Multi-Agent Flag Coordination Games

David Kohan Marzagão*
King's College London, UK
david.kohan@kcl.ac.uk

Nicolás Rivera
King's College London, UK
nicolas.rivera@kcl.ac.uk

Colin Cooper
King's College London, UK
colin.cooper@kcl.ac.uk

Peter McBurney
King's College London, UK
peter.mcburney@kcl.ac.uk

Kathleen Steinhöfel
King's College London, UK
kathleen.steinhofel@kcl.ac.uk

ABSTRACT

Many multi-agent coordination problems can be understood as autonomous local choices between a finite set of options, with each local choice undertaken simultaneously without explicit coordination between decision-makers, and with a shared goal of achieving a desired global state or states. Examples of such problems include classic consensus problems between nodes in a distributed computer network and the adoption of competing technology standards. We model such problems as a multi-round game between agents having flags of different colours to represent the finite choice options, and all agents seeking to achieve global patterns of colours through a succession of local colour-selection choices.

We generalise and formalise the problem, proving results for the probabilities of achievement of common desired global states when these games are undertaken on bipartite graphs, extending known results for non-bipartite graphs. We also calculate probabilities for the game entering infinite cycles of non-convergence. In addition, we present a game-theoretic approach to the problem that has a mixed-strategy Nash equilibrium where two players can simultaneously flip the colour of one of the opponent's nodes in the bipartite graph before or during a flag-coordination game.

Keywords

Consensus Protocols, Graph Colouring, Flag Coordination, Multi-agent Coordination

1. INTRODUCTION

Many multi-agent coordination problems involve a collection of agents choosing autonomously from the same finite set of options using only local information, while sharing a common desire for a global state. For example, users of a new technology choosing between alternative technical standards each face the same choice of possible options, but make their choices without necessarily knowing the choices of others. In the case of *network goods* [14], the utilities of each option to any one user depend on the choices made by the other users; in the classic example, a fax machine is only of value to one company if the organisations with which that company communicates also have fax machines. Hence, potential adopters may choose the option they believe others will choose [15]. Even for non-technology products, such as clothes and food, consumers may gain additional benefits from purchasing products or services which they believe have been chosen (or not chosen) by other consumers, over and above any perceived benefits of the good or service itself.

In these cases, agents may wish to all adopt the same choice as one another, so that the desired shared global state is one of *consensus*. In other cases, the global state may be a different pattern, for example, a sequence of alternating states. For instance, in a robot bucket brigade, each robot in a line would need to be either in a giving state or in a receiving state at each time step, and in the complementary state to each of its neighbours at that time step. At each subsequent time step, each robot would need to switch to the other state.

We can model such situations as an abstract multi-agent game of flag-colouring, where the different flag colours represent the different decision-options each agent faces. While there are applications where the desired global state of the system needs to be achieved in a single step [10], we consider only cases where the agents proceed in a sequence of rounds, making individual choices simultaneously at each step. If at any step, a desired global state is achieved, the game ends. Otherwise, it continues.

Because there are many possible variations on this general situation, in this paper we make certain assumptions to fix ideas. These assumptions are:

1. We assume a finite set of autonomous agents, with a shared clock, and with each empowered to decide between a finite set of decision options at each successive time-step. These options are the same for every agent. Decisions are made synchronously, at successive time steps. For simplicity, the decision options are represented by flags of different colours.

2. Agents are connected via a network, and after each time-step each agent is able to see the decisions made by its immediate neighbours, i.e., those agents to whom it is directly linked. No agent is able to see beyond that. Agents do not communicate in any other way with one another.

3. Agents do know the decision-option they themselves choose at each time-step but they are not assumed to have any memory of previous choices, of themselves or of other agents, or of previous global states.

*Suported by CNPq (206390/2014-9).

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

4. Agents are assumed to know the network topology and their place within it.

5. Agents all share a desired set of global goal states (possibly just one state) for the collective set of agents. This set of shared global goal states could be, for example, consensus (all agents choose the same decision-option) or a global state in which no two connected agents have made the same choice (e.g., alternating flag colours).

6. We assume that, between one time-step and the next, agents are not informed whether or not their previous decisions achieved one of the desired goal states. Instead, in this work, the individual agent decision algorithms we consider result in the global goal states being stable. If and when a stable goal state is achieved, we say the sequential decision process ends. Otherwise, the process continues, possibly for ever.

7. Agents are assumed to be well-intentioned (i.e., not malicious nor whimsical), and bug-free.

EXAMPLE 1 (MOTIVATING EXAMPLE). *Consider a set of twenty agents playing a colouring game in a circle, with the initial configuration as shown in figure 1. At each round, each node chooses one neighbour at random and copies its colour. For example, node v_5, for the next round, is choosing blue with 50% probability, and red otherwise. At the same time, node v_7 will be turning red with probability 1. All changes are made synchronously. We say that the game ends either when there is only one colour left or when the game is trapped in an infinite loop (if odd nodes are coloured with one colour and the even ones with another). In these conditions, which of the three colours is most likely to win this game?*[1]

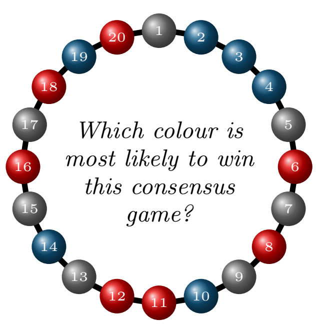

Figure 1: Consensus Game on a Cycle with 3 Colours.

In this paper we articulate a formal model for a flag-colouring game based on these assumptions, with the purpose of answering the following questions:

1. Given a defined network topology, a defined decision algorithm, a defined global goal state, and a random initial state, what is the probability that the sequential decision process will converge to a global goal state?

2. Given a defined network topology, a defined decision algorithm, a defined set of global goal states, and a random initial state, what is the expected number of decision rounds (time-steps) to reach a global goal state?

3. Given a defined network topology, a defined decision algorithm, a defined global goal state, and a random initial state, what is the probability that the sequential decision process will enter an infinite cycle that does not converge to a global goal state (i.e., an infinite cycle of non-convergence)?

4. What is the influence of the network topology on the probability and expected duration of achieving a global goal state?

5. What is the influence of the decision algorithm on the probability and expected duration of achieving a global goal state?

Building on prior work on general graphs by Hassin and Peleg [7], we present results for the first three questions in this paper for bipartite graphs. These are graphs where the nodes may be divided naturally into two mutually-exclusive types, for example, buyers and sellers in an online marketplace. The solution to the problem in Example 1 is interesting precisely because the cycle in that problem is a bipartite graph. To a limited extent, we also explore Questions 4 and 5. These questions will be the subject of future work.

In addition to this analysis, we present a game-theoretic approach to the situation where two players can simultaneously flip the colour of one of their opponents' nodes in the bipartite graph before or during a flag-coordination game, in Section 3. This game has a mixed-strategy Nash equilibrium.

2. FLAG COORDINATION GAMES

2.1 Notation and Basic Definitions

We let G denote a finite, undirected graph $G = (V, E)$ in which V is the set of nodes of the graph and E stands for the set of edges of G. If G is a bipartite graph, then we denote it as $G = (U, W, E)$ in which U and W represent the two partitions of the nodes in G, this is, $U \cup W = V$, $U \cap W = \emptyset$ and $\forall \{v_1, v_2\} \in E$, either $v_1 \in U$ and $v_2 \in W$ or $v_2 \in U$ and $v_1 \in W$.

We define $\mathcal{C} = X^V$ the set of functions $s : V \to X$ that colour the graph $G = (V, E)$, where X is a given set of colours. We say that s is a proper colouring of the nodes in G if no neighbouring vertices are assigned the same colour under s.

2.2 Flag Coordination Games

We now propose an abstract framework in which we can study coordination in multi-agent systems in a broader way. A flag coordination game can be seen as a multi-round game between agents having flags of different colours to represent the finite choice options, and all agents seeking to achieve a global state or states of colours through a succession of local colour-selection choices. We start by giving a general definition of such games.

DEFINITION 1 (FLAG COORDINATION GAME). *Let G, $G = (V, E)$, be a graph and $\beta : V \to \mathbb{P}(X)$ be a function that associates each vertex to a particular set of colours, or flags, at its disposal. Let also $\phi : V \to \mathbb{P}(V)$ be the function that determines the visibility of each node v, i.e., the set of vertices (and their labels) whose colours can be observed by v at at any time. Finally, the set $\mathcal{C} = X^V$ is the set of functions that colour V. Let $\Gamma = \{\gamma_1, \ldots, \gamma_n, \ \ldots\} \subset \mathcal{C}$ be the set of goal configurations of such game. Note that Γ is public and common to all nodes.*

[1]See Section 5 for the solution.

Let $s_0 \in \mathcal{C}$ be a configuration when the game starts. This game is played in turns and in each of these turns, nodes decide their own colour synchronous and independently (we also include 'no action' as a possible decision option). For each $v \in V$, there is an algorithm α_v, where α_v depends on $\beta(v)$ and $\phi(v)$, that makes a decision at each round. We call $\mathcal{A}$ the collection of all such algorithms.

*In sum, we define the tuple $\mathcal{F} = (G, X, \beta, \mathcal{C}, \Gamma, \mathcal{A}, \phi)$ as the **set of rules** of a flag coordination game, the pair $(\mathcal{F}, s_0)$ as the **initial configuration of a game**[2] and, finally, the infinite sequence of rounds $S = (s_0, \dots, s_k, \dots)$ as the **trace of a game**.*

Given a set of rules of a flag coordination game, we might be interested, for example, in the expected number of rounds until a goal is reached given an initial configuration, or even in the probability that a given game ends successfully: that it eventually reaches a configuration $\gamma \in \Gamma$. One particular example of such games is where the set Γ contains all the proper colourings of a given graph G given a set of colours X. For a given node $v \in V$, a simple possible algorithm α_v is for v to choose randomly among the least common colours of its neighbours. Note this particular algorithm might not be stable over proper colourings or efficient in terms of number of rounds for a goal to be achieved.

Example 2 (The Muddy Children Problem). *The commonly-studied Muddy Children Problem [2] can now be framed as a flag coordination game as long as we allow nodes to select the option of 'no action' in any round. We can have $G = (V, E)$ as, for instance, the complete graph with n nodes, where n is the number of children in the game. The set of colours is $X = \{mud, no\text{-}mud, mud\text{-}detected\}$. The initial configuration may have the nodes coloured with any of the two first colours, but we only allow the children the options of 'mud-detected' and 'no-action' (note that 'no-action' is not a colour, but the choice for the node to not change their current colour). We need to restrict the visibility of each agent to all other agents except themselves, so that: $\phi(v) = V \setminus \{v\}$, $\forall v \in V$.*

Our desired algorithm for v is 'wait' (i.e., take no action) until round i, where i is the number of 'mud' nodes that v can see. If no agent changes to 'mud-detected' until round i, then chose 'mud-detected' for round $i + 1$. In order to be consistent with our model, we have to define a public set of goal states Γ. Since we cannot simply give away the desired configuration to the nodes based on the number of 'mud'-coloured ones, we can define $\Gamma = \{\gamma \text{ s.t. } \gamma(v) \neq mud \; \forall v \in V\} \setminus \{all\ mud\text{-}detected\}$. This way, the set Γ does not give the nodes any new information as well as preventing them from arbitrarily choosing 'mud-detected' in the first round, because if they all do so they are trapped in the non-winning all-'mud-detected' state.

Finally, observe that consensus protocols in distributed systems can also be seen as flag coordination games. We now define a slightly broader class of consensus games, in which not only monochromatic goal states can be achieved.

Definition 2 (Generalised Consensus). *Consider $\mathcal{F} = (G, X, \beta, \mathcal{C}, \Gamma, \mathcal{A}, \phi)$ to be the set of rules of a flag coordination game where $X = \{x_0, \dots, x_{r-1}\}$ also, the set Γ has r elements, $\Gamma = \{\gamma_0, \gamma_1, \dots, \gamma_{(r-1)}\}$, such that, for a given pair (v, x), where $v \in V$ and $x \in X$, there is exactly one $\gamma \in \Gamma$ with $\gamma(v) = x$. We define $\beta(v) = X$ for all $v \in V$. The visibility of each vertex v is the set of neighbours of v, $\mathcal{N}(v)$. Finally, for each v, the algorithm α_v consists in choosing on round s_i a neighbour of v at random[3], say w then observing which $\gamma \in \Gamma$ is such that $\gamma(w) = s_i(w)$. We then define the value $s_{i+1}(v) = \gamma(v)$.*

The algorithm above is well defined because, for each pair (v, x), where $x = s(v)$, there is only one goal configuration in which v takes colour x. We use the term *generalised consensus* because, assuming the nodes know where they are and which other nodes they can see, they adhere to the winning configuration that the randomly chosen neighbour belongs to. In particular, if $\gamma_i(v) = i$, $\forall v \in V$ and $0 \leq i < r$, then we have a consensus problem in the usual way.

Definition 3 (Games on Bipartite Graphs). *Let us denote by $\mathcal{F}_b = (G, X, \beta, \mathcal{C}, \Gamma, \mathcal{A}, \phi)$ the rules of a generalised consensus flag coordination game played on a bipartite graph $G = (U, W, E)$, with $V = U \cup W$. We also define what is a monochromatic partition in a more broader way, in line with Definition 2: we say partition U is monochromatic if $\exists \gamma \in \Gamma$ such that $\forall u \in U$, $s_i(u) = \gamma(u)$. For short, we say that U is γ-monochromatic.*

2.3 Single-partition Games

In this section, we define single-partition games, games in which there is only one reachable winning configuration (Proposition 3). Alternatively, these games always have a non-randomising partition: a partition whose nodes have a deterministic behaviour.

In order to provide a motivation for the **split** function defined later in this section (see Definition 9) (and thus also for the definition of single-partition games), we show an interesting connection between annihilating random walks on cycles (see [6]) and flag coordination games. For other approaches on consensus and random walks on graphs, see [4].

Consider a flag coordination game $(\mathcal{F}_b, s_0)$ as in Definition 3, where $X = \{blue, red\}$ and G is not only bipartite but also a cycle. For simplicity, we assume the goal states are the standard consensus configurations: all-blue and all-red. In a given round s_i, we say that a vertex is a *non-randomising node* if it has deterministic behaviour, that is, both neighbours of the same colour (e.g., node v_1 in Figure 2). Otherwise, we have a *randomising node.* These nodes are going to chose *blue* or *red* with 50% chance each (e.g., node v_4 in Figure 2).

Independently, consider G a n-cycle, n even, and also $2k$ random walking particles each positioned in a different node of G. At each round of this game, each particle walk clockwise or counter-clockwise with probability 50% each. They all move synchronously. If two particles meet, both disappear. The game ends where there are no particles left. Note that particles that start within an odd distance between each other will never meet.

We claim that both games described above are equivalent. Given a flag coordination game $(\mathcal{F}_b, s_0)$, we generate an equivalent random walks game by positioning the random

[2]We sometimes use simply **game** to denote the pair $(\mathcal{F}, s_0)$.

[3]Here we can also add edge weights for more general definition. For simplicity, in this paper we consider that all edges have the same weight.

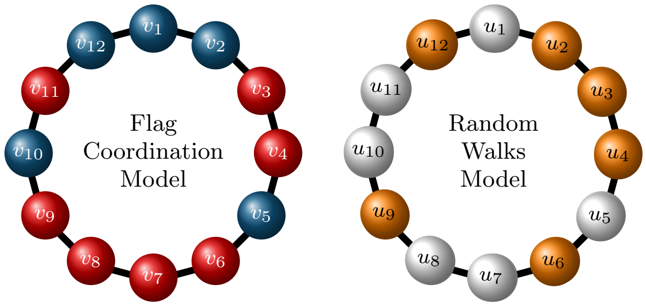

Figure 2: A Random Walks Game Generated from a Flag Coordination Game Considering Randomising Nodes.

walking particles at the nodes that are randomising nodes in $(\mathcal{F}_b, s_0)$. The expected duration of both games is the same. We can see an example in Figure 2. Orange nodes in the random walks model (the ones in which there are random walking particles) correspond to randomising nodes in the flag coordination model.

That connection, together with the fact that random walking particles in one partition never meet the ones in the other, motivates us to study each partition of the graph in a flag coordination game independently. Hence we state the following definition:

Definition 4 (Single-partition round and game). *Let $(\mathcal{F}_b, s_0)$ be a game as in Definition 3. We define a **single-partition round** of $(\mathcal{F}_b, s_0)$ as a round s_i in which the behaviour of all nodes in at least one partition of G is deterministic.*

*Moreover, we define a **single-partition game** as a game in which all rounds are single-partition rounds.*

Note that the corresponding random walks model of a single-partition flag coordination game has particles in one partition only. We will now show that if round s_0 is a single partition round, then $(\mathcal{F}_b, s_0)$ is a single-partition game.

Proposition 1. *Let $(\mathcal{F}_b, s_0)$ be a game as in Definition 3 where G is connected. If s_i is a single-partition round, then there is at least one partition, say U, that is monochromatic.*

Proof. Let W be the non-randomising partition on such round. Then, for each $w \in W$ all $u \in \mathcal{N}(w)$ are coloured according to the same $\gamma_w \in \Gamma$. Since G is connected, all such γ_w must coincide (otherwise there would be a w with neighbours in two different γ_w, which is not possible). We call that common colouring γ. Then, U is γ-monochromatic. □

Proposition 2. *A game that eventually reaches a single-partition round has all its subsequent rounds also single-partition. In particular, if s_0 is a single partition round, $(\mathcal{F}_b, s_0)$ is a single-partition game.*

Proof. Say U is γ-monochromatic partition in a single-partition round s_i. Then, in round s_{i+1}, all nodes in W will have been adhered to γ, thus W will be γ-monochromatic and so s_{i+1} is also a single-partition round. By induction, $(\mathcal{F}_b, s_0)$ is a single-partition game. □

Does this proposition imply anything regarding the possible final configurations of single-partition games? Indeed, the next proposition shows that there is only one possible winning state for such games.

Proposition 3 (Ending of Single-partit. Games). *Let $\gamma \in \Gamma$ be such that there is a γ-monochromatic partition on the initial round of a single-partition game $(\mathcal{F}_b, s_0)$. Then, in the case the game reaches consensus (it might not), such consensus must be γ.*

We now define a function that colours edges according to whether the colour of the nodes it connects belong to the same colouring ("black" edge) or not ("green" edge).

Definition 5 (Edge-colouring Function f). *Let $(\mathcal{F}_b, s_0)$ be a single partition game and $\mathcal{C}_E = \{green, black\}^E$ be the collection of all $2^{|E|}$ possible colourings for the edges in G. We define $f : \mathcal{C} \to \mathcal{C}_E$, $f(s) = r$ as the function that colour each edge $e = (u, w)$ in the following way:*

$$r(e) = \begin{cases} black, & \text{if } (\exists \gamma \in \Gamma)[s(u) = \gamma(u) \wedge s(w) = \gamma(w)] \\ green, & \text{otherwise} \end{cases}$$

In other words, an edge is green[4] if and only if the current colours of the nodes it links do not agree. Note that a game ends successfully when all edges are black. We can now give the probability of success based on the initial configuration of a single-partition game. We first formally define what we mean by "success".

Definition 6 (Winning Game). *We say that a game $(\mathcal{F}, s_0)$ is **successful**, or it is a **winning game**, if there exists $i \geq 0$ st $\forall j \geq i$, $s_j \in \Gamma$ where $(s_0, \ldots, s_k, \ldots)$ is the trace of such a game. We then define $P_{(\mathcal{F},s_0)}$ as the probability that $(\mathcal{F}, s_0)$ is a successful game More generally, we define $P^{\gamma}_{(\mathcal{F},s_0)}$ as the probability that the game reaches colouring γ. Since for the generalised consensus problem winning colourings are stable, we have $P_{(\mathcal{F},s_0)} = \sum_{\gamma \in \Gamma} P^{\gamma}_{(\mathcal{F},s_0)}$*

Definition 7 (Black Edges Counter). *Let G be a connected graph, $(\mathcal{F}_b, s_0)$ be a single partition game on G and $r_0 = f(s_0)$. We define $(Y_i)_{i>0}$ as the random variable that counts the number of black edges in r_i.*

Note that, for single-partition games on connected graphs, if $Y_0 = |E|$, then $s_0 \in \Gamma$ and therefore the game is certainly a winning game. On the other hand, if $Y_0 = 0$, then Y_1 is also zero and indeed $Y_i = 0$ for $i \geq 0$. We can show this by induction. Assume $Y_i = 0$. Then, there is one partition, say U, that is γ-monochromatic on round s_i. Therefore, W will be γ-monochromatic on round s_{i+1} and also no node in U will keep their colour, i.e., $s_{i+1}(u) \neq \gamma(u)$, because no node in U on round s_i have a neighbour in γ (since $Y_i = 0$). Thus, $Y_{i+1} = 0$ and, by induction, the game will never reach colouring γ.

Definition 8 (Duration of a Game). *For games of the form $(\mathcal{F}_b, s_0)$, we now define the duration d of the trace of a game being the smallest i such that $Y_i \in \{0, |E|\}$. In other words, we are considering both winning and losing games in our definition of duration. We define $D_{(\mathcal{F}_b,s_0)}$ to be the expected duration of a game with initial configuration $(\mathcal{F}_b, s_0)$.*

[4]For easier image reading, green edges are also dashed.

Theorem 1 (Probability of success). *Let $(\mathcal{F}_b, s_0)$ be a single partition game on a connected graph G. Let $\gamma \in \Gamma$ be the colouring for the monochromatic partition in $(\mathcal{F}_b, s_0)$. Then the probability of success of $(\mathcal{F}_b, s_0)$ is given by:*

$$P_{(\mathcal{F}_b,s_0)} = \frac{Y_0}{|E|} \quad (1)$$

Note that there is only one winning state in a single-partition game: the state which the nodes on the randomising partition are in.

Note that this result is similar to the one by Hassin and Peleg (see Section 4), but now instead of considering the entire graph, we consider only one partition: say $(\mathcal{F}_b, s_0)$ is a game in which W is γ-monochromatic. Then, defining $U_\gamma = \{u \in U \mid s_0(u) = \gamma(u)\}$ we have

$$P_{(\mathcal{F}_b,s_0)} = \sum_{u \in U_\gamma} \frac{\deg u}{|E|} \quad (2)$$

Theorem 2 (Expected Duration - Upper-bound). *Let $(\mathcal{F}_b, s_0)$ be a single partition game on a connected graph G, where $|V| = n$ and $|E| = m$. If $Y_0 = 0$ or $Y_0 = m$, then the duration of the game is zero. Otherwise, let $\gamma \in \Gamma$ be the colouring of the monochromatic partition in this initial state. Denote $Z_i(v)$ as the number of black edges connected to v on round i. Finally, let V_i be the monochromatic partition on round i. Then, we have*

$$mY_0 - Y_0^2 = \mathbb{E}\left(\sum_{i=0}^{\infty} \sum_{v \in V_i} Z_i(v)\left(\deg v - Z_i(v)\right)\right) \quad (3)$$

Thus, because the internal sum is greater than or equal to 1 for the duration of the game we have that the expectation of the duration of the game $(\mathcal{F}_b, s_0)$ until there are either no black edges left (the game is a losing game) or only black edges left (colouring γ wins) is bounded by:

$$D_{(\mathcal{F}_b,s_0)} \leq mY_0 - Y_0^2 \quad (4)$$

The proof of this theorem is a direct application of the following lemmas (see Appendix for proofs).

Lemma 1. $\mathbb{E}(Y_\infty^2) = mY_0$.

Lemma 2. *For each $i \geq 0$, we have*

$$\mathbb{E}(Y_{i+1}^2) - Y_0^2 = \sum_{s=0}^{i} \mathbb{E}(Z_s^2). \quad (5)$$

Lemma 3. *For each $s \geq 0$ we have that*

$$\mathbb{E}(Z_s^2) = \mathbb{E}\left(\sum_{j=1}^{n} Z_i^{v_j}\left(\deg v_j - Z_i^{v_j}\right)\right) \quad (6)$$

Example 3. *Consider the initial configuration of a game $(\mathcal{F}_b, s_0)$ in Figure 3. Here $\Gamma = \{b, r\}$ where b and r colour all nodes in blue and red, respectively.*

We have that $|E| = 12$ and $Y_0 = 7$, therefore, by Theorem 1, $P_{(\mathcal{F}_b,s_0)} = P^b_{(\mathcal{F}_b,s_0)} = \frac{7}{12}$. Also, by Theorem 2, the expected duration of the game is bounded by $D_{(G,s_0)} \leq$ $84 - 49 = 35$.

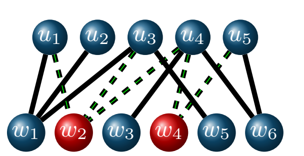

Figure 3: Game $(\mathcal{F}_b, s_0)$.

2.4 General bipartite graphs

The previous results were somehow very restrictive given that single-partition games are rare if we consider the initial configuration to be random. In order to be able to solve the problem for an arbitrary initial configuration, we define a function that splits the original problem into two single-partition games. Then, based on the results of the two new games, we can fully determine what happens on the original one.

Definition 9 (Split function). *We let $(\mathcal{F}_b, s_0)$ be a game on a connected graph G and $\textbf{split}_G$ be the function that takes a colouring $s \in \mathcal{C}$ and outputs two colourings $\rho, \sigma \in \mathcal{C}$ such that one colouring copies the colours of s in partition U and where the other colouring copies colours in W, colouring the remaining nodes according to the same given winning colouring γ. Formally, we have that*

$$\textbf{split}_G : \mathcal{C} \times \Gamma \to \mathcal{C} \times \mathcal{C}$$
$$\textbf{split}_G(s, \gamma) = (\rho, \sigma)$$

Where $\rho \restriction_U = s \restriction_U$ and $\rho \restriction_W = \gamma \restriction_W$, also $\sigma \restriction_W = s \restriction_W$ and $\sigma \restriction_U = \gamma \restriction_U$.

Example 4. *Let us consider the initial configuration of a game $(\mathcal{F}_b, s_0)$ as shown in Figure 5. Let $\Gamma = \{b, r\}$, where b and r are the monochromatic colourings blue and red, respectively. Applying the split function $\textbf{split}_G(s_0, b) = (\rho_0, \sigma_0)$ we get Figure 4.*

Note that the split function is solely a concrete way to visualise the independence of the behaviour of the two partitions in such games.

Theorem 3. *Let $(\mathcal{F}_b, s_0)$ be a flag coordination game as in Definition 3 and let $(\rho_0, \sigma_0) = \textbf{split}_G(s_0, \gamma)$, where $\gamma \in \Gamma$ is any given winning configuration. In these conditions,*

$$P^\gamma_{(\mathcal{F}_b,s_0)} = P_{(\mathcal{F}_b,\rho_0)} P_{(\mathcal{F}_b,\sigma_0)} = \sum_{u \in U_\gamma} \frac{\deg u}{|E|} \sum_{w \in W_\gamma} \frac{\deg w}{|E|} \quad (7)$$

In other words, we have that goal γ is the winning configuration of $(\mathcal{F}_b, s_0)$ if and only if both $(\mathcal{F}_b, \rho_0)$ and $(\mathcal{F}_b, \sigma_0)$ are winning games (note that, according to Proposition 3, they can only reach one winning configuration, and that is γ). Alternatively, denoting Y_0 and X_0 as the number of black edges in $(\mathcal{F}_b, \rho_0)$, respectively, then $P^\gamma_{(\mathcal{F}_b,s_0)} = \frac{Y_0 X_0}{|E|^2}$.

If we are not interested in the winning colouring but solely whether the game ends successfully, since winning colourings are stable, we have that

$$P_{(\mathcal{F}_b,s_0)} = \sum_{\gamma \in \Gamma} P^\gamma_{(\mathcal{F}_b,s_0)} \quad (8)$$

Proof. The idea of the proof is straightforward: the behaviour of nodes in W in $(\mathcal{F}_b, \rho_0)$ is the same as the ones in W in $(\mathcal{F}_b, s_0)$. That is because a node $w \in W$ see the same

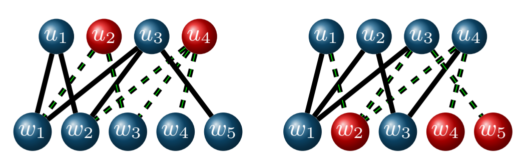

Figure 4: Games $(\mathcal{F}_b, \rho_0)$ and $(\mathcal{F}_b, \sigma_0)$.

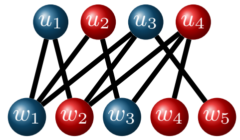

Figure 5: Game $(\mathcal{F}_b, s_0)$.

set of colours in both games. At the same time, the behaviour of U in $(\mathcal{F}_b, \sigma_0)$ is the same as U in $(\mathcal{F}_b, s_0)$. During the next round, the same is true but now for the opposite partitions. Moreover, all nodes v in non-randomising partitions (the ones looking at vertices all in γ) will have a deterministic behaviour: to choose the colour $\gamma(v)$.

The core of the proof relies on the fact that the behaviour of the nodes in a given partition, say U, in game $(\mathcal{F}_b, s_0)$ on even rounds will never depend on decisions these same nodes took on previous odd rounds. That happens because bipartite graphs have no cycles with odd length. All the 'information' contained in partition U on s_i is captured by partition W on s_{i+1}, and only by partition W. That information will go back to U on s_{i+2}. Therefore, the **split** function captures that behaviour by generating two independent games whose nodes in randomising partitions make decisions as nodes in U and W on $(\mathcal{F}_b, s_0)$ do. □

PROPOSITION 4 (EXPECTED DURATION - GENERAL). *Let $(\mathcal{F}_b, s_0)$ be a game on a connected graph G. Then the game will finish in $O(n^3 \log n)$ rounds.*

PROOF. This is an direct application of the results in [7]. We have that, for each of the games $(\mathcal{F}_b, \rho_0)$ and $(\mathcal{F}_b, \sigma_0)$, where $\mathbf{split}_G(s_0, \gamma) = (\rho_0, \sigma_0)$, the expected time is bounded by $O(n^3 \log n)$. Therefore the expected number of rounds for the general game is also bounded by $O(n^3 \log n)$. □

It is left for future work to find a way to use the upper-bound found for single-partition games (see Theorem 2) on a solution for general games on bipartite graphs. We know we cannot just take the greater of the two bounds for $(\mathcal{F}_b, \rho_0)$ and $(\mathcal{F}_b, \sigma_0)$ to estimate the bound for $(\mathcal{F}_b, s_0)$. As an illustration of this, consider the problem of expected times in dice tossing: although the expected number of tosses to get a face, say "4", in one die is 6, the expected number of rounds, on the other hand, for two dice (both tossed in each round) in order to get a "4" in both, not necessarily at the same time, is $\frac{96}{11}$ which is greater than 6.

All previous results take into account that every node knows the position of the neighbours they see in the graph G. If we relax that condition determining that nodes do see the colours of their neighbours, but not their labels, then we cannot solve the generalised consensus problem in the same way. In a non-bipartite graph, the standard consensus problem can be solved, as shown in [7]. Moreover, in bipartite graphs, not only can the standard consensus problem be solved, but also the proper colouring problem. Nodes do not have to know the partition they are in nor the labels of the nodes whose colours they are looking at, as long as they know they are in a bipartite graph and whether they seek standard consensus or proper colouring of the graph. That is the case because for both problems all neighbours of a given node are coloured the same in each of the goal states $\gamma \in \Gamma$.

	Single-partition games	General games on bipartite graphs
Probability of colouring γ winning	$\frac{Y_0}{\lvert E \rvert}$	$\frac{Y_0 X_0}{\lvert E \rvert^2}$
Upper-bound for expected duration	$mY_0 - Y_0^2$	$O(n^3 \log n)$

Table 1: Summary of results

3. GAME-THEORETIC APPLICATION

Let us consider the game $(\mathcal{F}_b, s_0)$ as shown in Figure 3. Say there is one player, player R, in this game and she wins if the consensus game played on this graph is successful **and** red is the winning colour. By Theorem 3, we can calculate the probability of this to happen: $P^r_{(\mathcal{F}_b, s_0)} = \frac{7}{10} \times \frac{4}{10} = \frac{28}{100}$.

Let us suppose now that the payoff for R in winning such game is \$100, whereas losing the game gives R a payoff of \$0. Say that, before the game starts, she is offered to freely choose a node and change its colour in exchange for \$10. Should she accept the offer? In this case, which node's colour should R change in order to maximise her gains? Finally, in general, what should be the fair price given by the function $\mathbf{cost}_{\mathcal{F}_b} : \mathcal{C} \to \mathbb{R}$ that associates a configuration $s \in \mathcal{C}$ to a real number $c \in \mathbb{R}$ that represents the maximum price player R is willing to pay for such an action?

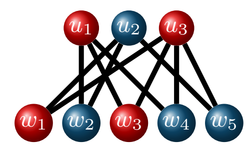

Figure 6: An Initial configuration $(\mathcal{F}_b, s_0)$.

We trivially answer the first question by showing that player R should accept the offer. Changing the colour of v_2, for example, would give a probability of winning of $\frac{10}{10} \times \frac{4}{10} = \frac{40}{100}$, therefore an increase of \$12 on average. So, by paying an extra \$10, she should expect an increase of her average final payoff by \$2. That, however, is not the best move for player R. Although v_2 is the most connected blue node in s_0, changing the colour of either of the other ones would contribute for a (equally) higher increase of payoff. We have $\frac{7}{10} \times \frac{6}{10} = \frac{42}{100}$, and therefore $\mathbf{cost}_{\mathcal{F}_b}(s_0) = 14$.

Let us now present the following simple result before exploring a similar problem with more than one player.

PROPOSITION 5. *Let $(\mathcal{F}_b, s_0)$ be a flag coordination game as in definition 3 with $X = \{red, blue\}$ and $\Gamma = \{r, b\}$, where $r(v) = red$ and $b(v) = blue$ $\forall v \in V$. Let R be a player that has the option to change the colour of a node before the flag coordination game starts in exchange for part of her payoff. Assume also that R's payoff is $L > 0$ if the game finishes in a consensus of red nodes and 0 otherwise. In these condition, the expected cost of the action of freely changing a node's colour is given by*

$$\mathbf{cost}_{\mathcal{F}_b}(s_0) = \frac{\max\{\Delta(U_b) \lvert E_{W_r} \rvert, \Delta(W_b) \lvert E_{U_r} \rvert\}}{E^2} L \tag{9}$$

Where $\Delta(V)$ stands for the maximum degree among the nodes in V and $\lvert E_V \rvert = \sum_{v \in V} \deg v$. Remind that $V_\gamma = \{v \in V \mid s_0(v) = \gamma(v)\}$.

PROOF. Firstly, let us recall from Theorem 3 that the probability of a given colour c winning a game is the product of two sums, each over nodes in a different partition. Thus, the particularities of connections and degrees do not matter as long as the sum of the degrees of nodes coloured c is each partition is the same. For that reason, the degree of the node R changes colour of *and* the partition it belongs to are only what counts towards the increase in her payoff.

Therefore, we just have to compare the (one of the) most connected blue nodes in each partition and choose the one that most increases $P^r_{(\mathcal{F}_b,s_0)}$. Since $|E|$ is constant, it is enough to take $\max\{\Delta(U_b)\,|E_{W_r}|\,,\Delta(W_b)\,|E_{U_r}|\}$, and thus the expected increase of R's payoff (and consequently maximum price to pay) is given by Equation 9. □

Let us now consider an alternative game played now by R and B based on a flag coordination game $(\mathcal{F}_b, s_0)$ in which the set of colours is $X = \{b, r\}$. Their payoffs according to the possible endings of the flag coordination game are:

	r wins	$(\mathcal{F}_b, s_0)$ unsuccessful	b wins
R	1	α	0
B	0	$1-\alpha$	1

Table 2: Payoffs for Players R and B Based on the Outcome of Game $(\mathcal{F}_b, s_0)$.

We assume, *w.l.o.g.*, $0 \le \alpha \le \frac{1}{2}$. If each player is allowed to synchronously change the colour of a node in their favour, what should be their choice? Again, it is not necessarily true that the most connected node in V is the best option. There might also not be a pure strategy Nash equilibrium for that scenario.

PROPOSITION 6. *Let players R and B be observers of a flag coordination game $(\mathcal{F}_b, s_0)$. Before the start of $(\mathcal{F}_b, s_0)$, player R can change the colour of a blue node to red, and player B can, at the same time, change the colour of a red node to blue. By proposition 5, we know that there are only two choices for each player: the most connected opponent nodes in each partition. B's expected payoffs are then described in table 3.*

B \ R	U_b	W_b
U_r	$Pay_B(U_r, U_b)$	$Pay_B(U_r, W_b)$
W_r	$Pay_B(W_r, U_b)$	$Pay_B(W_r, W_b)$

Table 3: Payoffs for B

Here, $Pay_B(P_r, Q_b)$ represents the expected payoff for B if he changes the colour of the one of the most connected red node in partition $P \in \{U, W\}$ and R changes the colour of one of the the most connected blue nodes in partition $Q \in \{U, W\}$. For example, $Pay_B(U_r, U_b) = \frac{1}{E^2}\big[\,|E_{W_b}|\,(|E_{U_b}| + \Delta(U_r) - \Delta(U_b))(1-2\alpha) + |E|\,(|E_{U_b}| + |E_{W_b}| + \Delta(U_r) - \Delta(U_b))\big]$

In these conditions, either one of the following applies. If $\alpha = \frac{1}{2}$, then the best strategy for both players is to choose the best increase in their probability of winning regardless of the other player's decision. Otherwise, if $Pay_B(U_r, U_b) \ge Pay_B(W_r, U_b)$ and $Pay_B(U_r, U_b) \ge Pay_B(W_r, U_b)$ (resp. for $\le$ and for Pay_A), then there is a pure strategy Nash Equilibrium for this game. Otherwise, there is a mixed strategy Nash equilibrium in which the probability of A choosing U_r is given by

$$p = \frac{K + \Delta(U_r)\Delta(W_b) + \Delta(W_r)\,|E_{U_b}| - \Delta(U_r)\,|E_{W_b}|}{\Delta(U_r)\Delta(W_b) + \Delta(W_r)\Delta(U_b)}$$

Where $K = \frac{|E|(\Delta(W_r) - \Delta(U_r))}{1-2\alpha}$. Then, the expected payoff for B is given by $Pay_B(U_r, U_b)p + Pay_B(U_r, W_b)(1-p)$.

The proof is given by direct calculation of the Nash Equilibrium for this game.

EXAMPLE 5. *For a real life application of these results, consider partitions U and W are groups of doctors and patients, respectively. An edge $\{u, w\}$ represents that u is a patient of doctor w. Let also R and B be two competing health insurance companies. We assume each node must adopt one and only one of the two insurance options at a given time. Of course, doctors want to accept the insurance of their patients and vice-versa. Finally, we assume that they all might switch to another company (or not) at a given common time every year, and that they make this decision based on their neighbours' current choices (as in Definition 2).*

We can then apply the previous results if R and B are willing to offer benefits for patients/doctor in order to persuade them to change companies.

4. RELATED WORK

The problem of distributed consensus in computational systems has been extensively studied, including specifically in multi-agent contexts; for reviews, see e.g., [11, 13]. If we consider communications protocols in which nodes base their decisions only on the colour of one of their neighbours (chosen at random), the probability of convergence for each colour and the complexity of the expected duration has been found by Hassin and Peleg [7] for any non-bipartite graph. In their model, nodes have a common clock and change their colours synchronously in rounds, until a consensus is reached. The probability of a given colour c to win that consensus game is $\sum_{v \in V_c} \frac{\deg(v)}{2|E|}$, where V_c is the subset of nodes that are coloured c. All edges are assumed to have the same weight. Note that our results extend the work of Hassin and Peleg to bipartite graphs, as well as proposing a solution for the generalised consensus problem (see Definition 2), provided nodes are aware of their neighbours' labels and of the graph structure.

Experiments with human participants for proper colouring of graphs on networks were conducted by Kearns *et al.* [9]. These authors explored different restrictions on the visibilities of the participating human agents and showed that more information does not necessarily lead to a better performance. There are two key differences between [9] and our model. Firstly, [9] does not assume that agents share a common clock, so that agents could change their selected colours at will, asynchronously. Secondly, the agents in question were actual humans in experiments who were able to use any decision algorithm, or combination of algorithms, or none at all, to select colours. Real humans may also have been whimsical or malicious.

A game-theoretic approach for graph colouring was studied by Panagopoulou and Spirakis in [12]. In their model, each node v chooses a colour and then receives a payoff equal to the number of nodes that have chosen the same colour, unless a neighbour of v is one of those nodes choosing the same colour, in which case the payoff to v is zero.

The authors prove that a Nash Equilibrium is always possible in this game. The key difference with our work is that Panagopoulou and Spirakis do not require nodes to choose their colours synchronously.

Other papers that consider different variants of this problem are [5, 10, 3]. In a nutshell, in [5], nodes make their decisions based on two random neighbours, not just one. In [10], one-round algorithms are studied instead. In [3], the number of available colours for the nodes is $\Delta + 2$, whereas in our work the number of colours is not a function of Δ (for example, we use 2 colours in any bipartite graph for the graph colouring problem for any Δ).

5. DISCUSSION AND CONCLUSIONS

We now return to the problem posed in Example 1. It may seem counter-intuitive, but we can now clearly see that the probability of *gray* being the winning colouring, although there are 7 *gray* nodes of the 20 nodes in total, is zero. Note there there is no *gray* node in the 'even' partition. By Theorem 3, we have that the probability of *blue* winning is given by $P^{blue}_{(\mathcal{F}_b,s_0)} = \frac{2}{20} \times \frac{12}{20} = 6\%$ and the probability of *red* winning is $P^{red}_{(\mathcal{F}_b,s_0)} = \frac{4}{20} \times \frac{8}{20} = 8\%$. Thus, in this case, the least common colour (also the colour with the fewest number of edges connected to nodes of that colour) is the most likely to win. However, note that the most likely outcome is not success, but that the game is a losing game, with probability 86%.

Such unexpected situations do not occur when G is non-bipartite: in these cases the most connected colour (considering the weights of edges) always has the highest probability of winning [7]. Also, the fact that non-bipartite graphs have at least one odd cycle implies that every game on such graphs is a winning game.

Note that this now generalised consensus flag coordination game for any graph G does not require that agents know their *current* colour in order to make a decision. Although each agent has to make a decision of a colour at each round, this decision may be forgotten immediately afterwards, and before deciding colours at the next round.

As mentioned in the Introduction, flag coordination problems arise in many areas of computer science, economics, and public policy. For consensus protocols, applications in distributed computing have been known for some time. The recent rise of distributed ledger technologies (also known as blockchains), which use consensus protocols and cryptographic methods to create variables having shared state across a set of autonomous nodes, creates another class of applications in many commercial domains [1].

In this paper, we have explored research Questions 1, 2 and 3 of Section 1, particularly when agents are connected via a bipartite graph. One avenue for future research is to consider these questions for n-partite graphs, for any integer n. A second area of future work is to obtain tighter upper bounds for single-partition games, perhaps through finer-grained analysis of the probability distributions of the colours. A third area is to consider agent decision-algorithms where agents are able to base their choice on the colours of more than one of their neighbours; doing this will permit consideration of larger sets of desired goal states. Understanding the effects of n-partite graphs and of more sophisticated decision-algorithms will allow the exploration of research Questions 4 and 5 of Section 1.

APPENDIX

This appendix presents proofs of Theorem 1 and Lemmas 1, 2 and 3 omitted in the text above.

PROOF OF THEOREM 1. We first prove that $(Y_i)_{i\geq 0}$ is a bounded martingale [8] with respect to $(s_i)_{i\geq 0}$ (note that by knowing s_i we also have $r_i = f(s_i)$). Denote also $\delta_i(v) = Z_{i+i}(v) - Z_i(v)$, where $Z_i(v)$ denotes the number of black edges connected to v on round i. Note that $\deg v$ stands for the number of neighbours of v.

If $Y_0 = |E|$, then $P_{(\mathcal{F}_b,s_0)} = 1$. On the other hand, $P_{(\mathcal{F}_b,s_0)} = 0$ if $Y_0 = 0$. Else, we call, V_i the monochromatic partition on round i. Then,

$$\begin{aligned}\mathbb{E}(Y_{i+1} \mid s_i) &= \mathbb{E}\left(\sum_{v\in V_i}(Z_i(v)+\delta_i(v)) \mid s_i\right)\\ &= \sum_{v\in V_i} Z_i(v) + \sum_{v\in V_i}\mathbb{E}(\delta_i(v) \mid s_i) =\\ &= Y_i + \sum_{v\in V_i}\Big[\mathbf{P}\{s_{i+1}(v) = s_i(v)\}(\deg v - Z_i(v))\\ &\quad + \mathbf{P}\{s_{i+1}(v) \neq s_i(v)\}(-Z_i(v))\Big] =\\ &= Y_i\end{aligned}$$

The last step follows from $\mathbf{P}\{s_{i+1}(v) = s_i(v)\} = \frac{Z_i(v)}{\deg v}$ and $\mathbf{P}\{s_{i+1}(v) \neq s_i(v)\} = \frac{\deg v - Z_i(v)}{\deg v}$.

Therefore, $(Y_i)_{i\geq 0}$ is a martingale with respect to $(s_i)_{i\geq 0}$. Since $0 \leq Y_i \leq |E|$, the martingale is also bounded and thus we can apply Doob's Optional stopping theorem to get $\mathbb{E}(Y_0) = \mathbb{E}(Y_\infty) = Y_d$, where d stands for the duration of the game $(\mathcal{F}_b, s_0)$ until all edges turn black or none is black anymore. Note that there are two absorbing states: 0 and $|E|$. Thus,

$$Y_0 = \mathbb{E}(Y_0) = \mathbb{E}(Y_\infty) = |E|\,\mathbf{P}(Y_d = |E|) + 0\mathbf{P}(Y_d = 0) \quad (10)$$

That concludes the proof. □

PROOF OF LEMMA 1. From $\mathbf{P}(Y_d = m) = \frac{Y_0}{m}$, we have

$$\mathbb{E}(Y_\infty^2) = m^2\mathbf{P}(Y_d = m) + 0^2\mathbf{P}(Y_d = 0) = mY_0 \quad □$$

PROOF OF LEMMA 2. It is clear that $Y_{i+1} = Y_i + Z_i$, where $Z_i = Y_{i+1} - Y_i$. Note that Z_i is the sum of $Z_i(v)$ for nodes v in one of the partitions of G. By Theorem 1, $\mathbb{E}(Z_i \mid s_i) = 0$. Then,

$$\mathbb{E}(Y_{i+1}^2 \mid s_i) = \mathbb{E}(Y_i^2 + 2Y_iZ_i + Z_i^2 \mid s_i) = Y_i^2 + \mathbb{E}(Z_i^2 \mid s_i)$$

By induction we have the result. □

PROOF OF LEMMA 3. We start by $\mathbb{E}(Z_i^2 \mid s_i)$. Using the notation $\delta_i(v) = Z_{i+1}(v) - Z_i(v)$ we have $Z_i = \sum_{v\in V_i}\delta(v)$. Since $\mathbb{E}(Z_i \mid s_i) = 0$, then $\mathbb{E}(Z_i^2 \mid s_i) = \mathrm{Var}(Z_i \mid s_i)$. The random variables $\delta_i(v)$ are independent, then

$$\mathrm{Var}(Z_i \mid s_i) = \sum_{v\in V_i}\mathrm{Var}(\delta_i(v)) = \sum_{v\in V_i} Z_i(v_j)(\deg v - Z_i(v))$$

because we have $(-Z_i(v))^2\,\frac{\deg v - Z_i(v)}{\deg v} + (\deg v - Z_i(v))^2\,\frac{Z_i(v)}{\deg v} = \mathrm{Var}(\delta_i(v))$.

Using $\mathbb{E}(Z_i^2) = \mathbb{E}(\mathbb{E}(Z_i^2 \mid s_i)) = \mathbb{E}(\mathrm{Var}(Z_i \mid s_i))$, we get

$$\mathbb{E}(Z_i^2) = \mathbb{E}\left(\sum_{v\in V_i} Z_i(v)(\deg v - Z_i(v))\right) \quad (11)$$

Which concludes the proof. □

REFERENCES

[1] Anonymous. The trust machine. *The Economist*, 31 October 2015.

[2] J. Barwise. Scenes and other situations. *Journal of Philosophy*, 78(7):369–397, 1981.

[3] K. Chaudhuri, F. C. Graham, and M. S. Jamall. A network coloring game. In *International Workshop on Internet and Network Economics*, pages 522–530. Springer, 2008.

[4] C. Cooper, R. Elsasser, H. Ono, and T. Radzik. Coalescing random walks and voting on connected graphs. *SIAM Journal on Discrete Mathematics*, 27(4):1748–1758, 2013.

[5] C. Cooper, R. Elsässer, and T. Radzik. The power of two choices in distributed voting. In *International Colloquium on Automata, Languages, and Programming*, pages 435–446. Springer, 2014.

[6] S. Y. Grigoriev and V. B. Priezzhev. Random walk of annihilating particles on the ring. *Theoretical and Mathematical Physics*, 146(3):411–420, 2006.

[7] Y. Hassin and D. Peleg. Distributed probabilistic polling and applications to proportionate agreement. *Information and Computation*, 171(2):248–268, 2001.

[8] J. Jacod and A. Shiryaev. *Limit Theorems for Stochastic Processes*, volume 288. Springer, 2013.

[9] M. Kearns, S. Suri, and N. Montfort. An experimental study of the coloring problem on human subject networks. *Science*, 313(5788):824–827, 2006.

[10] F. Kuhn and R. Wattenhofer. On the complexity of distributed graph coloring. In *Proceedings of 25th ACM Symposium on Principles of Distributed Computing*, pages 7–15. ACM, 2006.

[11] R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of IEEE*, 95(1):215–233, 2007.

[12] P. N. Panagopoulou and P. G. Spirakis. A game theoretic approach for efficient graph coloring. In *International Symposium on Algorithms and Computation*, pages 183–195. Springer, 2008.

[13] W. Ren, R. W. Beard, and E. M. Atkins. A survey of consensus problems in multi-agent coordination. In *Proceedings of 2005 American Control Conference, 2005*, pages 1859–1864. IEEE, 2005.

[14] C. Shapiro and H. Varian. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Review Press, Cambridge, MA, USA, 1998.

[15] T. Weitzel. *Economics of Standards in Information Networks*. Physica, Heidelberg, Germany, 2004.