Host Selection through Collective Decision

FABRICE SAFFRE
ETISALAT BT Innovation Centre (EBTIC)
Khalifa University

and

AISTIS SIMAITIS
School of Computing
University of Leeds

In this paper, we present a collective decision-making framework inspired by biological swarms and capable of supporting the emergence of a consensus within a population of agents in the absence of environment-mediated communication (stigmergy). Instead, amplification is the result of the variation of a confidence index, stored in individual memory and providing each agent with a statistical estimate of the current popularity of its preferred choice within the whole population. We explore the fundamental properties of our framework using a combination of analytical and numerical methods. We then use Monte Carlo simulation to investigate its applicability to host selection in the presence of multiple alternatives, a problem found in application migration scenarios. The advantages of self-organisation and the use of statistically predictive methods in this context are also discussed.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems—*algorithms*; *design*; *experimentation*

General Terms: Algorithms, Design, Experimentation

Additional Key Words and Phrases: Agent-based systems, collective decision-making

1. INTRODUCTION

Collective decision-making, or the process whereby a group of individuals reach consensus about selecting a common course of action in the absence of any clear leadership or hierarchical structure, has been extensively studied in social and biological sciences. Over the last two decades, the usefulness of such a mechanism in technology has gradually become more evident, as complex systems of interacting autonomous entities have become more widespread. An alternative to centralised architecture, such systems tend to be more robust, more adaptive and more scalable, mostly due to having lower maintenance and management requirements. These properties make them a particularly attractive option for large-scale deployments in dynamic environments, be it sensor networks, fleet of mobile robots or, more pragmatically, distributed applications or services.

However, relying on distributed complex adaptive systems to perform a useful

Author's address: F. Saffre, ETISALAT BT Innovation Centre (EBTIC), Khalifa University, Abu Dhabi P.O. Box 127788, Abu Dhabi, UAE.

task in the absence of central control requires the assurance that their constituents will be capable of operating harmoniously as a whole, i.e. that an efficient and appropriate global response will emerge in response to a particular problem or challenge (e.g. a specific configuration of the environment). Clearly, collective decision-making is an example of such emergent, fully decentralised yet system-wide phenomenon leading to coordinated action.

In one of its simplest forms, collective decision-making involves reaching agreement on selecting one of many options exclusively through peer-to-peer interactions. A concrete example of this can be found in the case of a population of autonomous software agents needing to choose a common hosting site so as to reduce communication delays. Assuming a plurality of potential hosts, how can the vast majority reach agreement as to which one to select and migrate to? More generally, what locally executable algorithm can lead to a common decision in a wide variety of circumstances? For instance, one variant of the problem might involve hosts of identical average value, making all solutions equivalent and therefore precluding the use of some quality-based reasoning. Another, somewhat opposite situation would involve sites of variable quality and agents who only have a very partial view of all available options: how can those who have successfully identified the best host (possibly a minority) be made capable of steering the collective decision into the right direction?

The primary objective of this paper is to make our new variant of a collective decision-making framework available to the research community by providing a clear and thorough description. Beyond that, we chose to focus on exposing the complex dynamics of the algorithm and on leveraging the understanding of these fundamental properties to inform the choice of an efficient design. We illustrate this methodology by applying it to the case of identical options and showing that the selection of suitable parameter values is critical to the success or failure of the collective decision process. We also demonstrate that a principled study of dynamical properties (e.g. stability analysis and systematic search for multistationarity) is necessary but not sufficient to predict the outcome of the decision process in a real-world application, due to the stochastic effects resulting from a discrete implementation.

The remainder of this paper is structured as follows: section 2 provides the overview of the related work, in section 3, we describe our variant of the local decision-making algorithm and translate it into a system of differential equations. In section 4, we use this mathematical framework to analyse system properties (influence of key parameter values) in a continuous approximation. In section 5, we present the results of a simulated implementation and show the influence of stochastic effects by comparing them with predictions from the continuous model. In section 6, we demonstrate how our methodology can be used to select appropriate parameter values when critical problem characteristics (e.g. number of available choices) are known at design time. Finally, in section 7, we discuss the meaning of our analytical and numerical results, as well as directions for future work.

2. RELATED WORK

The phenomenon of collective decision-making has been studied in many social or gregarious organisms such as ants [Beckers et al. 1990; Mallon et al. 2001], bees [Seeley et al. 1991; 1999; 2003], cockroaches [Ame et al. 2006], spiders [Saffre et al. 2000], and mammals (e.g. wildebeest [Holdo et al. 2009; Gueron and Levin 1993] and reindeer [Bergerud 2000]). The need for collective decision-making arises from different triggers in different species, for example, foraging site selection [Seeley et al. 1991; Beckers et al. 1990], nest site selection [Seeley and Morse 1978; 2004], and migration [Dyer and T.D. 1994; Holdo et al. 2009]. To this diversity of objectives or functions corresponds a variety of decision-making mechanisms [Conradt and Roper 2005]. Some groups of animals rely on clear leadership [Couzin et al. 2005] while others exploit emergent properties within a community of 'identical' agents in the absence of hierarchy or centralization. Our focus has been on the latter, i.e. on systems in which the decision emerges solely from socially equivalent individuals interacting with each other and with their environment. Such a behaviour has been observed, for instance, in nest selection by honey bee swarms [Seeley and Visscher 2004; 2003]. After the swarm leaves its former hive, it clusters on a tree branch until it eventually moves to a new home. Martin Lindauer [1951; 1953; 1955] was the first to discover that a collegial decision in favour of a particular nest site emerged from scout bees performing waggle dances on the swarm surface.

The decision-making process of honey bees must have three main properties [Seeley and Visscher 2004]: firstly, the decision must be *accurate* - a colony's success depends critically on it occupying a cavity that is sufficiently spacious; secondly, it must achieve a *speedy* decision - additional time that a swarm spends exposed lowers its energy reserves and increases its chances to be destroyed by rain; thirdly, it must achieve a *unified* decision - a split decision would lead to swarm fragmentation which would be disastrous in most cases.

These properties of the natural collective decision-making process found in social insects make it a very attractive proposition to address a variety of problems encountered in Information Technology (IT) such as, for instance, process migration in computational grids, resource allocation in server farms, and peer-to-peer networks.

Other aspects of decentralized consensus-building have been investigated since the 1980s, such as the necessary conditions for reaching agreement about distributed information within a partly unreliable set of communicating agents [Pease et al. 1980; Fischer et al. 1985]. More recently and more closely related to the present work, other authors have proposed various algorithms and protocols to optimally map application or service components onto a population of potential hosts [Koshi et al. 2009], including using ant algorithms [Musunoori and Horn 2007; Csorba et al. 2009]. However, these recent advances put more emphasis on performance, relative placement and substrate-based communication (stigmergy) than on reaching consensus and comparing the intrinsic quality of prospective hosts.

Many researchers in computer science have addressed resource allocation [Ardagna et al. 2007], load balancing [Montresor et al. 2002; Chow and Kwok 2002; Wolf and Yu 2001], routing [Heusse et al. 1998], process migration [Gupta and Srimani 2003; Fu and Xu 2005], sorting and clustering [Handl and Meyer 2002], and system management [Messig and Goscinski 2007; Heimfarth and Janacik 2006; Shen et al. 2002] problems using multi-agent systems and autonomous agents. They have shown that the strength of these systems lies in their ability to adapt to a changing environment and in their built-in fault tolerance [Heusse et al. 1998; Schaerf et al. 1995]. These properties emerge from the fact that the success of the system does not depend on the choices and/or actions of one individual (single point of failure) but on the many local interactions between 'identical' agents which support the self-organization of the group as a whole. There have been numerous studies of self-organisation in both biological [Bonabeau et al. 1997; Sumpter and Pratt 2009; Ame et al. 2006; Gueron and Levin 1993] and artificial systems [Saffre et al. 2009; Dorigo and Stuetzle 2004] which all come to the similar conclusion that group dynamics can provide a very robust and efficient framework for resource allocation and selection processes [Dorigo and Stuetzle 2004; Kennedy et al. 2001].

In this paper we propose a simple social insects-inspired agent-based framework for collective decision-making which could be effectively applied to resource allocation, process and service migration, and similar problems. It is worth noting that our method is closely related to that which was used by other authors to support consensus building in collective robotics [Parker and Zhang 2009] or quorum sensing (a somewhat different problem) in mobile ad-hoc networks [Peysakhov et al. 2006]. However, to the best of our knowledge, this is the first time that such a 'stigmery-less' technique is applied to host selection by a population of software agents, a problem which, unlike the above examples, is not fundamentally defined by spatial constraints but may involve measuring other evaluation criteria such as site quality. This is why we explicitly included the corresponding parameters (η and λ) and variable (Q) in the formal presentation of the model even though their influence was not investigated in this paper. The exhaustive exploration of our model's properties is a substantial endeavor that will require considerable future work by ourselves or by other authors.

3. SYSTEM MODEL FOR COLLECTIVE DECISION MAKING

In this section we describe our agent-based framework for collective decision-making. As already mentioned in the related work section, we use the principles of collective decision-making found in social insect colonies because they have desirable generic properties for real world systems such as accuracy, speed and fault tolerance. Furthermore, and critically for a fully decentralised multi-agent system, it is consensual, i.e. the whole colony eventually chooses in favour of a single option, avoiding dispersion. Biological examples demonstrate that this can be achieved using relatively simple interaction mechanisms (e.g. variations of waggle dance behaviour depending on quality of/confidence in the resource or stigmergetic modification of the environment).

This section is structured as follows: the 'terms' section provides the necessary definitions to understand the terminology used in the framework; the 'agent behaviour' section goes into the details of the rules governing the behaviour of individual agents; the 'system description' section presents the generic differential equations framework that was used to model the collective decision process.

3.1 Terms

—*Scout agent* - an agent whose role is to explore the environment and promote the choice of what it perceives as the best hosting site among its peers (recruitment).

—*Environment* - the network of hosting sites accessible to scout agents.

—*Hosting site* - a site capable of hosting software agents. Each hosting site has specific characteristics that determine its quality.

—*Site quality* - a measure of a hosting site's attraction for a scout agent (can be agent-specific, for instance if functional requirements dictate that different characteristics of the host are relevant to different agents).

—*Confidence level* - a measure of how confident a scout agent is that its preferred hosting site should be selected by the colony.

—*Exploration probability* - probability that a scout agent chooses to make an 'exploratory move', as opposed to initiating an interaction with another scout agent.

—*Non-linearity parameters* - there are three non-linearity parameters used in equations describing agent behaviour (and in the corresponding local decision rules):

  —η - affects the relative perceived quality of sites when a scout agent is comparing one to the other (since the remainder of this paper deals exclusively with the case in which all sites are of equal quality, this parameter is only included for completeness and its value is irrelevant).

  —λ - affects the relative perceived quality when two scout agents are comparing their preferred sites (same comment as above).

  —γ - affects the relative weight of the confidence level when two scout agents are comparing their preferred sites.

3.2 Agent behaviour

The goal of the scout agents is to explore the environment, gathering information about the quality of potential hosting sites, and to collectively select the best one through a fully distributed interaction mechanism. In line with the parsimony of the algorithm, the behavioural repertoire of scout agents is limited to performing two actions:

—*Explore the environment.*
The agent chooses one hosting site from the environment at random, migrates there, then evaluates its quality. If the target of this exploration move is its current favourite, the scout simply updates its record to reflect the new perceived quality. If the site is different, the scout will change its preferred site with probability P_x given by (1) where Q_x is the quality of the hosting site currently under evaluation, Q_y is the quality of the scout agent's preferred site and η is a non-linearity parameter:

$$P_x = \frac{Q_x^\eta}{Q_x^\eta + Q_y^\eta} \tag{1}$$

—*Compare preferred hosting sites with another scout agent.*
Agent A with preferred hosting site x randomly contacts scout agent B with favourite hosting site y. If agents have the same preference ($x = y$), they both increment their confidence level. If agents have different preferred hosting sites

($x \neq y$) they will *both* choose in favour of site x with probability P_A given by (2) where Q_x and Q_y is the perceived quality of their respective favourite site x and y, l_A and l_B are the confidence levels of scouts A and B, and λ, γ are non-linearity parameters:

$$P_A = \frac{Q_x^\lambda l_A^\gamma}{Q_x^\lambda l_A^\gamma + Q_y^\lambda l_B^\gamma} \quad (2)$$

Symmetrically, they will *both* choose in favour of site y with probability $1 - P_A$. The interaction concludes with the scout whose site has been selected incrementing its confidence level and the scout who has 'lost' the contest resetting its confidence level to the lowest possible value. In the current implementation, this interaction can only take place between scouts who have both chosen not to explore the environment on this time-step.

It may be worth emphasising at this stage that no global knowledge or centralised infrastructure is required for this framework to operate. Indeed, it is not necessary for a scout to have access to *all* potential hosting sites or to know the current address of *all* other agents for the collective decision process to function. The exploration of the environment could easily take the form of a random walk, thereby requiring only the ability to access or query a small sub-set of all potential hosts at any given time. As for information exchanges, although they would of course require a shared communication medium, this does not imply that interacting scouts need to know each other's explicit address.

In practice, we envisage that both exploration and comparison could take form of a query that is probabilistically propagated or answered every time that it reaches another peer in a P2P framework (said peer being either a potential host or another scout, respectively). The local bias introduced by such method could easily be alleviated (if never completely eliminated) by making the probability of answering low enough to ensure that the message would statistically propagate sufficiently deep into the overlay. Although this would obviously create some communication overhead, there is no reason to think that it would be any worse than in the case of existing, working P2P applications. If anything, it would be made substantially lower by the fact that the query could advantageously be propagated on a point-to-point basis (each relay forwarding it to just one randomly selected neighbour) instead of by flooding.

However, for the sake of clarity and to make the model analytically tractable, such practical implementation considerations will be disregarded in the remainder of this paper. In other words: we will make the simplifying assumption that every site and every scout in the system is *in principle* equally likely to be selected as the target of an exploration or comparison action.

The pseudocode for scouts' behaviour is provided in Algorithm 1.

Algorithm 1 Scout behaviour

```
procedure Run
    if scout is at the nest site then
        random ← random number from (0,1)
        if random ≤ P_exp then
            move to random site
        else
            select another scout at random
            if same site preferred then
                increase confidence levels
            else
                compare notes using Equation 2
                update favourite sites
                update confidence levels
            end if
        end if
    else
        compare the site with the favourite using Equation 1
        update favourite site
        update confidence level
    end if
    Run()
end procedure
```

3.3 System description

The system can be described using a set of parameters and variables representing an environment consisting of a number of hosting sites and a scout population consisting of a number of scout agents.

3.3.1 *Environment parameters*

—k - number of hosting sites.

—$\{C_1, C_2, \ldots, C_k\}$ - set of characteristics of every hosting site in the environment.

3.3.2 *Scout agent parameters*

—n - the number of scout agents in the system (scout population).

—P_{exp} - exploration probability: the probability that an agent chooses to migrate to and evaluate the quality of a randomly selected hosting site. The probability that it will choose to compare preferred hosting sites with another scout is therefore $1 - P_{exp}$.

—l - the confidence level of the scout agent ($l \in \{1, 2, \ldots, l_{max}\}$).

—l_{max} - the maximum possible value for the confidence level. After a scout's confidence level has reached l_{max} it can no longer be incremented.

—$f(C)$ - the function used by scout agents to measure the quality Q of a hosting site ($Q_x = f(C_x)$)

Having defined the parameters for the environment and scout agents, the system state S can now be described by (3):

$$S = \{x_{11}, \ldots, x_{1l_{max}}, \ldots, x_{k1}, \ldots, x_{kl_{max}}\} \tag{3}$$

Where x_{ij} is the number scout agents preferring hosting site i having confidence level j and $\sum x_{ij} = n$ ($i \in \{1, 2, \ldots, k\}$, $j \in \{1, 2, \ldots, l_{max}\}$).

Using equations (1) and (2) we can now construct the differential equation for every x_{ij} ($i \in \{1, 2, \ldots, k\}$, $j \in \{1, 2, \ldots, l_{max}\}$). The differential equation for x_{ij} in its general form is shown in (4):

$$\begin{aligned}\frac{dx_{ij}}{dt} = & -(P_{x_{ij} \to x_{ij+1}} + \sum_{z}^{k(z \neq i)} P_{x_{ij} \to x_{z1}}) x_{ij} \\ & + P_{x_{ij-1} \to x_{ij}} x_{ij-1} \\ & + \sum_{z}^{k(z \neq i)} \sum_{l}^{l_{max}} P_{x_{zl} \to x_{ij}} x_{zl}\end{aligned} \quad (4)$$

—$P_{x_{ij} \to x_{ij+1}}$ is the probability that a scout agent which is currently in favour of site i and has a confidence level of j will increase its confidence level. As per (5) a scout agent will increase its confidence level when:
  - —It is exploring a potential hosting site in the environment, but the outcome of the evaluation conducted using (1) has led the agent to keep its currently preferred site i as its favourite.
  - —It contacts (or is contacted by) another scout agent which also prefers hosting site i.
  - —It contacts (or is contacted by) another scout agent in favour of a different hosting site, but the resolution of the contest by (2) has led to site i being chosen by both scouts (i.e. 'win').

$$\begin{aligned}P_{x_{ij} \to x_{ij+1}} = & \frac{P_{exp}}{k} \sum_{z}^{k(z \neq i)} \frac{Q_i^\eta}{Q_i^\eta + Q_z^\eta} \\ & + \frac{1 - P_{exp}}{n} (\sum_{l}^{l_{max}} x_{il} + \sum_{z}^{k(z \neq i)} \sum_{l}^{l_{max}} \frac{Q_i^\lambda j^\gamma}{Q_i^\lambda j^\gamma + Q_z^\lambda l^\gamma} x_{zl})\end{aligned} \quad (5)$$

Please note that $P_{x_{ij} \to x_{ij+1}} = 0$ when $j = l_{max}$ and $P_{x_{ij-1} \to x_{ij}} = 0$ when $j = 1$.

—$P_{x_{ij} \to x_{z1}}$ is the probability that a scout agent will change its preferred hosting site from i to z. As per (6) a scout will change its preferred site when:
  - —It is exploring potential hosting site z and, based on the outcome of the evaluation conducted using (1), it changes its preferred site from i to z.
  - —It contacts (or is contacted by) another scout agent in favour of hosting site z, and the resolution of the contest by (2) has led to site z being chosen by both scouts (i.e. 'lose').

$$P_{x_{ij} \to x_{z1}} = P_{exp} \frac{Q_z^\eta}{Q_z^\eta + Q_i^\eta} + \frac{1 - P_{exp}}{n} \sum_{l}^{l_{max}} \frac{Q_z^\lambda l^\gamma}{Q_z^\lambda l^\gamma + Q_i^\lambda j^\gamma} x_{zl} \quad (6)$$

4. MODEL ANALYSIS

In this paper we focus on the role played by agent interactions in collective decision making, i.e. on the mechanisms leading to the emergence of a 'majority vote' or consensus. In order to decouple these dynamics from the influence of external factors, we consider the special case in which all hosting sites are identical ($C_1 = C_2 = ... = C_k$) and all scout agents share the same evaluation function, leading to all perceived qualities of all hosting sites being equal throughout the scout population ($Q_1 = Q_2 = ... = Q_k$). It is important to emphasise that, even though it can be regarded as unrealistic with respect to practical implementation, the use of such a hypothetical scenario is critical to the in-depth understanding of the fundamental properties of complex distributed systems. In turn, this understanding underpins our ability to make quantitative statistical predictions about the collective behaviour of a population of agents individually governed by specific local rules and parameter values. In doing so, it creates the right conditions for an informed use of self-organising properties, which could otherwise be perceived as too unreliable to support a real-world application. Furthermore, in this particular case, not having an external bias guiding the decision process (such as one more favourable hosting site) arguably makes for a tougher problem to solve collectively and so for a better test of the framework's capability.

4.1 Emergence of the decision

The simplest possible case involves choosing between two hosting sites of identical quality (binary choice). Figure 1 shows the results of the simultaneous solution of the system of differential equations presented in the 'system description' section, using the Forward Euler Method (FEM). In this case, the maximum confidence level l_{max} of the scout agents was set to 2, which can be interpreted as them having only two possible states: '1=early adopter' or '2=confirmed supporter' of a particular site. We then studied the effect of the γ parameter on the emergence of a decision, for 3 different values of P_{exp}. In initial conditions, scout agents are distributed evenly between the four possible states, i.e. two sites multiplied by two confidence levels ($x_{11} = x_{12} = x_{21} = x_{22}$). A small quantity ϵ representing noise is then added to x_{11} and the system of equations is iterated using FEM until the system reaches steady state, i.e. the invariant distribution had been found ($\forall x_{ij} : \frac{dx_{ij}}{dt} = 0$).

In the absence of any difference in site quality, γ is the key parameter influencing the outcome of scout interaction, as its value determines the advantage to the agent with the highest confidence level, i.e. the only possible bias susceptible to facilitate the emergence of a consensus. The results shown in figure 1 are therefore not surprising as they confirm the intuitive conclusion that the higher the chances of a 'confirmed supporter' winning over an 'early adopter', the clearer the decision - the larger the fraction of scouts in favour of the winning site at steady state. Similarly, one can easily understand why, as P_{exp} increases, it becomes harder for a clear majority to emerge, simply because of the lower frequency of interaction and higher 'spontaneous' dispersion resulting from random exploration. The key finding here is clearly the quantitative study of the interplay between P_{exp} and γ, showing which higher value of the latter is needed to compensate for a given increase in the former in order to achieve a target majority, or indeed any decision at all (bifurcation threshold).

4.2 Decision threshold

Figure 2 illustrates an effect complementary to that discussed in the previous section, namely the presence of a finite basin of attraction around the homogeneous distribution equilibrium for larger environments ($k > 2$) and for given values of parameters P_{exp} and γ. This effectively means that there are regions of the $k/\gamma/P_{exp}$

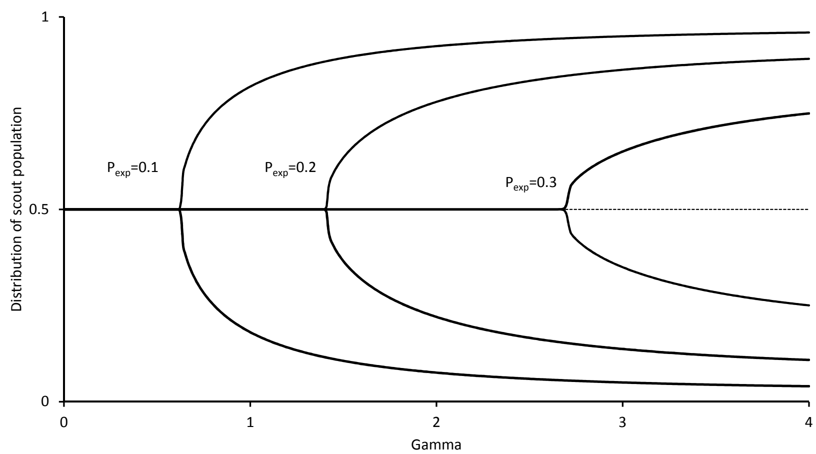

Fig. 1. Invariant distribution of the scout population as a function of γ for three values of the exploration probability P_{exp}.

parameter space where two stable solutions coexist, one centered on the homogeneous distribution (decision failed) and one on a distribution biased in favour of a winning site (successful decision). Figure 2 indicates that the divider between the two basins of attraction (unstable equilibrium), expressed in terms of the ratio between the population in favour of one site and the average number of scouts in favour of other sites, is a linear function of system size k, and that the slope is inversely proportional to the value of γ. This result is important because it demonstrates that, in a real (discrete) implementation involving a realistic number of hosting sites ($k \gg 2$), a designer would need to ensure that there is an adequate level of noise for this threshold to be statistically reachable, failing which the scouts will be unable to reach consensus.

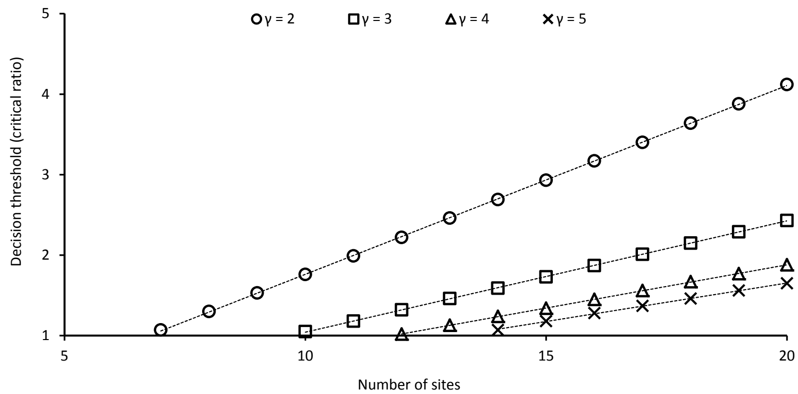

Fig. 2. The minimal ratios between the scout population preferring the 'winning' site at the beginning with which the decision emerges (winning site stays the winner).

5. SIMULATED IMPLEMENTATION AND STOCHASTIC EFFECTS

In this section, we stop focusing on fundamental properties in terms of steady state and direct our attention towards the dynamics of the decision-making process and their influence on system history. This is clearly of paramount importance for any real implementation as a set of parameter values which, e.g., lead to a clear and accurate choice but only after a prohibitively long convergence time would likely be unusable in practice. Figure 3 and 4 illustrate the decision-making dynamics within a population of scouts faced with a choice between three hosting sites of identical quality, for changing values of γ (recruitment non-linearity) and P_{exp} (exploration probability) respectively. The diagrams show the frequency distribution of observations as a function of the fraction of the population on all three sites (one data point per site). As expected, for low values of γ (fig. 3) and high values of P_{exp} (fig. 4), the single peak centred on 1/3 indicates that the scouts are evenly distributed between all potential sites, i.e. failed to make a collective decision. Conversely, at the opposite end of the range of tested parameter values, the clear U-shaped distribution (i.e. two peaks, centered on $P_{exp}/2$ and 1-P_{exp}, in figure 3) indicates a clear-cut decision, with the two losing sites being visited on average by half of the residual explorers each ($P_{exp}/2$) and the winning host (one third of all data points) rassembling the remaining fraction (1-P_{exp}). Compared to the information contained in fig. 1, the results of this Monte Carlo simulation provide a better insight into the 'shape' of the transition. For instance, even though it is very sharp for increasing values of γ (fig. 3), it appears slower for decreasing values of P_{exp}, exhibiting a very gradual 'flattening' of the single peak before reaching the decision threshold (fig. 4).

As for the evolution over time, results allow us to better understand the interplay between the parameter values affecting the decision functions themselves (in

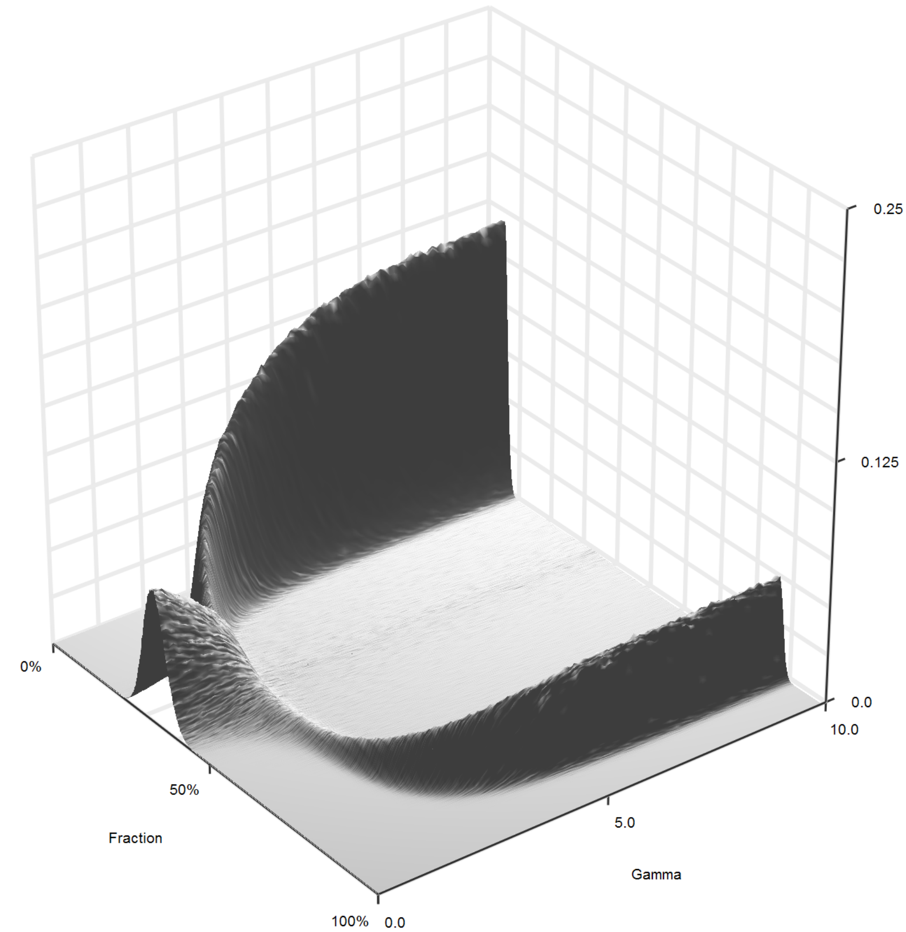

Fig. 3. Frequency distribution of observations as a function of the fraction of the population in favour of any one of three sites, for increasing values of γ ($k = 3$, $P_{exp} = 0.1$, $l_{max} = 2$). See text for details.

the case of sites of identical value, γ and P_{exp}) and other system characteristics, such as scout population size n or number of sites k. For instance, figure 5 illustrates the effect of using fine ($n = 1000$) or coarse-grained ($n = 100$) discretisation levels in a larger environment ($k = 100$). It clearly demonstrates that reducing the number of scouts is an efficient way of speeding-up the decision process when the identification and selection of a high quality site is not an issue (as in the present scenario). However, it is obvious that a lower number of scouts increases the noise level and impacts negatively on their collective sampling ability, potentially introducing conflicting requirements that would need to be balanced in any real deployment.

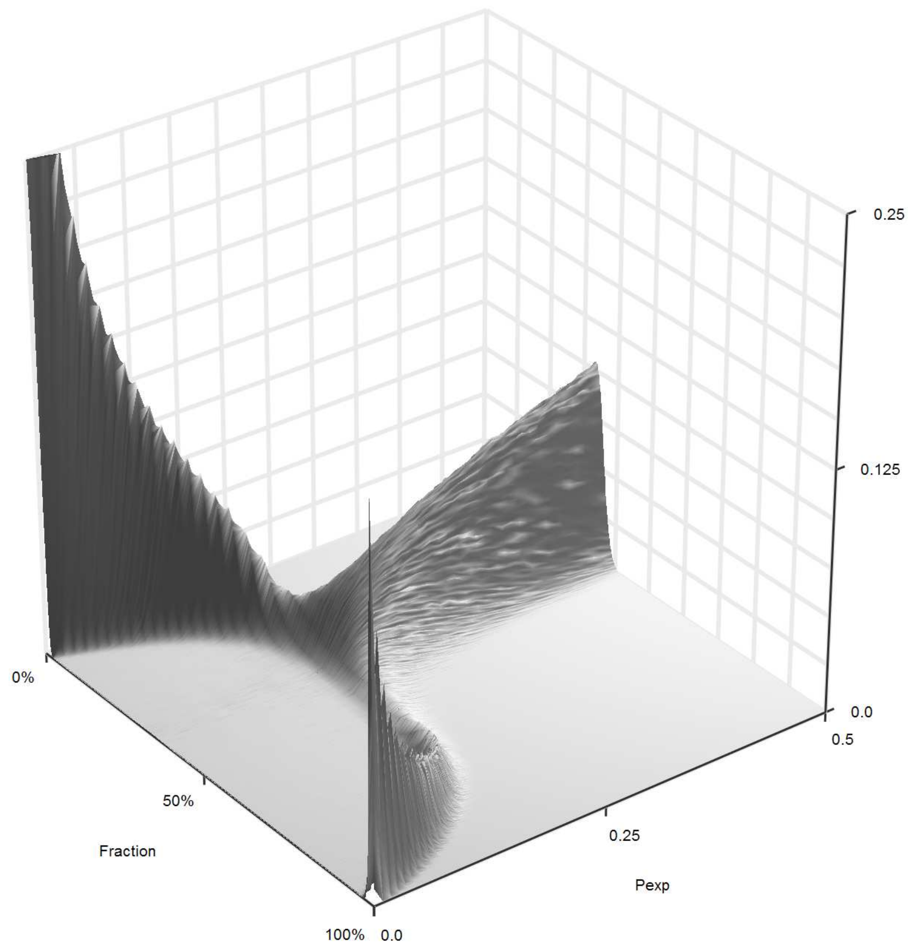

Fig. 4. Frequency distribution of observations as a function of the fraction of the population in favour of any one of three sites, for increasing values of P_{exp} ($k = 3$, $\gamma = 2$, $l_{max} = 2$). See text for details.

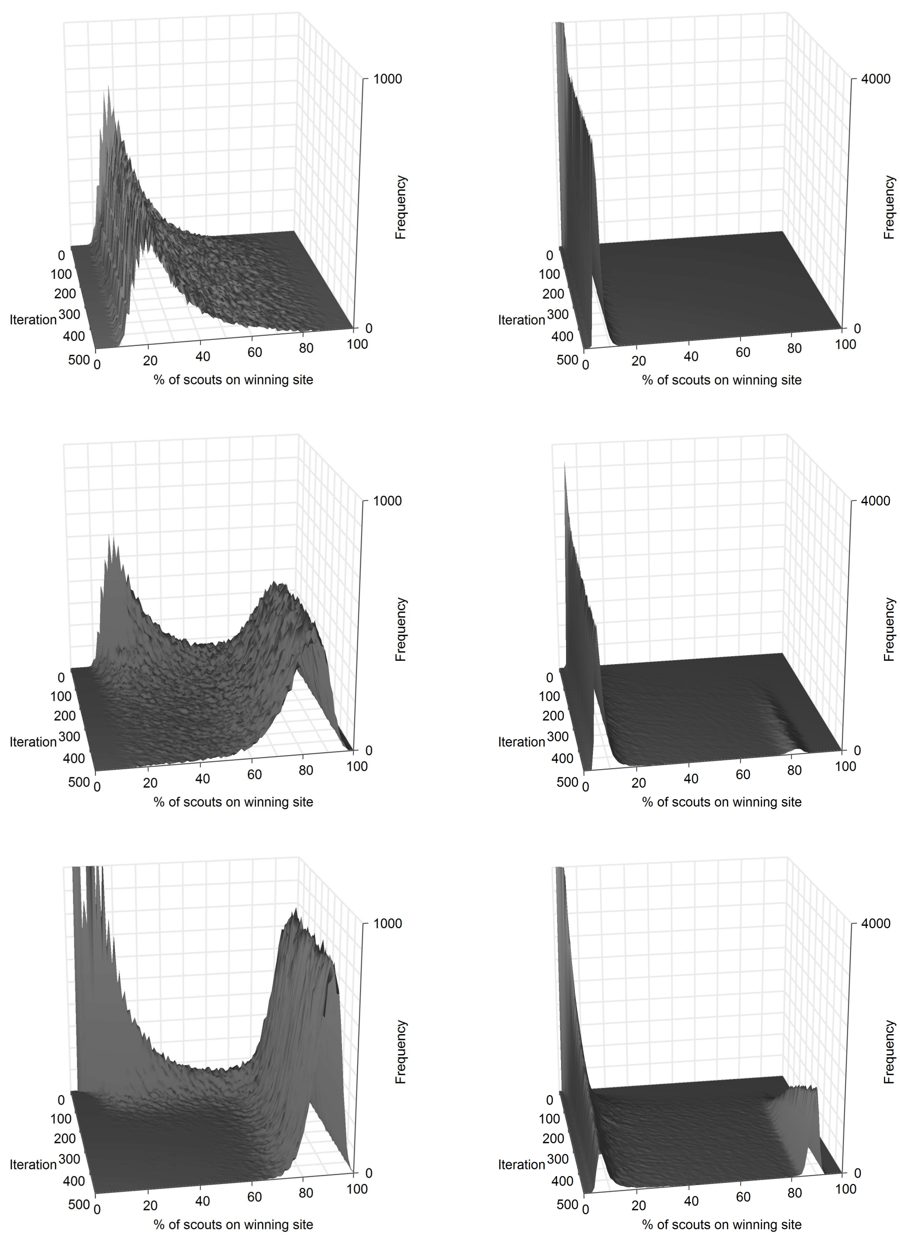

Fig. 5. Frequency distribution of observations as a function of the percentage of scouts in favour of the winning site. Evolution over time for multiple values of γ and n ($k = 100$, $P_{exp} = 0.1$, $l_{max} = 2.0$). Left column $n = 100$, right - $n = 1000$. First row $\gamma = 1$, second - $\gamma = 3$, third - $\gamma = 3$.

6. SYSTEM PROPERTIES AND PARAMETER VALUES

In this section we examine two fundamental performance measures that will determine our framework's ability to tackle real-world problems, namely convergence speed and decision stability. In doing so, we also demonstrate how it is possible to model and statistically predict system behaviour based on known characteristics (in this case, population size), which would allow for efficient testing of multiple candidate sets of parameter values prior to deployment.

Figure 6 shows that the probability not to have reached a decision is a decreasing sigmoid function of time, the characteristics of which ('half-life' and slope) are jointly determined by the population and environment sizes (n and k) and by the value of the parameters affecting scout behaviour (γ, P_{exp} and l_{max}).

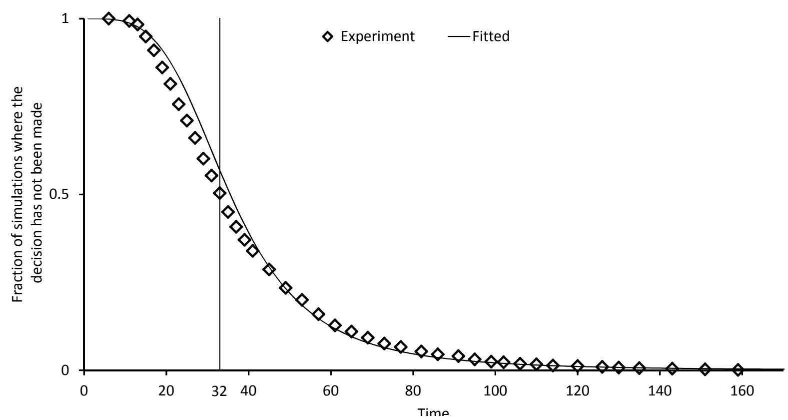

Fig. 6. Decision speed ($k = 100$, $n = 50$, $P_{exp} = 0.1$, $\gamma = 2.0$, $l_{max} = 2.0$). See text for details.

Another key property is the stability of the decision, i.e. the typical duration of the consensus which, in a finite population of scouts featuring stochastic behaviour, will always eventually disappear as a result of random fluctuations. It is intuitively obvious that decision will be faster if there are fewer scouts: for instance, in the trivial case of a single agent, decision is instantaneous. However its choice will also change almost instantly as the scout keeps exploring other hosting sites and the decision cannot be stabilised by interaction with others (positive feedback). We therefore sought to characterise the stability of the decision, i.e. the duration of the time interval between the emergence of the first consensus and its replacement by another one. Figure 7 shows the 'survival curve' of numerical experiments as a function of time, the condition for termination being that the first collective 'change of mind' has occurred. It clearly exhibits the signature of an exponential decay, i.e. there is a typical 'half-life' but variability is very high.

Having identified a suitably simple set of variables to describe system behaviour, namely the characteristic 'half-life' of a decision and 'half-life' and slope of the

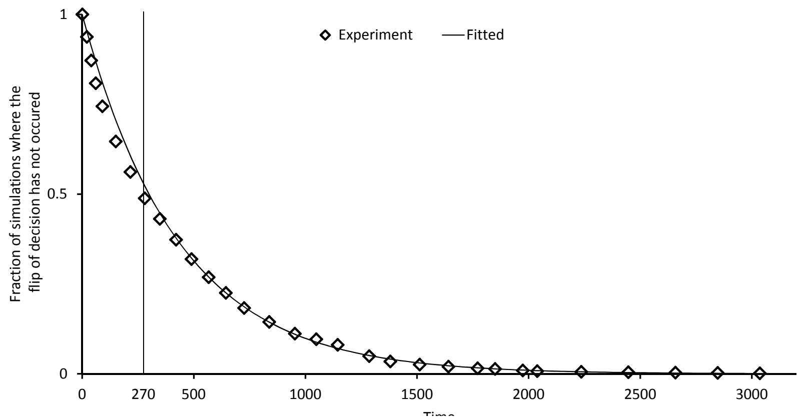

Fig. 7. Decision stability ($k = 100$, $n = 50$, $P_{exp} = 0.1$, $\gamma = 2.0$, $l_{max} = 2.0$). See text for details.

convergence period, we can use these to quantify the performance of the whole framework as a function of some other parameters. For instance, figure 8 shows the evolution of the two 'half-life' variables as a function of the population size n. This particular result illustrates a very desirable property of our framework, namely that the stability of the decision increases exponentially with the number of scouts, while the convergence time only does so linearly. In other words: by increasing population size, it is possible to promote a clear and long-lived consensus in exchange for a comparatively small delay in the emergence of a decision.

7. DISCUSSION

In this paper, we have explored the potential of a novel algorithm for collective decision-making. Our framework is able to support the emergence of a consensus in favour of a randomly selected site in the absence of any quality bias. However, the model also incorporates the means to take into account such a bias when present and so we anticipate that the probability of a correct decision will be maximised when alternative options of different values are available. Because host quality can vary over time (for instance as a result of exogenous fluctuations in the availability of resources such as memory, bandwidth etc.), taking this aspect into account bears the possibility of studying the influence of environmental dynamics on the decision process. At this stage, we can only speculate that it will depend essentially on relative time-scales (the static environment used here can be regarded as an approximation of the case in which consensus can be reached over a short period compared to that which characterises fluctuations in host quality). These questions will be the subject of future work.

Another promising area of research is to draw further inspiration from a variety of biological models, especially with respect to the loss of consensus, i.e. the collective change of mind whereby the colony ceases to favour one particular site and relocates

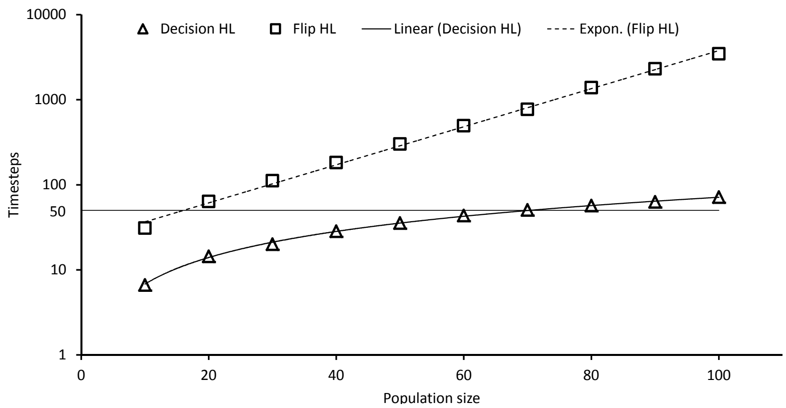

Fig. 8. Decision delay and consensus longevity as a function of scout population ($k = 100$, $P_{exp} = 0.1$, $\gamma = 2.0$, $l_{max} = 2.0$). See text for details.

to a new one (what we refer to as a flip) or scatters in the environment (which cannot happen in the current version of the algorithm but could present some advantages). Also, insect societies offer interesting alternative behavioural models for scouts. For instance, the positive feedback underpinning nest selection in *Temnothorax* seems to be modulated by the frequency of encounters with nest-mates *at* the prospective site [Pratt et al. 2005], which represents an intriguing (if not necessarily practical or scalable in the presence of many such sites) alternative to its explicit identification.

One of the key differentiators of our framework is that, unlike other swarm-based algorithms, it does not use stigmergetic communication, since no information is deposited in the environment. Instead, we use internal memory in the form of a confidence index to facilitate convergence towards a consensus. Although we do not wish to imply in any way that this approach is intrinsically better than stigmergy, we believe that it may offer certain advantages in some specific circumstances. For instance, there could be security concerns in the case where software agents involved in the collective decision do not have an established trust relationship with the hosting infrastructure. Or there could be several colonies simultaneously scouting the environment, in which case it may be desirable to avoid leaving behind information that could be exploited by a competitor.

We have demonstrated that well-known analytical and numerical techniques can be used to statistically predict the outcome of the collective decision process, i.e. time to convergence, population distribution at steady state, longevity of a solution etc. We argue that the ability to infer the likely behaviour of the system as a whole from the value of the parameters governing that of its constituents is critical to the usability of any multi-agent framework. Indeed, this makes it possible to determine, at design time, which combination of such values will yield the best results in any particular deployment scenario. Similarly, this predictability could allow the designer to rule out collective decision-making as a suitable mechanism if the quality of the solution falls below the requirements of a particular use-case.

The unavoidable presence of noise, e.g. in the form of residual exploration after a consensus has been reached, will always be regarded by some as a fundamental flaw of self-organising systems when it comes to real-world applications. However, in many cases, this problem can be alleviated by some ad-hoc techniques. For instance, in a service migration scenario, the actual relocation of a software component to a new host could be made to obey a secondary filtering mechanism designed to prevent the premature implementation of what could be a short-lived consensus among scouting agents. Note that some of the cited references [Parker and Zhang 2009; Peysakhov et al. 2006; Pratt et al. 2005] do address more specifically the problem of the termination of a collective decision, which lies beyond the scope of our paper.

More fundamentally, we believe that this zero-tolerance attitude will eventually become obsolete as the advantages of self-organised design principles, in terms of plasticity, simplicity, resilience and reduced management overhead, increasingly outweigh the cost of partial randomness.

REFERENCES

Ame, J., Halloy, J., Rivault, C., Detrain, C., and Denebourg, J. 2006. Collegial decision making based on social amplification leads to optimal group formation. In *Proceedings of the National Academy of Sciences of the USA*. Vol. 103. 5835–5840.

Ardagna, D., Trubian, M., and Zhang, L. 2007. Sla based resource allocation policies in autonomic environments. *Journal of Parallel and Distributed Computing 67*, 259–270.

Beckers, R., Deneubourg, J., Goss, S., and Pasteels, J. 1990. Collective decision making through food recruitment. *Insectes Sociaux 37*, 258–267.

Bergerud, A. 2000. *Ecology and Management of Large Mammals in North America*. Prentice Hall, New Jersey, Chapter Caribou.

Bonabeau, E., Theraulaz, G., Denebourg, J., Aron, S., and Camazine, S. 1997. Self-organization in social insects. *Trends in Ecology and Evolution 12*, 188–193.

Chow, K. and Kwok, Y. 2002. On load balancing for distributed multi-agent computing. *IEEE Transactions on Parallel and Distributed Systems 13*, 787–801.

Conradt, L. and Roper, T. 2005. Consensus decision making in animals. *Trends in Ecology and Evolution 20*, 449–456.

Couzin, I., Krause, J., Franks, N., and Levin, S. 2005. Effective leadership and decision-making in animal groups on the move. *Nature 433*, 513–516.

Csorba, M., Meling, H., Heegaard, P., and P., H. 2009. *Foraging for Better Deployment of Replicated Service Components (Lecture Notes in Computer Science)*. Vol. 5523. Springer Berlin, Heidelberg.

Dorigo, M. and Stuetzle, T. 2004. *Ant Colony Optimization.* MIT Press, Cambridge, MA.

Dyer, F. and T.D., S. 1994. Colony migration in the tropical honey beeapis dorsata f. (hymenoptera: Apidae). *Insectes Sociaux 41*, 129–140.

Fischer, M. J., Lynch, N. A., and Paterson, M. S. 1985. Impossibility of distributed consensus with one faulty process. *J. ACM 32,* 2, 374–382.

Fu, C. and Xu, C. 2005. Service migration in distributed virtual machines for adaptive grid computing. In *The International Conference on Parallel Processing.* IEEE Computer Society, 358–365.

Gueron, S. and Levin, S. 1993. Self-organization of front patterns in large wildebeest herds. *Journal of Theoretical Biology 165*, 541–552.

GUPTA, S. AND SRIMANI, P. 2003. Adaptive core selection and migration method for multicast routing in mobile ad hoc networks. *IEEE Transactions on Parallel and Distributed Systems 14*, 27–38.

HANDL, J. AND MEYER, B. 2002. Improved ant-based clustering and sorting in a document retrieval interface. In *Proceedings of Parallel Problem Solving from Nature*. Vol. 7.

HEIMFARTH, T. AND JANACIK, P. 2006. Ant based heuristic for os service distribution on ad-hoc networks. *Biologically Inspired Cooperative Computing*, 75–84.

HEUSSE, M., GUERIN, S., SNYERS, D., AND KUNTZ, P. 1998. Adaptive agent-driven routing and load balancing in communication networks. *Advances in Complex Systems 1*, 234–257.

HOLDO, R., HOLT, R., AND FRYXELL, J. 2009. Opposing rainfall and plant nutritional gradients best explain the wildebeest migration in the serengeti. *American Naturalist 173*, 431–445.

KENNEDY, J., EBERHART, R., AND SHI, Y. 2001. *Swarm intelligence*. Morgan Kaufmann, San Francisco.

KOSHI, K., HILTUNEN, M., AND JUNG, G. 2009. Performance aware regeneration in virtualized multitier applications. In *DSN09 Workshop on Proactive Failure Avoidance, Recovery and Maintenance (PFARM)*. IEEE Computer Society.

LINDAUER, M. 1951. Bienentnze in der schwarmtraube. *Naturwissenschaften 38*, 509–513.

LINDAUER, M. 1953. Bienentnze in der schwarmtraube ii. *Naturwissenschaften 40*, 379–385.

LINDAUER, M. 1955. Schwarmbienen auf wohnungssuche. *Zeitschrift Fuer Vergleichende Physiologie 37*, 263–324.

MALLON, E., PRATT, S., AND FRANKS, N. 2001. Individual and collective decision-making during nest selection by the ant leptothorax albipennis. *Behavioral Ecology and Sociobiology 50*, 352–359.

MESSIG, M. AND GOSCINSKI, A. 2007. Autonomic system management in mobile grid environments. In *Fifth Australasian Symposium on Grid Computing and e-Research*. Vol. 68. 49–58.

MONTRESOR, A., MELING, H., AND BABAOGLU, O. 2002. Messor: Load-balancing through a swarm of autonomous agents. In *Agents and Peer-to-Peer Computing*. 125–137.

MUSUNOORI, S. B. AND HORN, G. 2007. Application service placement in stochastic grid environments using learning and ant-based methods. *Multiagent and Grid Systems 3,* 1, 19–41.

PARKER, C. A. AND ZHANG, H. 2009. Cooperative decision-making in decentralized multiple-robot systems: the best-of-n problem. *IEEE/ASME Transactions on Mechatronics 14,* 2, 240–251.

PEASE, M., SHOSTAK, R., AND LAMPORT, L. 1980. Reaching agreement in the presence of faults. *J. ACM 27,* 2, 228–234.

PEYSAKHOV, M., DUGAN, C., JODI, P. J., AND REGLI, W. 2006. Quorum sensing on mobile ad-hoc networks. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*. 1104–1106.

PRATT, S. C., SUMPTER, D. J., MALLON, E. B., AND FRANKS, N. R. 2005. An agent-based model of collective nest choice by the ant temnothorax albipennis. *Animal Behaviour 70,* 5, 1023–1036.

SAFFRE, F., MAILLEUX, A., AND DENEUBOURG, J. L. 2000. Exploratory recruitment plasticity in a social spider (anelosimus eximius). *Journal of Theoretical Biology 205*, 37–46.

SAFFRE, F., TATESON, R., HALLOY, J., SHACKLETON, M., AND DENEUBOURG, J. L. 2009. Aggregation dynamics in overlay networks and their implications for self-organised distributed applications. *The Computer Journal 52,* 4, 397–412.

SCHAERF, A., SHOHAM, Y., AND TENNENHOLTZ, M. 1995. Adaptive load balancing: A study in multi-agent learning. *Journal of Artificial Intelligence Research 2*, 475–500.

SEELEY, T. 2003. Consensus building during nest-site selection in honey bee swarms: the expiration of dissent. *Behavioural Ecology and Sociobiology 53*, 417–424.

SEELEY, T. AND BUHRMAN, C. 1999. Group decision making in swarms of honey bees. *Behavioral Ecology and Sociobiology 45*, 19–31.

SEELEY, T., CAMAZINE, S., AND SNEYD, J. 1991. Collective decision-making in honey bees: how colonies choose among nectar sources. *Behavioral Ecology and Sociobiology 28*, 227–290.

Seeley, T. and Morse, R. 1978. Nest site selection by the honey bee apis mellifera. *Insectes Sociaux 25*, 323–337.

Seeley, T. and Visscher, P. 2004. Group decision making in nest-site selection by honey bees. *Apidologie 35*, 101–116.

Shen, K., Tang, H., and Yang, T. 2002. A flexible qos framework for cluster-based network services.

Sumpter, D. and Pratt, S. 2009. Quorum responses and consensus decision making. *Philosophical Transactions of The Royal Society B 364*, 743–753.

Wolf, J. and Yu, P. 2001. On balancing the load in a clustered web farm. *ACM Transactions on Internet Technology 1*, 231–261.

8. ACKNOWLEDGMENTS

The authors wish to thank Jean Louis Deneubourg and Jose Halloy from the Universite libre de Bruxelles for their useful comments and continued feedback during the course of this work. At the time when this research was conducted, both authors were employed by BT Innovate and Design (Ipswich, UK), where it was supported by Mark Shackleton and Mike Fisher under the strategic research programme on business-class open services.