2014 47th Hawaii International Conference on System Science

Validation and Analysis of a Distributed, Agent-Based Metaheuristic for Negotiation of Consensus Inspired by Honeybee Nest Site Selection Behavior

Alexander S. Mentis
Auburn University
asm0015@auburn.edu

Levent Yilmaz
Auburn University
yilmale@auburn.edu

Abstract

In this paper, we validate and analyze the ability of an agent-based metaheuristic to facilitate the negotiation of consensus among distributed, networked agents. Our metaheuristic is based on the process honeybees use to achieve consensus in selecting a new nest site. We show that our metaheuristic successfully guides systems to a consensus in a high percentage of cases and that the quorum size parameter controls the trade-off between optimality of choice versus time to consensus and failure rate. Despite agents having communication with only their local neighbors and the absence of centralized data aggregation, coordination, or mediation, our metaheuristic frequently results in the same consensus as that which would be returned by one or more well-known voting algorithms that require global knowledge and centralized tallying to generate a solution; when it does not, the returned consensus is usually an alternative similar in quality to the globally-informed option.

1. Introduction

Social communities and multi-agent systems sometimes face the problem of having to reach a consensus on a single choice selected from among several discrete options. While individual members of these groups might differ in their preferred ranking of the available choices and/or the perceived quality of each choice, it may nevertheless be in the best interest of the entire community that it reach a consensus on only one of the available choices. Resolving the conflict between individual disagreements and the need for collective agreement on the final outcome depends on both social choice theory, which is concerned with the aggregation of individual preferences to derive a group preference, and negotiation protocols, which are concerned with processes that individuals with conflicting preferences can use to effect compromise.

Most current voting and negotiation protocols use a centralized mediator and/or operate under the assumption that each individual can communicate directly with every other. In practice, however, it is not always possible to have full interconnection of individuals, and centralized mediation may cause an undesirable bottleneck or single point of failure affecting scalability and robustness. Essentially, we would like a way to determine the choice that a collective would prefer if each member voted for its favored choice, without actually having to hold an election, which requires all votes to be consolidated for tallying. Instead, we desire that the collective reach an agreement by interacting with and persuading only their local neighbors.

Potential applications for this type of algorithm include, among other possibilities, traffic control software in which road segments negotiate a desired speed limit based on perceived traffic flow or devices connected to a smart electric grid that must agree upon production and consumption rates or schedules. As is typical of metaheuristic techniques; however, the iterative nature of this algorithm can lead to lengthy execution times. One of the findings we present in this paper shows how the quorum size parameter can be adjusted to manipulate the algorithm's average speed to consensus; nevertheless, this metaheuristic is likely to be more suitable for offline negotiation, perhaps in conjunction with machine learning techniques.

In this paper, we present the experimental performance results of our decentralized, unmediated, agent-based, metaheuristic process for consensus negotiation that approximates the negotiation protocol used by honeybee scouts to reach consensus on a new nest site. Results from experiments indicate that, more than 94% of the time, the metaheuristic yields collective agreement on a single choice that matches the outcome that would have resulted from an election using plurality, range, and/or Borda voting. Unlike these voting systems, however, our metaheuristic is able to produce a consensus without the need for a centralized mechanism for vote tallying. Additionally, agents are able to reach compromises resulting in global consensus despite negotiating with only their local neighbors rather than requiring full interconnectedness and

978-1-4799-2504-9/14 $31.00 © 2014 IEEE
DOI 10.1109/HICSS.2014.42

individual negotiation with all other agents. In the rare instances that our metaheuristic does not reach a consensus predicted by one of the three voting systems above, the decision reached is usually very similar in quality to the outcome predicted by the voting systems.

The remainder of the paper is structured as follows. In Section 2 we provide background on the consensus negotiation problem, describe the bee behavior upon which our metaheuristic is based, and give an overview our metaheuristic technique. In Section 3 we discuss the metaheuristic parameters and methodology we used in our experiments. In Section 4 we analyze and discuss the results. In Section 5 we present our conclusion and suggest future work.

2. Background

2.1. Related Work

A significant amount of the research conducted in the development of automated negotiation has resulted in protocols that require the use of a centralized mediator that passes proposals and counter-proposals between negotiating agents. Examples of this trend can be found in [4, 7, 8, 10, 11, 13, 26, 28], all of which rely upon some sort of centralized mediation. In addition to being dependent on a centralized mediator, none of these protocols place restrictions on which pairs of agents are allowed to negotiate through the mediator; they assume that all agents are fully-connected, or at least connected to a mediator, so they can potentially negotiate with any other agent. Agents situated in the physical world, however, may not operate in conditions compatible with these assumptions; there may be physical barriers to direct communication with other agents in the system, and for the sake of robustness, it is often desirable to avoid potential single points of failure, such as a centralized mediator.

A number of natural systems have evolved to handle the problem of arriving at a consensus without full interconnection or dedicated mediators. Examples include groups of animals choosing a collective direction of travel [18] and insect colony nest site selection [14, 20, 27]. In studying these systems, it has been shown that groups can arrive at decentralized consensus decisions using simple rules, even when individuals participating in the decision making are ignorant of global knowledge, such as the current majority preference or the quality of their own information [2]. In this paper, we focus specifically on the techniques used by honeybee swarms to form consensus on selection of a new nest site location.

Much like the combination of cooperation and competition evident in most agent-based negotiations [28], the negotiation process by which honeybees form a consensus on a new nest site relies upon a balance of independent, competitive proposals and interdependent, cooperative persuasion [6, 12]. Unlike many other negotiation scenarios, however, failure of a honeybee swarm to come to a sufficiently rapid consensus on a new nest site location can result in the death of the swarm. If a new nest site is not located quickly enough, the swarm may die of exposure, whereas reaching a split decision can result in the splitting of the swarm and loss of the queen, which is similarly fatal to the colony [16].

Given such dire consequences of failure, the honeybee decision making process, and therefore our proposed metaheuristic, prioritizes compromise over individual stubbornness. This prioritization is not without precedent in negotiation protocols where it has been observed that early and frequent concessions are key to successful negotiation of complex contracts [10]. Like in [10], we treat consensus formation as an adaptive optimization problem where the objective is to iteratively improve the social well-being of the collective through an iterative, stochastic process; however, whereas the negotiation protocol in [10] uses simulated annealing as its foundation, we borrow features from particle swarm optimization [9], particularly the ideas of local neighbor interactions and the competition between social and cognitive influences, as the foundation for our metaheuristic approach.

2.2. Honeybee Decision Making Overview

To understand the rationale for our metaheuristic negotiation protocol, one must first have a basic understanding of honeybee nest site selection behavior. We provide enough detail here to understand the mechanics of our negotiation protocol. For a more comprehensive description of the behavior, see [12, 16, 17, 20-22, 25, 27], from which the following description is derived.

When a honeybee swarm must find a new nest site, the swarm leaves the old nest and forms a cluster nearby. A relatively small sub-population of the swarm, the scout bees, then depart the swarm in search of candidate nest sites. Scouts are either successful in finding a potential site and return to the swarm with a site preference, or else they are unsuccessful and return to the swarm as uncommitted scouts.

A scout that finds a potential site performs a "waggle dance" at the swarm to tell its sister scouts where the candidate site is located, and the enthusiasm of her dance is proportional to the scout's perceived quality of the candidate site. Between dances for a site, the scout will make trips between the swarm and candidate site. How long the scout will continue to dance for and visit the candidate site is also proportional to the scout's perceived quality of the site. A scout will spend more time dancing for a high-quality site than for one of mediocre quality.

An uncommitted scout may choose to follow the dance of a sister scout that found a candidate site or return to exploration. This choice is based probabilistically on the number of dancing scouts it encounters at the swarm. This makes sense intuitively; if a large number of scouts are dancing for the same site, it is more likely that an uncommitted scout will encounter one of these dances and be persuaded by it. On the other hand, if very few scouts are dancing for sites at the swarm, the uncommitted scout is much more likely to give up on finding a dance to follow and go exploring on its own.

If an uncommitted scout chooses to follow a dance, she performs her own independent assessment of the candidate site, and, if it meets her minimum criteria, she will return to the swarm and dance for the site as described above. As more scouts are recruited to a particular site, the number of scouts visiting the site between dances increases. Once a threshold of visitors to a candidate site is reached, a quorum for that location is formed and bees returning to the swarm from that site augment their dances with a "piping" signal. When the piping signal reaches a threshold value, a consensus is considered to exist, the swarm lifts off, and the scouts guide the swarm to the agreed upon site.

2.3. Swarm Decision-Making Metaheuristic Overview

In order to provide a foundation for understanding our methodology and results, this section provides a brief overview of the key elements in our metaheuristic approach to consensus negotiation and maps algorithmic steps to the honeybee decision making process. For a more detailed description of the algorithm, see [15].

In adapting the honeybee behavior to a generic negotiation protocol, we make some adjustments. First, we do not require agents to go exploring for potential options at the outset. Instead, an agent begins deliberations with a set of utility values for the set of potential outcomes, and agents begin deliberations committed to the option that produces the highest perceived utility for themselves. Also, whereas a honeybee will only dance for sites that meet a minimum level of quality and otherwise remain uncommitted, our agents maintain a utility value for all options, even those for which the agent's perceived utility is very low. This ensures that no agents stubbornly refuse any option and eventually become committed to one of the options, as described next.

It has been shown in [3] that uninformed individuals are an important part of moderating the influence of a strongly opinionated minority in forming democratic consensus. While none of our agents begin deliberations uninformed, we do have a mechanism whereby agents become progressively less attached to their current preference, just like honeybee enthusiasm for a site decays over time. The length of time that an agent will stay committed to a particular choice is a function of its perceived utility for that choice. An evaporation rate is controlled by a model parameter, and when an agent's commitment dips below a model-specified value the agent becomes neutral and open to committing to new options based on the aggregate of information provided by its neighbors. This behavior is required to achieve the "expiration of dissent" described in [19] and allow agents to make the concessions required for successful negotiation [10].

For the results presented in this paper, the number of simulation ticks an agent will remain committed to a choice is equal to the integer value of the preference weight for that choice. For example, if an agent's preference distribution over three choices is $\langle 0.40, 0.30, 0.30 \rangle$, this would indicate that the agent prefers choice 1 over both choices 2 and 3 by 10% but it is indifferent with respect to choice 2 versus choice 3, and once committed to choice 1, it remains so committed for 40 simulation ticks. In these experiments, we used a linear enthusiasm decay function, with the commitment level decaying by one point for every simulation tick. We chose to use a linear decay function because it matches the observed enthusiasm decay of honeybees in nest site selection. [20] Thus, in this example, by the 40th simulation tick after making its choice, the agent's commitment level reaches 0 and the agent becomes uncommitted and open to selecting a new choice as described next.

As previously mentioned, agents have a set of neighbors, represented by a graph, with which they can communicate. The information provided to an agent by a neighbor consists of a set containing the other agents in the collective that agree with the neighbor on a particular choice (which it either knows about directly or has learned about from its own neighbors) along with each of that set's agents' remaining commitment durations for that choice. It should also be noted that, as this information is passed from agent to agent, the remaining commitment level for each agent in the set is decayed in accordance with the previously-described decay function. Thus the commitment level associated with an agent's membership in the set stays consistent with the agent's actual commitment level at each simulation tick.

At each iteration of the deliberation process, an agent that is currently committed to a particular choice only aggregates the information of neighbors that it agrees with. This is equivalent to how honeybees visiting a candidate nest site will only see other bees that are visiting the same site. On the other hand, if an agent's commitment to a choice has expired, it must determine a new choice to which to commit. After aggregating its neighbors' information, if there is only one option that has achieved a number of members sufficient to meet the quorum threshold, the agent automatically commits to it. This is essentially a form of preferential attachment and allows an option with a sufficient lead to assert its dominance; however, if there is no quorum, or if there is more than one quorum, the agent uses roulette wheel selection to probabilistically make a new commitment based on the percentage of its neighbors committed to each option.

In the case of the honeybee, the end of deliberations is signaled by an audible piping signal. To replicate this in an artificial setting, the model would require a broadcast channel. Since we are assuming that agents cannot necessarily communicate with every other agent and are restricted to communicating with only their neighbors, we cannot model any such channel. Instead, as is common in adaptive optimization algorithms, we simply run our algorithm for a specified number of iterations. As we will see in the results section, this number does not need to be too large to achieve good consensus results.

3. Methodology

We conducted a total of 10,000 trials to determine how successful the metaheuristic was at reaching an acceptable consensus among 200 agents on one of five possible choices. The 10,000 trials were split into four groups of 2,500 trials, each, using a quorum threshold size of 5, 15, 25, and 35, respectively. Our choices of values for agent count and quorum sizes come from numbers typical in empirical observations of honeybee swarms and the values used to model honeybee nest site selection in [17]. The quorum sizes we used bracket the quorum size of 20 typically used by honeybees [17, 21].

For each of the four quorum threshold levels, we ran 50-trial batches on each one of 50 different randomly-generated agent preference distributions. While an agent preference distribution was the same within a single batch of 50 trials, each trial within the batch used a differently-seeded pseudorandom number generator to shuffle the agent neighbor interconnections and drive the stochastic agent activation order and choices made by the agents at each time tick.

The random generation of agent preferences was performed in two stages. First, a target probability distribution (p_{target}) for the five choices was generated from the uniform distribution $\mathcal{U}(0,1)$ using the method described in [24]. This determined what probability each choice had of being the preferred choice of an agent. For example, a p_{target} of $\langle 0.16, 0.04, 0.09, 0.68, 0.03 \rangle$ would indicate that approximately 68% of the agents should assign Choice 4 their highest preference, whereas approximately 3% should assign Choice 5 their highest preference. We then generated a range of preference weights for each agent, again from the uniform distribution $\mathcal{U}(0,1)$ using the method described in [24], and we used a roulette wheel selection based on p_{target} to assign the generated preferences to each of the five choices. This methodology allowed us to randomly generate preferences across all agents that resulted in collective preference sets that could generate winners, losers, and ties, depending on the social rules by which they were evaluated. Had we not generated the p_{target} distribution for preference assignment, the uniformly random agent preferences would have been distributed so uniformly that each of the five options would have received about 20% of the total utility.

Agents were connected in an augmented ring topology where they were connected to their twenty nearest neighbors both forward and backward along the circumference of the ring. More formally, this is the 20-regular circulant graph $Ci_{200}(1,\ldots,10)$. In order to judge the impact of the location of particular agents on the consensus-forming process, we shuffled agent neighbors on each trial.

The acceptability of an outcome can be interpreted in different ways. We follow [5] in choosing to evaluate our outcomes with respect to the resultant social welfare of our artificial multi-agent society. A number of possible social welfare metrics are presented in [1, 23, 29]; we use three in the presentation of our results: plurality voting, Borda voting, and range voting (which is similar to utilitarian social welfare, but uses the average utility rather than the sum of the utilities). At this point, it should be noted that agents were allowed to rank options equally if they both had the same utility. In plurality voting, this resulted in the total tally of first-place votes across all preferences sometimes exceeding the number of agents, and this allowance departs from what is traditionally allowed in Borda and range voting.

4. Results and Discussion

For our trials, we consider an appropriate consensus result to be one that agrees with at least one of the three calculated social welfare metrics mentioned above. In practice, the majority of our random preference distributions (61%) *always* result in unanimous decisions where the all three social welfare metrics return the same choice as socially "best," and generally speaking, about 85% of our trials result in a unanimous decision. In these cases, matching at least one outcome is the same as matching all of them and is therefore a highly appropriate consensus by our metrics. We consider a failure to be any case that does not result in a consensus within 3,000 simulation ticks, and we consider a sub-optimal result any case that results in a consensus that is not one of the results returned by any one of the social welfare metrics.

Figure 1 shows the 95% confidence intervals for our trial results. On average, for our four chosen quorum sizes, we find that our metaheuristic successfully converges to a socially-appropriate consensus over 94% of the time, regardless of the locations of agents in the social network, and it fails to result in any consensus at all less than 1% of the time. Figure 2 shows the average number of simulation ticks to reach consensus with 95% confidence intervals.

From these figures we can see that the success rate of reaching a consensus increases as the quorum size increases, but the result is less likely to be an optimal one. Varying the size of the quorum allows us to balance the trade-off between speed to consensus and quality of the result. These results may help explain why honeybees have evolved to detect quorums at a threshold size of about 20, as observed in [17, 21].

A question is why raising the quorum threshold, which would be intuitively expected to make it harder to reach a consensus, actually lowers the failure rate. We suspect that increasing this value (up to a point) actually has the effect of allowing clear majorities to trigger the end of deliberations sooner. When the quorum size is small, quorums for several alternatives form easily, creating a situation where several quorums must compete for selection. Further work is required to study the dynamics underlying this behavior.

It remains to evaluate the "badness" of the sub-optimal consensus results. Should our technique be faulted for the cases when it comes to a consensus that is not in the set of acceptable options returned by our social welfare functions? We have not developed quantifiable metrics for making this evaluation, but Table 1 presents data on all of the probability distributions that resulted in sub-optimal consensus 20% or more of the time for one or more quorum sizes.

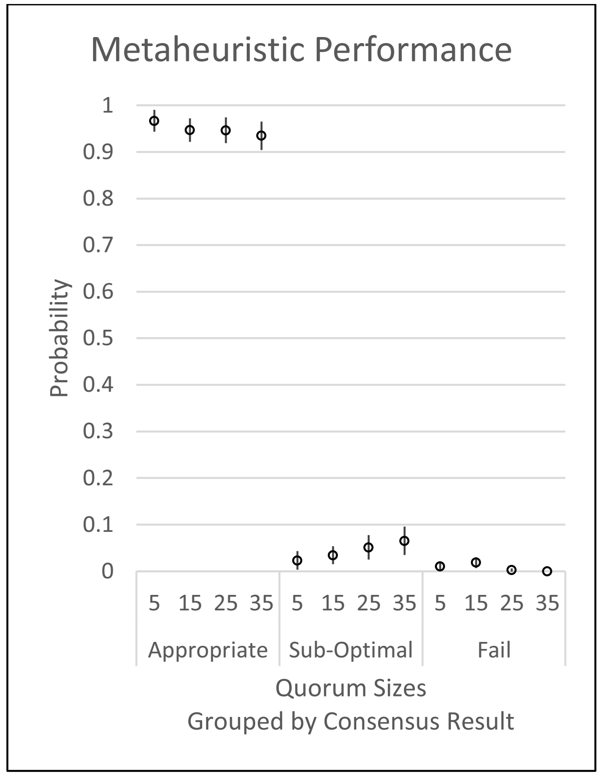

Figure 1. Probability of consensus results for different quorum sizes, shown with 95% confidence intervals.

Table 1. Average choice utility values across all agents for probability distributions resulting in >20% sub-optimal selection.

Probability Distribution Seed	Average Utility					5		15		25		35	
	Choice					Sub-Optimal	Fail	Sub-Optimal	Fail	Sub-Optimal	Fail	Sub-Optimal	Fail
	1	2	3	4	5								
4	0.21	0.17	0.23	0.15	0.24	0.06	0.00	0.30	0.00	0.26	0.00	0.34	0.00
6	0.35	0.10	0.33	0.09	0.13	0.06	0.10	0.20	0.02	0.28	0.08	0.30	0.00
20	0.30	0.16	0.05	0.27	0.21	0.02	0.00	0.06	0.00	0.08	0.00	0.22	0.00
25	0.25	0.21	0.10	0.25	0.19	0.42	0.00	0.24	0.08	0.20	0.02	0.30	0.00
26	0.15	0.27	0.10	0.24	0.22	0.00	0.00	0.10	0.02	0.10	0.00	0.20	0.00
32	0.21	0.21	0.23	0.09	0.25	0.00	0.00	0.14	0.00	0.22	0.00	0.36	0.00
35	0.25	0.25	0.14	0.18	0.17	0.12	0.00	0.16	0.12	0.38	0.00	0.38	0.00
44	0.24	0.27	0.06	0.15	0.28	0.20	0.08	0.18	0.06	0.32	0.00	0.16	0.00
48	0.17	0.12	0.25	0.25	0.21	0.22	0.04	0.10	0.10	0.20	0.00	0.16	0.00

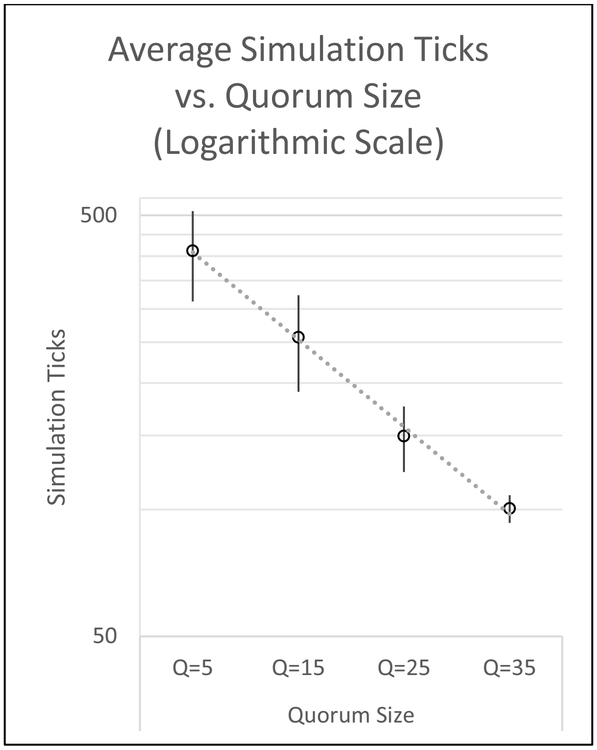

Figure 2. Average simulation time ticks to reach consensus for varying quorum sizes, shown with 95% confidence intervals on logarithmic scale.

As we can see, all of these cases have preference distributions where the average utilities for two or more choices are very close. In fact, it is these similarly-valued choices that are being selected in the cases when a "sub-optimal" consensus is reached, so, in fact, this result might be better named "nearly-optimal." This is not a surprising result when we consider that our metaheuristic is based on the behavior of individual bees that, for reasons of survival, value reaching a unanimous decision quickly over possibly never reaching a perfect decision, especially when the compromise in quality required to reach a unanimous decision is very small. It also suggests that our metaheuristic generally operates according the social rules prescribed by utilitarian social welfare.

A final point of analysis is the speed with which our metaheuristic is able to achieve its results. Figure 2 depicts how the number of simulation ticks required to reach consensus, which corresponds to negotiation rounds, relates to quorum size. It is interesting to note that the average number of negotiation rounds required to reach consensus decreases logarithmically with respect to increased quorum size. Some of the difference in performance between quorum sizes can be attributed to the higher failure rate of the smaller quorum thresholds where there are more batches with a maximum tick count of 3,000. Nevertheless, even when the batches with failures are not considered, the higher quorum threshold is still faster at generating consensus than the lower quorum threshold. We suspect that this result is due to the same underlying mechanism as that which results in lower failure rates for the higher threshold value.

5. Conclusions and Future Work

In this paper, we have shown that a metaheuristic negotiation protocol based on the simple rules honeybees use to negotiate consensus on a new nest site produces highly-successful results for negotiation of single-attribute, distributed consensus when considered from a social welfare perspective. Unlike many other utility-based voting and negotiation protocols, our technique requires neither dedicated mediators nor global knowledge of social preferences to produce results similar to those that would result from global preference aggregation techniques. Our data indicate that our protocol primarily adheres to the social rule set of utilitarian social welfare. We have also shown that quorum-sensing is an effective technique for balancing the influences of conflict and compromise in negotiation protocols and that the correct tuning of a quorum size threshold parameter can yield faster and higher success rates in forming consensus at the cost of decreased optimality, but that the selected alternative is still close to the optimal choice in average utility. Positions of individuals in the social network do not appear to have a significant impact on performance in our augmented ring topology.

This negotiation protocol provides many avenues for interesting future work in terms of variation and scalability. It may very well be the case that our chosen network topology is an idealized case that skews our results toward high success rates. It also may not be the most likely topology for real-world networks. Therefore, we have plans to perform further evaluation with respect to different network topologies, such as hierarchical and small-world networks or dynamic topologies. We expect this algorithm to handle dynamic topologies especially well, since this would replicate the dynamic neighborhoods honeybees encounter on the swarm as they travel back and forth between the swarm and nest sites.

We also plan to evaluate this protocol's scalability in both number of negotiating agents and number of choices. We have already conducted a limited number of experiments which suggest that the algorithm reaches consensus in a similar number of simulation ticks as the results presented here, even when the number of agents is increased by an order of magnitude. Another interesting scalability aspect to explore would be to determine if we can scale this technique to handle multi-parameter negotiation.

In addition, we would like to evaluate the robustness of our protocol, such as how our protocol handles dynamic group membership and its resilience to manipulability by selfish agents. We expect that the quorum threshold parameter and the periodic generation of neutral agents will be effective in mitigating selfish behavior. Finally, we believe we can we improve our protocol's ability to overcome deadlocks that lead to failure to reach consensus by incorporating simple learning behavior in the agents.

6. References

[1] Chevaleyre, Y., Dunne, P., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-Aguilar, J., and Sousa, P., "Issues in Multiagent Resource Allocation", *Informatica*, 30, 2006.

[2] Couzin, I.D., Krause, J., Franks, N.R., and Levin, S.A., "Effective Leadership and Decision-Making in Animal Groups on the Move", *Nature*, 433 (7025), 2005, pp. 513-516.

[3] Couzin, I.D., Ioannou, C.C., Demirel, G., Gross, T., Torney, C.J., Hartnett, A., Conradt, L., Levin, S.A., and Leonard, N.E., "Uninformed Individuals Promote Democratic Consensus in Animal Groups", *Science*, 334 (6062), 2011, pp. 1578-1580.

[4] Ehtamo, H., Verkama, M., and Hamalainen, R.P., "How to Select Fair Improving Directions in a Negotiation Model over Continuous Issues", *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 29 (1), 1999, pp. 26-33.

[5] Endriss, U., Maudet, N., Sadri, F., and Toni, F., "Negotiating Socially Optimal Allocations of Resources", *J. Artif. Int. Res.*, 25 (1), 2006, pp. 315-348.

[6] Galla, T., "Independence and Interdependence in the Nest-Site Choice by Honeybee Swarms: Agent-Based Models, Analytical Approaches and Pattern Formation", *Journal of Theoretical Biology* (262), 2010, pp. 186-196.

[7] Heiskanen, P., "Decentralized Method for Computing Pareto Solutions in Multiparty Negotiations", *European Journal of Operational Research*, 117, 1999, pp. 578-590.

[8] Heiskanen, P., Ehtamo, H., and Hämäläinen, R.P., "Constraint Proposal Method for Computing Pareto Solutions in Multi-Party Negotiations", *European Journal of Operational Research*, 133 (1), 2001, pp. 44-61.

[9] Kennedy, J., and Eberhart, R., "Particle Swarm Optimization", IEEE International Conference on Neural Networks 1995, 1995, pp. 1942-1948 vol.1944.

[10] Klein, M., Faratin, P., Sayama, H., and Bar-Yam, Y., "Protocols for Negotiating Complex Contracts", *Intelligent Systems, IEEE*, 18 (6), 2003, pp. 32-38.

[11] Li, M., Vo, Q.B., and Kowalczyk, R., "Searching for Fair Joint Gains in Agent-Based Negotiation", Autonomous Agents & Multiagent Systems/Agent Theories, Architectures, and Languages, 2009, pp. 1049-1056.

[12] List, C., Elsholtz, C., and Seeley, T.D., "Independence and Interdependence in Collective Decision Making: An Agent-Based Model of Nest-Site Choice by Honeybee Swarms", *Philosophical Transactions of the Royal Society Biological Sciences* (364), 2009, pp. 755-762.

[13] Lopez-Carmona, M.A., Marsa-Maestre, I., and Klein, M., "Consensus Policy Based Multi-Agent Negotiation", 4th International Workshop on Agent-based Complex Automated Negotiations, 2011.

[14] Mallon, E.B., Pratt, S.C., and Franks, N.R., "Individual and Collective Decision-Making During Nest Site Selection by the Ant *Leptothorax Albipennis*", *Behavioral Ecology and Scoiobiology*, 50, 2001, pp. 352-359.

[15] Mentis, A.S., and Yilmaz, L., "Adapting a Natural System Simulation Model to a General-Purpose Metaheuristic: Toward Engineering Emergent Distributed Decision-Making", Summer Simulation Multi-Conference 2013, 2013.

[16] Niven, J.E., "How Honeybees Break a Decision-Making Deadlock", *Science*, 335 (6064), 2012, pp. 43-44.

[17] Passino, K.M., and Seeley, T.D., "Modeling and Analysis of Nest-Site Selection by Honeybee Swarms: The Speed and Accuracy Trade-Off", *Behavioral Ecology and Sociobiology*, 59 (3), 2006, pp. 427-442.

[18] Reynolds, C.W., "Flocks, Herds and Schools: A Distributed Behavioral Model", *SIGGRAPH Comput. Graph.*, 21 (4), 1987, pp. 25-34.

[19] Seeley, T., "Consensus Building During Nest-Site Selection in Honey Bee Swarms: The Expiration of Dissent", *Behavioral Ecology and Sociobiology*, 53 (6), 2003, pp. 417-424.

[20] Seeley, T.D., and Buhrman, S.C., "Group Decision Making in Swarms of Honey Bees", *Behavioral Ecology and Sociobiology*, 45, 1999, pp. 19-31.

[21] Seeley, T.D., and Visscher, P.K., "Choosing a Home: How the Scouts in a Honey Bee Swarm Perceive the Completion of Their Group Decision Making", *Behavioral Ecology and Scoiobiology*, 54, 2003, pp. 511-520.

[22] Seeley, T.D., Visscher, P.K., Schlegel, T., Hogan, P.M., Franks, N.R., and Marshall, J.a.R., "Stop Signals Provide Cross Inhibition in Collective Decision-Making by Honeybee Swarms", *Science*, 335 (6064), 2012, pp. 108-111.

[23] Shoham, Y., and Leyton-Brown, K., *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge University Press, Cambridge, 2009.

[24] Smith, N.A., and Tromble, R.W., "Sampling Uniformly from the Unit Simplex", Johns Hopkins University, 2004.

[25] Stefan, J., Middendorf, M., and Beekman, M., "Honeybee Swarms: How Do Scouts Guide a Swarm of Uninformed Bees", *Animal Behaviour*, 70, 2005, pp. 349-358.

[26] Verkama, M., Ehtamo, H., and Hämäläinen, R., "Distributed Computation of Pareto Solutions in *N*-Player Games", *Mathematical Programming*, 74 (1), 1996, pp. 29-45.

[27] Visscher, P.K., and Camazine, S., "Collective Decisions and Cognition in Bees", *Nature*, 397 (6718), 1999, pp. 400-400.

[28] Vo, Q.B., Padgham, L., and Cavedon, L., "Negotiating Flexible Agreements by Combining Distributive and Integrative Negotiation", *Intelligent Decision Technologies*, 1 (1-2), 2007, pp. 33-47.

[29] Wooldridge, M., *An Introduction to Multiagent Systems*, Wiley, 2nd edn, United Kingdom, 2009.