Self-Organized Collective Decision Making: The Weighted Voter Model

Gabriele Valentini†, Heiko Hamann‡ and Marco Dorigo†‡
†IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
{gvalenti, mdorigo}@ulb.ac.be
‡Department of Computer Science, University of Paderborn, Paderborn, Germany
{heiko.hamann, marco.dorigo}@uni-paderborn.de

ABSTRACT

Collective decision making in self-organized systems is challenging because it relies on local perception and local communication. Globally defined qualities such as consensus time and decision accuracy are both difficult to predict and difficult to guarantee. We present the weighted voter model which implements a self-organized collective decision making process. We provide an ODE model, a master equation model (numerically solved by the Gillespie algorithm), and agent-based simulations of the proposed decision-making strategy. This set of models enables us to investigate the system behavior in the thermodynamic limit and to investigate finite-size effects due to random fluctuations. Based on our results, we give minimum requirements to guarantee consensus on the optimal decision, a minimum swarm size to guarantee a certain accuracy, and we show that the proposed approach scales with system size and is robust to noise.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: intelligent agents, multiagent systems

General Terms

Algorithms, Performance, Theory

Keywords

collective decision making, voter model, consensus, swarm intelligence, swarm robotics, self-organization, modeling

1. INTRODUCTION

The problem of finding a collective agreement over the most favorable choice among a set of alternatives, namely the *best-of-n* decision problem, is a general and abstract cognitive challenge for both natural [8, 9, 24, 25, 29, 33] and artificial [2, 20, 22, 26, 31] self-organized systems. In this paper, we focus on artificial systems and we describe a control algorithm that solves the best-of-n decision problem with a self-organized approach. Our control algorithm relies on a positive feedback mechanism inspired by the house hunting behavior of honeybee swarms [9, 24, 25, 33]. We apply methods from opinion dynamics [4, 10] to define mathematical models that reliably predict the system behavior as a function of a number of relevant parameters. The key idea is to develop a control algorithm along with a set of models that allows designers to give guarantees on the expected system performance.

Collective decision making is a significant challenge for artificial self-organized systems. Notably, for those systems designed to provide a valuable alternative to classical centralized solutions (e.g., robotic swarms [3, 16, 17, 20, 22, 23], wireless sensor networks [19, 27, 28], virtual agents operating in high-dimensional spaces [6, 7, 11, 12]). Artificial self-organized systems achieve high degrees of scalability, flexibility, and robustness by relying on limited perception and communication capabilities (e.g., few and noisy sensors, only neighbor-to-neighbor communication). Depending on the particular scenario of interest, the collective decision making problem may demand for problem-dependent requirements that are either loose (e.g., a large majority in favor of an alternative suffices) or stringent (e.g., consensus on the best available alternative within given time constraints). When studying solutions to the best-of-n decision problem, we focus on scenarios characterized by strict problem-dependent requirements. Our goal is to define a control algorithm capable to guarantee unanimous agreement among the agents, high accuracy of the decision and predictable performance.

With the purpose of satisfying the above mentioned requirements, researchers took inspiration either by looking at natural systems such as honeybee swarms and ant colonies [16, 17, 22, 23], or by adapting opinion formation models borrowed from the field of opinion dynamics [2, 20, 26, 31].

In the first case, designers devise bio-inspired control algorithms largely by considering the *house hunting* behavior characteristic of honeybee swarms [9, 24, 25, 29, 33] and ant colonies [8, 9]. These systems offer interesting examples of natural, self-organized solutions to the best-of-n decision problem that are both accurate and reliable, thus appealing to designers. For swarms and colonies it is essential to make good choices, either periodically or occasionally, on where to relocate the nest — choices that are often decisive for the survival of the whole population. They evolved therefore self-organized decision mechanisms that are flexible enough to cope with the speed and accuracy trade-off [24] demanded by the particular scenario and that enable them to reliably assert which of many candidate sites represents a

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

high-quality choice for the population. These natural solutions are based on simple positive feedback mechanisms [9, 24, 25, 29, 33] and do not rely on individual agents' cognitive skills to directly compare and discriminate sites [8, 33].

In the second case, researchers take advantage from theoretical frameworks available in the field of opinion dynamics. Opinion dynamics is a branch of statistical physics that focuses on the dynamics of social systems [4, 10] (e.g., opinion dynamics, spread of disease, democratic voting). Generally, opinion formation models consist in a round-based stochastic process where, at each round, a randomly chosen agent from the population applies a particular decision rule over a discrete set of equivalent opinions. In opinion dynamics, researchers consider agents as adaptive entities rather than rational agents, and the focus of studies is on communication structures instead of particular decision strategies [4]. Even though they rarely take the quality of alternatives into consideration [4, 10], the theoretical frameworks built by physicists can be adapted and the mathematical results can be exploited to guarantee performance.

In this paper we use a hybrid approach that extends the classic voter model — a simple and general model of democratic voting. To consider opinion qualities, we introduce a positive feedback mechanism inspired by the house hunting behavior of honeybee swarms. In the classic voter model [5, 18], agents are distributed over a static lattice and they interact only with their neighbors. At each round, a randomly picked agent adopts the opinion of a random neighbor. The evolution of the process continues until consensus is eventually reached [5, 18]. We extend the classic voter model with a number of changes and we devise a self-organized decision making strategy referred to as the *weighted voter model*. First, we consider agents with motion capabilities whose neighborhood changes over time; hence the decision process operates on a dynamic interaction network. Second, to drive the system toward consensus on the best opinion, we allow agents to participate in the decision process at different rounds for a time proportional to the qualities of their opinions. This implements a positive feedback mechanism similar to the duration of honeybees' waggle dance (cf. [33]). Finally, as observed in honeybee swarms [9, 24, 25, 33], agents temporarily leave the decision pool after every application of the decision rule in order to survey the quality of their current opinion.

2. WEIGHTED VOTER MODEL

At the abstract level, the best-of-n decision problem consists of a set of n alternatives and a collection (or swarm) of N agents. Each alternative $a_i \in \{a_1, \ldots, a_n\}$ is characterized by a given quality $\rho_i \in (0, 1]$. Agents in the swarm occasionally perceive the quality of alternatives, and, at all time, have a preference for a certain alternative (henceforth referred to as *opinion*). Agents change opinions by applying a given decision rule. The best-of-n decision problem is considered successfully solved if (i) the swarm reaches consensus on a particular opinion and (ii) this opinion is associated to the alternative of highest quality.

We study self-organized solutions to the best-of-n decision problem within the context of embodied agents, i.e., swarms of autonomous robots [1]. We consider agents acting within a bounded, two-dimensional environment which is divided in a number of regions. In the following, we restrict our study to binary decision problems ($n = 2$) and we refer to the two available alternatives as A and B. These alternatives correspond to particular regions in the environment called *sites*. As a consequence, agents' opinions are preferences for these spatially defined sites. In addition to sites, the environment is characterized also by a third region, the *nest*, where all agents are initially located and that functions as a hub for the decision-making process. Agents travel between nest and sites and when they are located at a certain site they can perceive its quality. Without loss of generality, we consider site A to have higher quality than site B and for the remaining of this study we fix $\rho_A = 1$ while varying the value of $\rho_B \in (0, 1]$. That is, the collective decision making problem is successfully solved if the swarm eventually reaches consensus on opinion A — the opinion associated to the optimal site.

2.1 Control Algorithm

To solve the best-of-n decision problem, we have devised a self-organized control algorithm following a behavior-based design approach [1]. Agents in the swarm are driven by the four-state probabilistic finite state machine [1] shown in Figure 1a. In the *waggle dance* state (either $\mathsf{W_A}$ or $\mathsf{W_B}$) agents advertise their own opinion about the best available site. In the *survey* states (either $\mathsf{S_A}$ or $\mathsf{S_B}$) agents estimate the quality of a particular site. The time spent by an agent in a certain state consists of two contributions. Initially, agents spend an unknown period of time to move to the proper region of the environment where to perform the activities defined by the current state (henceforth referred to as *traveling time*). Next, once in the right region, agents act according to the current state for a period of time defined by a control parameter. We choose to adopt exponentially distributed time periods. Thanks to its lack of memory, the exponential distribution facilitates our successive mathematical modeling phase and enhances the predictability of the proposed strategy. Other options would be constant time periods or stochastic periods with different probability distributions. Nonetheless, these latter alternatives are characterized by less favorable mathematical properties.

As soon as agents enter the waggle dance state (either $\mathsf{W_A}$ or $\mathsf{W_B}$) they start performing a random walk within the boundaries of the nest. In the meanwhile, they advertise their own opinion about what they currently consider to be the best site. Before their transition to the survey state, that is, as soon as the waggle dance time expires, agents reconsider their own opinion about the best available site. As in the classic voter model [5, 18], agents first poll the opinion of neighboring agents within a limited interaction range, and then, they adopt a randomly picked opinion from this poll (resulting in a probability Σ_A to adopt opinion A). Similarly to honeybees' decisions (cf. [33]), the agents' decision rule does not take into consideration any information concerning advertised sites. Once agents have deliberated on their opinion, they move to the survey state (either $\mathsf{S_A}$ with probability Σ_A or $\mathsf{S_B}$ with probability $1-\Sigma_A$, see Figure 1a).

To drive the system towards the optimal decision, we introduce a positive feedback mechanism to the waggle dance state which is similar to the waggle dance of honeybees. In honeybees' waggle dance, the duration of the dance is proportional to the quality of the advertised site [9]: the longer the dance, the higher the chances to influence other nestmates. Similarly, in the proposed control algorithm, the time spent by an agent in the waggle dance state is propor-

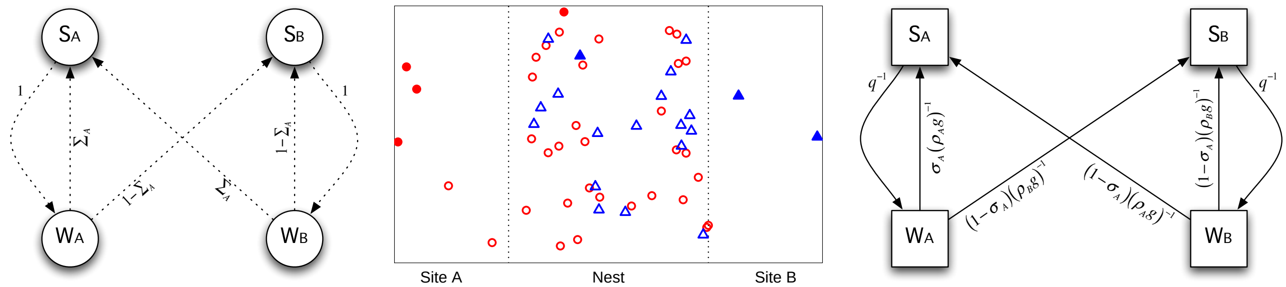

(a) (b) (c)

Figure 1: Different modeling abstractions of the agents' control algorithm: (a) probabilistic finite state machine of the individual agent; (b) screen shot of an agent-based simulation (symbols $\circ, \bullet, \triangle$ and $\blacktriangle$ identify agents in states $\mathsf{W_A}, \mathsf{S_A}, \mathsf{W_B}$ and $\mathsf{S_B}$); and (c) flows of proportions of agents between pairs of control states.

tional to the quality of its opinion. In particular, we consider (as a control parameter) a mean time g for the duration of the waggle dance. This mean time g is then weighted by the quality ρ_i of the agent's opinion and used as rate parameter of the exponential distribution, $1/(\rho_i g)$.

As soon as agents enter the survey state (either $\mathsf{S_A}$ or $\mathsf{S_B}$) they depart from the nest toward the site associated with their opinion. Once arrived at the proper site, agents first determine the duration of the surveying phase by drawing an exponentially distributed random time with mean duration q. Note that the surveying time is independent of the quality of the currently surveyed site, i.e., q is a control parameter (possibly with the constraint of an application-dependent minimum). Next, they explore the site by moving randomly. During this time, agents evaluate the characteristics of the site through their sensors and finally estimate the site quality. Once the surveying time has expired, agents return to the nest and enter the waggle dance state.

2.2 Agent-based Simulation

For the purpose of studying the dynamics of the weighted voter model, we implemented a simple agent-based simulations. In our simulations, agents are represented as massless particles, i.e., points moving at constant velocity in a bounded, two-dimensional space. As a consequence, we do not consider a particular metric or scale for the size of the environment but employ dimensionless units.

Agents are positioned in a 150×50 rectangular arena, shown in Figure 1b. The arena is partitioned into three regions: two 40×50 regions at the two ends of the arena represent the sites, respectively, site A at the left-most side and site B at the right-most side; and a 70×50 region centered between the two sites represents the nest. Agents are equipped with a digital compass that, when necessary, allows them to reorient toward a particular region of the environment. In Figure 1b, we represent agents' opinions by colored symbols using red circles for opinion A and blue triangles for opinion B. Empty symbols represent agents in the waggle dance state (either $\mathsf{W_A}$ or $\mathsf{W_B}$) and filled symbols correspond to agents in the survey state (either $\mathsf{S_A}$ or $\mathsf{S_B}$).

In our simulations, agents perform the control algorithm described in Section 2.1. Their motion is determined by a random walk implemented as follows. Agents move straight for a normally distributed amount of time; next, they uniformly choose a new orientation and they resume a straight motion. Being dimensionless points, we do not consider collisions between agents. However, agents do collide with the boundaries of the arena. A collision with a wall changes the agent's direction of motion by mirroring the incidence angle.

Finally, to change opinion, agents firstly pool the opinions of neighboring agents within a given interaction range r. Secondly, they randomly choose one of the opinions within their pool. In the case that an agent has no neighbors, thus being unable to survey other opinions, it keeps its current opinion prior to move to the survey state (either $\mathsf{S_A}$ or $\mathsf{S_B}$).

3. THERMODYNAMIC PROPERTIES

In the thermodynamic limit, i.e., when the number of agents tends to infinity ($N \to \infty$), random fluctuations that characterize self-organized systems vanish and the system itself approaches a deterministic behavior. Such an asymptotic perspective allows us to gain insights into the dynamics of the weighted voter model irrespectively of its actual system size. The object of interest here is the development of consensus, thus, we look at the dynamics of the opinions in space (nest and sites) and time. We employ dynamical systems theory and we define, under the assumption of null traveling times, a system of ordinary differential equations (ODEs) to describe the weighted voter model dynamics.

We define quantities w_A and w_B as proportions of agents in the swarm that are dancing for their favorite site, respectively, with opinions A and B. Besides, we denote proportions of agents surveying site A with s_A and site B with s_B. The evolution in time of w_A, w_B, s_A and s_B is given by the solution of the following system of ODEs:

$$\begin{cases} \frac{d}{dt}w_A = -\frac{1}{\rho_A g}w_A + \frac{1}{q}s_A \\ \frac{d}{dt}w_B = -\frac{1}{\rho_B g}w_B + \frac{1}{q}s_B \\ \frac{d}{dt}s_A = \sigma_A \frac{1}{\rho_A g}w_A + \sigma_A \frac{1}{\rho_B g}w_B - \frac{1}{q}s_A \\ \frac{d}{dt}s_B = (1-\sigma_A)\frac{1}{\rho_A g}w_A + (1-\sigma_A)\frac{1}{\rho_B g}w_B - \frac{1}{q}s_B, \end{cases}$$

where $\sigma_A = w_A/(w_A + w_B)$ is the probability[1] for an agent to adopt opinion A by randomly choosing a neighbor's opinion. The flows of agents between control states and opinions

[1]We compute the probability σ_A under the well-mixed assumption [21], i.e., we assume a uniform spatial distribution of the opinions among the agents in the nest.

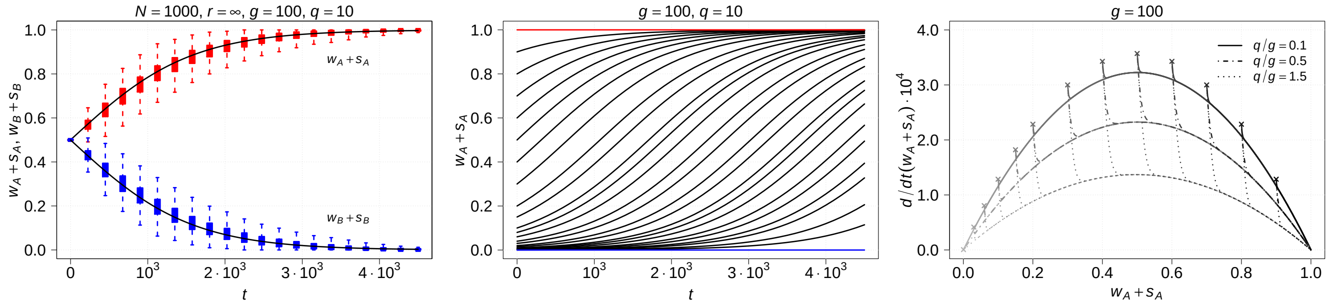

(a) (b) (c)

Figure 2: Predictions of the ODE model $\{\rho_A = 1.0, \rho_B = 0.875\}$: (a) time evolution of consensus (lines) against agent-based simulations (box-plots) ; (b) time evolution of consensus for a number of different initial conditions; (c) proportion of agents favoring opinion A against its first derivative in time.

can be understood by looking at Figure 1c. First, we focus on the first two equations of the above system. The proportion of agents dancing for a site, either w_A for site A or w_B for site B, increases at a rate q^{-1} due to agents returning from the survey states. This proportion also decreases at a rate $(\rho_i g)^{-1}$ as agents leave the waggle dance state. Next, consider the proportion s_A of agents surveying site A (the reasoning is equivalent for site B). This quantity depends on the application of the decision rule underlying the voter model, and thus, from probability σ_A. In particular, s_A increases at rates $\sigma_A(\rho_A g)^{-1}$ and $\sigma_A(\rho_B g)^{-1}$, respectively, for the proportions of agents w_A and w_B leaving the waggle dance state; and it decreases at a rate q^{-1} due to agents that finished the surveying of site A.

There are a number of points to notice in this mathematical model. Firstly, the system of ODEs describing the weighted voter model is defined at a mesoscopic scale. That is, we look at intermediate quantities given by proportions of opinions within different regions of the environment. However, our final interest is in the aggregated information of the overall proportion of opinion A. As a consequence, in the following analysis we mostly focus on the aggregated macroscopic dynamics of $w_A + s_A$. Secondly, given the dependence on the probability σ_A, the system of ODEs is non-linear, and thus, we use standard numerical methods to solve it. Finally, in the definition of this system of ODEs, we are neglecting the temporal delay associated to agents' state transitions, and resulting from the time necessary for agents to move between different regions of the environment.

In the ODEs model, the magnitude and the ratio of control parameters g and q determine the duration of the collective decision process. The longer the time agents spend at the sites, the longer is the consensus time. In particular, for the weighted voter model the consensus time increases linearly with the ratio q/g (data not shown). From an engineering perspective of minimizing the consensus time, a designer should thus prefer values for control parameters such that $g \gg q$. In Figure 2a, we compare predictions of the ODEs model (lines) against agent-based simulations of $N = 10^3$ agents having $r = \infty$ (box-plots). We set sites' qualities to $\rho_A = 1.0$ and $\rho_B = 0.875$. The difference between opinion qualities drives the system toward consensus on the best opinion ($w_A + s_A = 1$ and $w_B + s_B = 0$). The agreement between the ODEs model and agent-based simulations (1000 independent runs) shown in Figure 2a is good.

In the thermodynamic limit, the weighted voter model guarantees consensus on the best opinion. Figure 2b depicts the time evolution of the proportion $w_A + s_A$ of agents in the swarm with opinion A for a number of different initial conditions. When $\rho_A > \rho_B$, every trajectory initially starting at $\{w_A \in (0,1], w_B = 1 - w_A\}$ eventually converges to a consensus on opinion A (that is, $w_A + s_A = 1$). As for the classic voter model [5, 18], the two macroscopic solutions, consensus either on opinion A or on opinion B, characterize the asymptotic behavior of the collective decision making process. Noticeably, for the assumption $\rho_A = 1$ and $\rho_A > \rho_B$, the consensus $w_A + s_A = 1$ is a stable fixed point, while the consensus $w_A + s_A = 0$ is an unstable fixed point. The system of ODEs is characterized by the two equilibria

$$w_A = \frac{g}{g+q},\ w_B = 0,\ s_A = \frac{q}{g+q},\ s_B = 0, \text{ and}$$

$$w_A = 0,\ w_B = \frac{\rho_B g}{\rho_B g + q},\ s_A = 0,\ s_B = \frac{q}{\rho_B g + q}.$$

The first equilibrium is asymptotically stable and the system converges to it in the regime $w_A \in (0,1]$; while the second equilibrium is unstable and the system reaches it only when initialized to a consensus on opinion B ($w_A = 0, w_B = 1$). This result might be useful in applications at design time because it allows to choose the final distribution of agents among nest and site (e.g., to optimize resource gathering rates in foraging tasks [20, 26]).

Finally, we consider the non-equilibrium dynamics (i.e., transient) of the weighted voter model by looking at the speed of change $d/dt(w_A + s_A)$ of opinion A as a function of $(w_A + s_A)$. Figure 2c depicts a number of trajectories for different values of the control parameter ratio q/g and various initial conditions $\{w_A \in (0,1], w_B = 1 - w_A, s_A = 0, s_B = 0\}$ (shaded lines). Initially, the value of $(w_A + s_A; d/dt\, w_A + d/dt\, s_A)$ is determined by the initial conditions of the system of ODEs and is independent of the ratio q/g (see crosses in Figure 2c with trajectories moving from left to right). The speed of change of opinion A as a function of itself decreases abruptly due to agents rapidly redistributing among nest and sites. At a later stage it converges toward a parabolic trajectory determined by the magnitude of q/g. The less time agents spend to survey the sites, the faster is the change in the proportion of opinion A (compare solid against dotted lines). Notice that the speed of change of opinion A reaches its peak at the unbiased conditions $w_A + s_A = 0.5$. Opinions among agents in the swarm depart rapidly from an unbiased situation of $w_A + s_A = 0.5$ and then converge toward consensus with decreasing speed.

4. FINITE-SIZE EFFECTS

In Section 3, we studied properties of the weighted voter model in the thermodynamic limit ($N \to \infty$) where the system behaves deterministically. However, despite self-organized systems are usually composed of a large number of agents, their finite size ($N < \infty$) often plays a crucial role in their dynamics. Toral and Tessone [30], for instance, show that for a number of collective decision making systems finite-size effects may produce unexpected dynamics that differ from those predicted in the thermodynamic limit. It is thus important to develop also mathematical models with dependency on the actual size of the swarm (e.g., time-homogeneous Markov chains as done in [14, 15, 31]).

To study how finite-size effects influence the weighted voter model, we define a macroscopic mathematical model using the formalism of (chemical) master equations [32], i.e., by means of stochastic differential equations. As in Section 3, we assume null traveling times and we neglect the influence of agents' displacement periods. We model the proposed collective decision making system as a set of coupled chemical reactions. In our settings, molecules play the role of agents, reaction rules define the agents' behavior for each state and reaction rates their duration. However, the resulting master equations are characterized by non-linearities that prevent the use of analytical approaches. We use therefore numerical methods, in particular the Gillespie algorithm [13].

Given a swarm of N agents, we define quantities W_A and W_B as the number of agents in the nest dancing, respectively, for site A and site B. Equivalently, we denote the number of agents surveying site A with S_A and those surveying site B with S_B. The weighted voter model is then defined by a set of reaction rules and corresponding reaction rates. Next we give equations for reactions concerning agents that favor opinion A (those for opinion B are equivalent). Within the nest, agents change opinion as a result of the application of the voter model. Such a change in the opinion is captured by reactions

$$W_A \xrightarrow{\Sigma_A(\rho_A g)^{-1}} S_A,$$

$$W_A \xrightarrow{(1-\Sigma_A)(\rho_A g)^{-1}} S_B,$$

where $\Sigma_A = W_A/(W_A + W_B)$. According to the weighted voter model, agents with opinion A reconsider their opinions at an overall rate of $(\rho_A g)^{-1}$. They keep opinion A with probability Σ_A and switch to opinion B with probability $(1-\Sigma_A)$. Probability Σ_A accounts for the interactions of deliberating agents with their neighbors and it is computed under the well-mixed assumption [21]. Agents ceasing to survey site A are modeled by the reaction

$$S_A \xrightarrow{q^{-1}} W_A.$$

Notice that, in contrast to the previous reaction rules, agents return from the survey state at a constant rate q^{-1}.

The above set of reaction rules (together with their equivalent for agents with opinion B) is sufficient to define a master equation following the methods described in van Kampen [32]. The solution of the master equation can be studied numerically applying the Gillespie algorithm — a Markov chain Monte Carlo method capable of generating statistically correct trajectories of stochastic equations [13]. Algorithm 1 depicts our particular formulation of the Gillespie algorithm according to the weighted voter model.

Algorithm 1: GILLESPIE ALGORITHM(N, g, q, ρ_A, ρ_B)

1 Initialize agents in the swarm W_A, W_B, S_A, S_B
2 Initialize time $t = 0$
3 **repeat**
4 Compute the total reaction rate: $\kappa = (\rho_A g)^{-1} W_A + (\rho_B g)^{-1} W_B + q^{-1} S_A + q^{-1} S_B$
5 Generate an exponentially distributed time t' with rate parameter κ and set $t = t + t'$
6 Randomly choose the next reaction to occur with probabilities proportional to reaction rates
7 Update W_A, W_B, S_A, S_B according to the outcome of **step 6**
8 **until** $W_A + S_A \in \{0, N\}$;

In the remaining of this section, we analyze the master equation by means of Algorithm 1 and we asses how finite-size effects influence the dynamics of the weighted voter model. For this purpose, we consider (i) the exit probability E_N, i.e., the probability that a swarm of N agents eventually reaches consensus over opinion A and (ii) the consensus time T_N, i.e., the time necessary to reach consensus on any opinion. We study the weighted voter model under the settings $g = 100$ and $q = 10$, while varying the value of parameters N, ρ_B and r. The numerical solutions of the master equation model are compared against the results of agent-based simulations both averaged over 2.5×10^4 independent runs.

4.1 Influence of Opinion Qualities

When the swarm size N is finite, the dynamics of the weighted voter model is not deterministic and shows stochastic behavior with dependency on N. Figure 3a depicts the exit probability as a function of the initial proportion of opinions A for $N = 100$ agents with unlimited interaction range $r = \infty$ (symbols give data from the agent-based simulations and lines give data from numerical solutions of the master equation obtained with the Gillespie algorithm). We find a good agreement between the numerical solutions of the master equation and the agent-based simulations. For equal opinion qualities ($\rho_A = \rho_B$), E_N resembles a straight line with slope 1; when the difference in quality increases ($\rho_A > \rho_B$) E_N increases as well for all initial conditions $\{W_A \in (0, N], W_B = N - W_A, S_A = 0, S_B = 0\}$ and eventually tends to a step function. Hence, the weighted voter model enables swarms to easily discriminate an inadequate site from a good one, while it correctly generates an unbiased behavior for sites of equal qualities.

As shown in Figure 3b, the master equation also provided a good approximation of the consensus time T_N with a small prediction error that increases with increasing values of T_N. This prediction error is due to our assumption of null traveling times underlying the master equation model. When T_N is maximal agents perform many visits to sites and the influence of traveling times appears more important. Nonetheless, traveling times only affect the transient dynamics of the systems, and therefore the overall consensus time (as in Figure 3b). In contrast, equilibrium dynamics given by the exit probability are independent of such delays (see Fig-

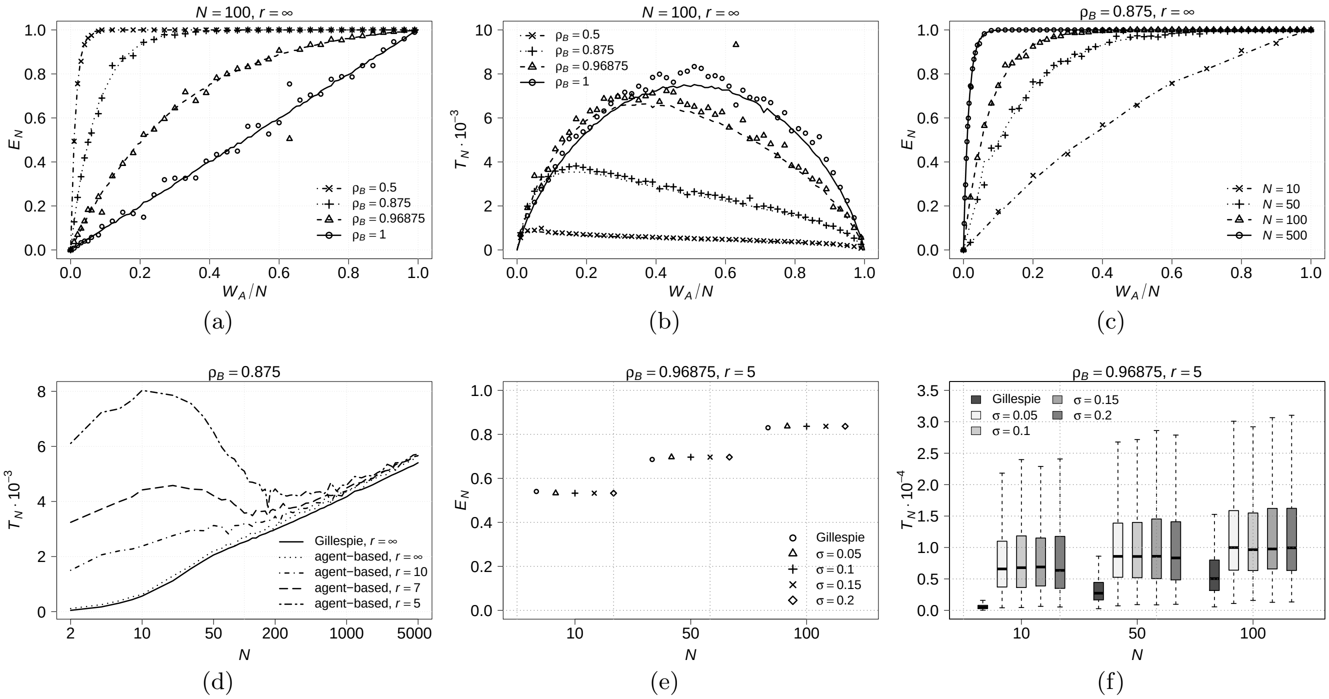

Figure 3: Master equation numerically solved through the Gillespie algorithm (lines) against agent-based simulations (symbols) $\{g = 100, q = 10\}$: (a) exit probability and (b) consensus time as a function of the initial proportion of opinions A; (c) scaling behavior of exit probability as a function of initial proportion of opinions A; (d) scaling behavior of consensus time for an initially unbiased system; (e) robustness to noise of exit probability and (f) of consensus time.

ure 3a). The weighted voter model requires longer times to discriminate between sites of similar qualities, while easier decision problems are solved with much smaller effort; furthermore, for equal sites' quality, $\rho_A = \rho_B$, the consensus time is symmetric (see Figure 3b).

4.2 Scalability with the Size of the Swarm

In the following we study the scalability of the weighted voter model using the master equation model and the agent-based simulations. We use the accuracy of the decision and the consensus time as performance measures.

The accuracy of a decision here is defined by the exit probability which gives the probability of reaching a correct consensus (i.e., on the opinion A). Figure 3c shows the exit probability E_N as a function of the initial proportion of agents favoring opinion A for different swarm sizes and unlimited interaction range $r = \infty$. The agent-based approach (symbols) and the master equation model (lines) show good agreement. When N is small (e.g., $N = 10$) the exit probability approaches a straight line with slope 1, resembling the initial proportion of opinions A. However, as the swarm size increases the exit probability rapidly grows and approaches a step function. Thus, the accuracy of the weighted voter model depends positively on the size of the swarm: Bigger swarms are more accurate. Notice in this context the agreement with the deterministic consensus on opinion A as predicted by the ODE model. In addition, using agent-based simulations, we determined that the accuracy of the weighted voter model is independent of the magnitude of the agents' interaction range r (data not shown).

Next we investigate scalability in terms of consensus time T_N. For unlimited interaction range $r = \infty$, the agent-based simulations behave similarly to numerical solutions of the master equation (compare dotted and solid lines in Figure 3d). The small prediction error of the master equation model is a result of neglecting the traveling time of agents between nest and sites. Notice that, according to the master equation model, the consensus time scales logarithmically with the swarm size N. In contrast, when the interaction range is finite (e.g., $r = 5$) the master equation model fails to predict the consensus time for small swarm sizes N. For larger swarm sizes N and larger interaction range (e.g., $r \in \{7, 10\}$) this discrepancy is reduced. While a finite r affects the consensus time, the prediction error is reduced for larger swarms as higher agent densities reduce the effects of small r. For these finite interaction ranges r, bigger swarms have faster decision processes as found in honeybee swarms [25]. Hence, contrary to the accuracy of the decision, the consensus time is affected by both the size of the swarm and the magnitude of the interaction range.

4.3 Robustness to Noise

In real-world applications, agents of artificial self-organized systems are most likely equipped with low-cost sensors suffering from noisy measurements. Under such conditions, the accuracy of the decision as well as the duration of the decision process may be largely affected, resulting in loss of performance and prediction errors. We study the robustness of the weighted voter model by adding a Gaussian noise with zero mean and deviation $\sigma \in \{0.05, 0.10, 0.15, 0.2\}$ to the values of the site qualities (respectively, $\rho_A = 1.0$ and $\rho_B = 0.96875$) when surveyed by the agents.

With this noise applied, Figure 3e shows the exit probability E_N for a swarm of N and with $r = 5$. The exit probability is equivalent for all tested magnitudes of noise and consistently predicted by the Gillespie algorithm. Thus, as observed for honeybee swarms [24, 25], the weighted voter model shows high robustness to noisy quality of sites.

Figure 3f gives the time T_N necessary to reach consensus under noisy conditions for swarms with finite interaction range $r = 5$. The consensus time seems not to be significantly influenced by noise except for a slightly greater variance for larger noise values. In agreement with the results in Section 4.2, the numerical solution of the master equation fails to predict the consensus time for a swarm with limited interaction range r. However, the prediction error decreases with increasing swarm size N and consequently the prediction is expected to be accurate for large swarms.

5. DISCUSSION

The field of control algorithms for collective decision making systems can be roughly separated into two groups: approaches based on opinion dynamics of either virtual or embodied agents and approaches based on aggregation behaviors of inherently embodied agents. The former are easier to model because they often result in system behaviors that have approximately well-mixed spatial agent distributions while the latter explicitly rely on inhomogeneous agent distributions. In addition, agents based on aggregation behaviors operate in continuous space [17] which defines a continuum of opinions (e.g., at which position to form a cluster) in contrast to the discrete number of alternatives in opinion dynamics [4, 10]. Although it is possible to reduce the opinion continuum of aggregation-based approaches to a discrete set of opinions [15], these reductions are coarse and the modeling of spatial inhomogeneous agent distributions remains a tough challenge. Instead, scenarios and design approaches based on opinion dynamics allow for stringent modeling and subsequently more accurate predictions and guarantees.

An example of an approach that is exclusively based on an aggregation behavior is that of Kernbach et al. [17] which is a behavioral model of young honeybees' aggregation behavior. The investigated decision making problem is to find a consensus on where to form a cluster; due to continuous space there is a continuum of options ('best-of-∞'). The control algorithm works without any direct robot-to-robot communication and uses only indirect communication (detection of close-by robots). While the decision making process is robust, it is also rather slow and particularly difficult to model due to the relevance of spatiality and an emerging, complex dynamics of cluster formation and cluster breakup [16].

The approach of Parker and Zhang [22] to the best-of-n decision problem is an aggregation-based approach that closely resembles the actual house hunting behavior of ant colonies. In their work, agents directly recruit peers. In contrast, our approach is based on an indirect recruitment mechanism inspired by the waggle dance of honeybees. In addition to the deliberation phase their algorithm includes an initial search phase and a final commitment phase. Due to their different objectives, the control algorithm of Parker and Zhang is more complex and has several control states. Also in contrast to our approach, their study is limited to an empirical investigation based on experiments with robot swarms of up to 15 robots [22].

Montes de Oca et al. [20] apply a design approach based on opinion dynamics and develop a control algorithm for a swarm of foraging robots. Differently from the weighted voter model, their control algorithm uses the majority rule [10] as the key feature to implement the collective decision making mechanism. Their approach also differs in the type of positive feedback. In our case, positive feedback is the result of agents that advertise a particular opinion for a time proportional to the opinion qualities. In Montes de Oca et al. [20] instead, positive feedback is an indirect effect of traveling times between regions of the environment that are associated with each opinion. The longer the traveling time the lower is the opinion quality. As a consequence, their method takes more time to discriminate between opinions of very different qualities ($\rho_B \ll \rho_A$), contrary to our method which solves such easier collective decision making problems faster (see Figure 3b). Similarly to our study, their control algorithm is supported by a thorough theoretical investigation [20, 26, 31]. Differently from the weighted voter model, consensus on the optimal opinion is only guaranteed for initial proportions of preferences for the best opinion that are greater than a certain threshold. Similar considerations apply also to the work of Brutschy et al. [2] which substitutes the majority rule used by Montes de Oca et al. with the k-unanimity rule.

6. CONCLUSIONS

In this paper we have introduced the weighted voter model which is a simple control algorithm for self-organized collective decision making in distributed systems. The main advantages of this approach are its increasing decision accuracy with increasing system size, logarithmic scalability of the consensus time, and robustness to noisy assessments of site qualities. We have reported an ODE model and a master equation model that allow precise predictions of critical performance characteristics such as decision accuracy and consensus time. Using the ODE model we are able to guarantee convergence to the best decision in the thermodynamic limit, while using Gillespie simulations of the master equation we are able to give guarantees for accuracy and consensus time for finite size systems. We empirically investigated the robustness of the weighted voter model using agent-based simulations. We conjecture that this robustness is a result of the distributed collective decision making process that does not rely on the individual agents' cognitive skill to directly compare and discriminate sites (as observed in natural swarms [8, 33]).

Currently, the weighted voter model is implemented and tested on a physical robot swarm with promising preliminary results. The models and simulations reported in this paper will be tested against this real-world implementation. Our conjecture about the robustness of the proposed control algorithm will be further investigated through an empirical evaluation against the direct comparison of opinions as well as other decision rules. Future theoretical work will be focused on decision problems with more alternative opinions.

7. ACKNOWLEDGMENTS

This work was partially supported by the European Research Council through the ERC Advanced Grant "E-SWARM: Engineering Swarm Intelligence Systems" (contract 246939). Marco Dorigo acknowledges support from the Belgian F.R.S.–FNRS.

8. REFERENCES

[1] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.

[2] A. Brutschy, A. Scheidler, E. Ferrante, M. Dorigo, and M. Birattari. "Can ants inspire robots?" Self-organized decision making in robotic swarms. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 4272–4273. IEEE Press, 2012.

[3] A. Campo, S. Garnier, O. Dédriche, M. Zekkri, and M. Dorigo. Self-organized discrimination of resources. *PLoS ONE*, 6(5):e19888, 2010.

[4] C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Rev. Mod. Phys.*, 81:591–646, 2009.

[5] P. Clifford and A. Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.

[6] G. A. Di Caro and M. Dorigo. AntNet: Distributed stigmergetic control for communications networks. *Journal of Artificial Intelligence Research*, 9:317–365, 1998.

[7] F. Ducatelle, G. A. Di Caro, and L. M. Gambardella. Principles and applications of swarm intelligence for adaptive routing in telecommunications networks. *Swarm Intelligence*, 4(3):173–198, 2010.

[8] S. C. Edwards and S. C. Pratt. Rationality in collective decision-making by ant colonies. *Proceedings of the Royal Society B: Biological Sciences*, 276(1673):3655–3661, 2009.

[9] N. R. Franks, S. C. Pratt, E. B. Mallon, N. F. Britton, and D. J. T. Sumpter. Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1427):1567–1583, 2002.

[10] S. Galam. Sociophysics: A review of Galam models. *International Journal of Modern Physics C*, 19(03):409–440, 2008.

[11] V. Gazi and K. M. Passino. Stability analysis of swarms. *IEEE Transactions on Automatic Control*, 48(4):692–697, 2003.

[12] V. Gazi and K. M. Passino. A class of attractions/repulsion functions for stable swarm aggregations. *International Journal of Control*, 77(18):1567–1579, 2004.

[13] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977.

[14] H. Hamann. A reductionist approach to hypothesis-catching for the analysis of self-organizing decision- making systems. In *7th IEEE Int. Conf. on Self-Adaptive and Self- Organizing Systems (SASO 2013)*, pages 227–236. IEEE Press, 2013.

[15] H. Hamann. Towards swarm calculus: Urn models of collective decisions and universal properties of swarm performance. *Swarm Intelligence*, 7(2-3):145–172, 2013.

[16] H. Hamann, T. Schmickl, H. Wörn, and K. Crailsheim. Analysis of emergent symmetry breaking in collective decision making. *Neural Computing and Applications*, 21(2):207–218, 2012.

[17] S. Kernbach, R. Thenius, O. Kernbach, and T. Schmickl. Re-embodiment of honeybee aggregation behavior in an artificial micro-robotic system. *Adaptive Behavior*, 17(3):237–259, 2009.

[18] T. M. Liggett. *Stochastic interacting systems: contact, voter and exclusion processes*, volume 324 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1999.

[19] K. L. Mills. A brief survey of self-organization in wireless sensor networks. *Wireless Communications and Mobile Computing*, 7(7):823–834, 2007.

[20] M. Montes de Oca, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari, and M. Dorigo. Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5:305–327, 2011.

[21] M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563, 2006.

[22] C. Parker and H. Zhang. Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem. *IEEE/ASME Transactions on Mechatronics*, 14(2), 2009.

[23] C. Parker and H. Zhang. Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence*, 4:199–220, 2010.

[24] K. M. Passino and T. D. Seeley. Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. *Behavioral Ecology and Sociobiology*, 59(3):427–442, 2006.

[25] T. M. Schaerf, J. C. Makinson, M. R. Myerscough, and M. Beekman. Do small swarms have an advantage when house hunting? The effect of swarm size on nest-site selection by *apis mellifera*. *Journal of The Royal Society Interface*, 10(87), 2013.

[26] A. Scheidler. Dynamics of majority rule with differential latencies. *Phys. Rev. E*, 83:031116, 2011.

[27] I. Schizas, G. Giannakis, S. Roumeliotis, and A. Ribeiro. Consensus in ad hoc wsns with noisy links — part II: Distributed estimation and smoothing of random signals. *IEEE Transactions on Signal Processing*, 56(4):1650–1666, 2008.

[28] I. Schizas, A. Ribeiro, and G. Giannakis. Consensus in ad hoc wsns with noisy links — part I: Distributed estimation of deterministic signals. *IEEE Transactions on Signal Processing*, 56(1):350–364, 2008.

[29] T. D. Seeley. *Honeybee Democracy*. Princeton University Press, 2010.

[30] R. Toral and C. J. Tessone. Finite size effects in the dynamics of opinion formation. *Communications in Computational Physics*, 2(2):177–195, 2007.

[31] G. Valentini, M. Birattari, and M. Dorigo. Majority rule with differential latency: An absorbing Markov chain to model consensus. In *Proceedings of the European Conference on Complex Systems 2012*, Springer Proceedings in Complexity, pages 651–658. Springer, 2013.

[32] N. G. van Kampen. *Stochastic processes in physics and chemistry*. Elsevier, Amsterdam, NL, 1992.

[33] P. K. Visscher and S. Camazine. Collective decisions and cognition in bees. *Nature*, 397(6718):400, 1999.