domination of \mathcal{C}'') Let β'', κ'' be positive integers, with $\beta'' + \kappa'' < n_1$. Let also $\mathcal{B}_{\mathcal{C}''}$ be a birth-death process, which has state space $\mathcal{S}_{\mathcal{B}_{\mathcal{C}''}} = \{T_0'', \dots, T_{n_1}''\}$, with a reflecting barrier at T_0'' , an absorbing state T_{n_1}'' and transition probability matrix Q'' , with $Q''(T_i'', T_{i+1}'') = 1$ for all $i \in \{0\} \cup \{\kappa'', \dots, n_1 - 1\}$, $\frac{Q''(T_i'', T_{i+1}'')}{Q''(T_i'', T_{i-1}'')} = \left(1 + \frac{1}{n_1 - \beta'' - \kappa''}\right) \frac{\kappa''+1}{n_1 - \beta'' - \kappa''}$ for all $i \in \{1, \dots, \kappa'' - 1\}$ and $Q''(T_i'', T_i'') = \Pr(\mathcal{C}''(\tau) = i | \mathcal{C}''(\tau-1) = i)$, for all $\tau \geq 1$ and for all $i \in \{1, \dots, \kappa'' - 1\}$. Then, given that the vertices of type b are at most β'' , the process \mathcal{C}'' is stochastically dominated by $\mathcal{B}_{\mathcal{C}''}$ in the following sense: $\Pr(\mathcal{C}''(\tau) > x | \mathcal{C}''(0) = |R_0''|) \leq \Pr(\mathcal{B}_{\mathcal{C}''}(\tau) \in \cup_{y>x} T_y'' | \mathcal{B}_{\mathcal{C}''}(0) = |R_0''|)$, for any time τ and $x \in \{0, \dots, n_1\}$, where $|R_0''|$ is the number of vertices of type r in K_{n_1} at time 0.

Continuing with the proof of Lemma 9, we note that the expected time needed for the protocol to reach consensus (i.e. reach a configuration where all vertices are of type r) is at least the time that it needs to reach either β'' vertices of type b or κ'' vertices of type r .

We now set $\alpha'' = \frac{n_1}{3}$, $\beta'' = \frac{n_1}{2}$ and $\kappa'' = \frac{n_1}{7}$. In that case, the forward bias of the birth-death process $\mathcal{B}_{\mathcal{W}''}$ at states $\{S_{\alpha''+1}'', S_{\alpha''+2}'', \dots, S_{\beta''-1}''\}$ is at most $r_{\mathcal{W}''} \stackrel{\text{def}}{=} \frac{6.1}{7} < 1$ (taking also into account that $n_1 \rightarrow \infty$) and the forward bias of the birth-death process $\mathcal{B}_{\mathcal{C}''}$ at states $\{T_1'', T_2'', \dots, T_{\kappa''-1}''\}$ is at most $r_{\mathcal{C}''} \stackrel{\text{def}}{=} \frac{2.1}{5} < 1$.

In particular, by Lemma 2(b) and by the domination Lemma 10 for process \mathcal{W}'' , the expected number of steps needed in order to reach β'' vertices of type b (given that initially there are no vertices of type b) is $\Omega\left(\frac{1}{\frac{1}{r_{\mathcal{W}''}} - 1} \left(\frac{1}{r_{\mathcal{W}''}}\right)^{\frac{n_1}{6}}\right) = \Omega\left(\left(\frac{7}{6.1}\right)^{\frac{n_1}{6}}\right)$. Similarly, by Lemma 2(b) and by the domination Lemma 11 for process \mathcal{C}'' , the expected number of steps in the time scale of \mathcal{C}'' needed for \mathcal{C}'' to reach κ'' (given that initially only vertex u is of type r) is at least $\Omega\left(\frac{1}{\frac{1}{r_{\mathcal{C}''}} - 1} \left(\frac{1}{r_{\mathcal{C}''}}\right)^{\frac{n_1}{7}}\right) = \Omega\left(\left(\frac{5}{2.1}\right)^{\frac{n_1}{7}}\right)$. By the first part of Lemma 3, this is at

least the expected time needed for the number of vertices of type r to reach κ'' . The proof of Lemma 9 is completed by noting that the expected time needed for the protocol to reach the configuration where all vertices are of type r is at least the minimum of these two exponential quantities. \square

Using Lemma 9, we can prove the main result of this subsection.

Theorem 5 Let G be an interaction graph on n vertices, which consists of 2 disjoint cliques on n_1 and n_2 vertices each and a single edge between them (see Figure 3). Suppose that initially, all vertices in the K_{n_1} clique are of type g , while all vertices in the K_{n_2} clique are of type r . Then the majority protocol of [4] needs an expected exponential number of steps to reach consensus.

Proof Let T_G be the number of steps needed for the protocol to reach consensus, starting from the initial configuration described in the Theorem. Let also A be the event that the initial majority eventually wins. We then have that

$$\begin{aligned} \mathbb{E}[T_G] &= \Pr(A)\mathbb{E}[T_G|A] + (1 - \Pr(A))\mathbb{E}[T_G|\bar{A}] \\ &\geq \Pr(A)\mathbb{E}[T_G|A]. \end{aligned}$$

Assume now w.l.o.g. that $n_1 \leq n_2$. Then, by a similar argument to the derivation of equation (3) in the proof of Theorem 2, the probability that the initial majority wins (i.e. eventually only vertices of type r remain), is at least $\frac{1}{2}$. The proof is concluded by noting that if T_H is defined as in Lemma 9 then $\mathbb{E}[T_G|A] \geq \mathbb{E}[T_H]$. \square

References

1. Alistarh, D., Gelashvili, R., Vojnović, M.: Fast and exact majority in population protocols. In: Proceedings of the 34th ACM Symposium on Principles of Distributed Computing (PODC), pp. 47–56 (2009)
2. Angluin, D., Aspnes, J., Diamadi, Z., Fischer, M.J., Peralta, R.: Computation in networks of passively mobile finite-state sensors. Distributed Computing **18**, 235–253 (2006)
3. Angluin, D., Aspnes, J., Eisenstat, D.: Fast computation by population protocols with a leader. Distributed Computing **21**(2), 183–199 (2007)
4. Angluin, D., Aspnes, J., Eisenstat, D.: A simple population protocol for fast robust approximate majority. Distributed Computing **21**(2), 87–102 (2008)
5. Aspnes, J., Ruppert, E.: An introduction to population protocols. In: B. Garbinato, H. Miranda, L. Rodrigues (eds.) Middleware for Network Eccentric and Mobile Applications, pp. 97–120. Springer-Verlag (2009)
6. Bénézit, F., Thiran, P., Vetterli, M.: Interval consensus: From quantized gossip to voting. In: Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASP), pp. 3361–3364 (2009)

7. Cook, M., Soloveichik, D., Winfree, E., Bruck, J.: Programmability of chemical reaction networks. In: A. Condon, D. Harel, J.N. Kok, A. Salomaa, E. Winfree (eds.) *Algorithmic Bioprocesses*, Natural Computing Series, pp. 543–584. Springer Berlin Heidelberg (2009)
8. De Marco, G., Pelc, A.: Randomized algorithms for determining the majority on graphs. *Combinatorics, Probability and Computing* **15**(6), 823–834 (2006)
9. DeGroot, M.H.: Reaching a consensus. *Journal of the American Statistical Association* **69**(345), 118–121 (1974)
10. Draief, M., Vojnović, M.: Convergence speed of binary interval consensus. *SIAM Journal on Control and Optimization* **50**(3), 1087–1109 (2012)
11. Fischer, M., Jiang, H.: Self-stabilizing leader election in networks of finite-state anonymous agents. In: *Proceedings of the 10th International Conference on Principles of Distributed Systems (OPODIS)* (2006)
12. Holley, R.A., Liggett, T.M.: Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability* **3**, 643–663 (1975)
13. Jukna, S.: *Extremal Combinatorics with Applications to Computer Science*. Springer-Verlag (2011)
14. Kearns, M., Tan, J.: Biased voting and the democratic primary problem. In: *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE)*, pp. 639–652 (2008)
15. Kurtz, T.G.: *Approximation of Population Processes*. Society for Industrial and Applied Mathematics (1987)
16. Lamport, L., Shostak, R., Pease, M.: The byzantine generals problem. *ACM Trans. Program. Lang. Syst.* **4**(3), 382–401 (1982)
17. Liggett, T.M.: *Interacting Particle Systems*. Springer-Verlag (2004)
18. Merzios, G.B., Nikolettseas, S.E., Raptopoulos, C.L., Spirakis, P.G.: Determining majority in networks with local interactions and very small local memory. In: *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming*, vol. 1, pp. 871–882 (2014)
19. Mizrahi, A.: Majority vote and monopolies in social networks. Master’s thesis, Department of Communication Systems Engineering, Faculty of Engineering, Ben-Gurion University of the Negev (2013)
20. Moran, P.A.P.: Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society* **54**(1), 60–71 (1958)
21. Mossel, E., Neeman, J., Tamuz, O.: Majority dynamics and aggregation of information in social networks. *Autonomous Agents and Multi-Agent Systems* **28**(3), 408–429 (2014)
22. Norris, J.: *Markov Chains*. Cambridge University Press (1998)
23. Nowak, M.A.: *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press (2006)
24. Perron, E., Vasudevan, D., Vojnović, M.: Using three states for binary consensus on complete graphs. In: *Proceedings of the 28th IEEE International Conference on Computer Communications (INFOCOM)*, pp. 2527–2535 (2009)
25. Preparata, F.P., Metze, G., Chien, R.T.: On the connection assignment problem of diagnosable systems. *IEEE Transactions on Electronic Computers* **16**, 848–854 (1967)
26. Ross, S.R.: *Stochastic Processes*. John Wiley & Sons, Inc. (1995)
27. Saks, M., Werman, M.: On computing majority by comparisons. *Combinatorica* **11**(4), 383–387 (1991)