

Economics of Cooperation: Social Foraging in Distributed Systems

Ashish Umre¹ and Ian Wakeman¹

¹ Software Systems Group, Department of Informatics,
University of Sussex, Falmer, Brighton BN1 9QH, United Kingdom
(ashishu, ianw)@sussex.ac.uk

Abstract. The sharing and collective processing of information by individuals in any social system is an attempt to reduce the uncertainty associated with key features of their environments by collecting and storing information. By sampling each of its options regularly, an individual gains from being able to exploit them when they are productive and avoid them otherwise. In this way, collection of information can be thought of as a solution to the uncertainty problem that maximises potential opportunities [3], [4]. Some group-living species have evolved effective social mechanisms for reducing uncertainties in their environments. However, doing so may entail certain costs with respect to attributes such as time, energy and attention. In this paper, we explore the cost/benefits of cooperation within the domain of distributed systems, where biologically inspired agents interact with each other using the environment to disseminate information about resources (foraging sites). In the sections that follow, we describe briefly the theory of cooperation, social foraging theory, the simulation model and some experiments to understand/analyse the dynamics of social foraging in stochastic environments.

1 Introduction: Social Foraging and Cooperation

To account for the manifest existence of cooperation and related group behaviour, such as Altruism and Restraint in competition, evolutionary theory has acquired two kinds of extension: Genetic kinship theory and reciprocity theory. If the agents are sufficiently closely related, altruism can benefit reproduction of the set, despite losses to the individual altruist. The evolution of the suicidal barbed sting of the honeybee worker could be taken as a paradigm for this line of theory [12].

Many of the benefits sought by living things are disproportionately available to cooperating populations. The problem lies with the fact that while an individual can benefit from mutual cooperation, each can also do so even better by exploiting the cooperative efforts of others. Over a period of time, the same individuals may interact again, allowing for more complex patterns of strategic interactions. [10] Argues that there are at least three ways that cooperation can evolve among unrelated individuals: reciprocity, group selection, and by-product mutualism. Though, kin selection is a fourth candidate.

As well as the existence of group, team and partitioned tasks in complex societies, another facet of higher-level functionality is a shift from individual to social/group foraging strategies. [18] identified six foraging strategies in ant colonies: (1) ‘*individual foraging*’ foraging without cooperation and communication with others; (2) ‘*tandem running*’ a scout guides one recruit to the food source with or without trail laying; (3) ‘*group mass recruitment*’ the scout guides a group of recruits to the source, usually laying a trail to the nest; (4) ‘*mass recruitment*’ the scout lays a trail while returning to the nest which guides recruits to the food source; (5) ‘*trunk trail*’ semi-permanent trails guide foragers to long-lasting food sources; and (6) ‘*group hunting*’ a group leaves the nest and forages collectively in a swarm along a well-defined trail system. These strategies also appear to be correlated with a decrease in the autonomy of the individual foragers themselves [19]. That is, there is a shift from information processing by individuals to emergent properties of a set of essentially probabilistically behaving individuals mediated through signals, i.e. a set of trail pheromones. For instance, in an individual foraging strategy the worker must rely on its own information, navigating back to the nest using the sun or other landmarks (e.g. the desert ant *Cataglyphis bicolor*).

In tandem running, a successful returning forager can recruit just one individual and passes on information of where the food source is by physically leading the recruit to the source (e.g. *Leptothorax*). However, with more complex strategies trail pheromones can pass the information not just to one other recruit but to many. There is no need for an individual to be able to navigate back to the nest using the sun or a prominent rock but can simply orient (‘smell’) their way along a chemical trail (e.g. *Atta*). Despite the apparent simplicity of this task, foragers experience a constant probability per unit distance of losing the trail. Seemingly counterintuitive, this apparently errant behaviour has been shown to be very adaptive at the group-level [20, 21]. Once lost, these workers become scouts who can search for new sites. However, it appears that the error rate is sufficiently tuned so that enough foragers do not lose the trail and thus can exploit the source whilst enough become scouts enabling a constant supply of new sources. (Parallel behaviour is known in honeybee foraging in which the directional information in waggle dances is imprecise) [22]. It seems that the complexity emerges at the level of the trail network (or group), which can adaptively adjust to fluctuating food dispersion or density. Thus, the foragers are a ‘group-level adaptive unit’ [5, 23], and also see [24].

2 Then again, how advantageous cooperation really is?

The acquisition and use of socially acquired information is commonly assumed to be profitable. But, there could be scenarios where the use of such information either provides no benefit or can actually incur a cost. It is suggested [2] that the level of incompatibility between the acquisition of personal and socially acquired information will directly affect the extent of profitability of the information, when these two sources of information cannot be acquired simultaneously, because of cognitive or

physical constraints. Also, a solitary individual's behavioural decisions will be based on cues revealed by its own interactions with the environment.

However, in many cases, for social animals the only socially acquired information available is the behavioural actions of others that expose their decisions, rather than the cues on which the decision was based. In such a situation it is thought that the use of socially acquired information can lead to *information cascades* that sometimes result in sub-optimal behaviour.

In our experiments, we look for results that suggest the presence of information cascades in the context of information sharing in distributed systems. Designing agents that rely both on individual foraging and shared information, or agents that just rely on shared information. Ongoing studies are focused on understanding whether this might happen in a highly dynamic environment; where there are constant changes in the flow of information about resources that undergo frequent updates.

2.1 Cost of cooperative efficacy

In any social group, individuals possess various behaviours that define the assortment of the interactions at all sorts of levels, individual, groups, cliques, teams etc. The social foraging theory suggests that, the functional consequence of an individual's foraging behaviour depends on both the individual's own actions and the behaviour of other foragers. There may be conflicts of interest between signallers and receivers. Where such a conflict exists, the receiver's need to acquire information may favour sensitivity to the cues provided by the behaviour and appearance of the signaller. In turn, this sensitivity may give rise to opportunities for manipulation and exploitation by the signaller.

It is understood that exploitative strategies are unlikely to persist in the long run, because they generate selection for a change in receiver responses. However, it is argued, that the evolution of exploitation may prove a recurrent, though, transient phenomenon. There are costs associated with broadcasting information publicly, as exemplified by the production of "*food vocalisations*" in many social animals. The issues that come under this context are, dangers of predation, and mass recruitment to a very less profitable resource may lead to starvation. This is equivalent to the "Slash Dot" effect that the Internet sometimes experiences.

Other costs within the context of a social system are cost of misinformation (lying), cost of accessing/using the resources and cost of signalling/cooperation. We use foraging games to analyse the economics of Kleptoparasitic¹ behaviour, to predict the ecological circumstances under which the behaviour is maintained. Other costs are expressed as survival rate; if an agent keeps failing/delaying to locate resources for

¹ *Kleptoparasitism* refers to all forms of exploitation of others' food discoveries or captures. It constitutes the information-sharing models in the Social Foraging Theory paradigm.

the requested processes/services it gets penalised and if this increases above a threshold, then the agent dies and a new agent replaces the old agent.

3 Model Overview

We implement a discrete-event simulation of cooperative (collaborative) agents, which share information (through the environment, Stigmergy²) about the location of resources. A process generator (P) generates processes/requests/tasks with Poisson distribution. Processes enter the system queue at the start of the simulation, where they wait to be allocated to N agents (which are initialised randomly). An agent gets allocated a process/task. Individual processes/tasks require a certain number of resources/services ($r_1, r_2 \dots r_n$) that it requires for the successful completion/execution of the process.

The resource generator (\mathcal{R}) generates a random number of resources for the successful execution/completion of a request. When an agent encounters some information about a resource/service, it probabilistically stores the information in its resource vector and/or publishes the information onto a “HotSpot”, if it decides to share it with others.

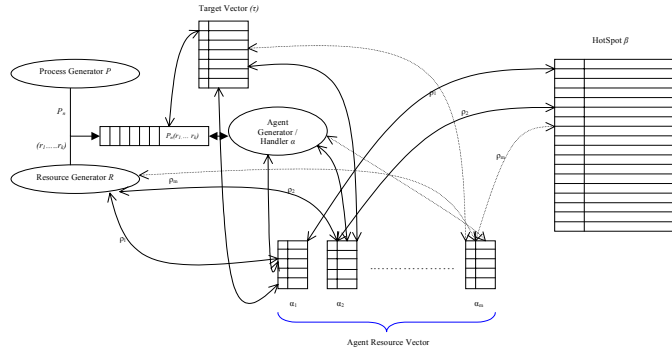


Fig. 1. Schematic representation of the information dissemination system.

If the agent encounters a resource it is searching for, it locks the resource, provided it's available at the time and marks the resource entry in the target vector (which contains the list of processes waiting to be finished and the status of the resources) under the specific process. Once all the required resources/services have been located, the process is executed. The agent can only lock the resource for a fixed time after

² The term “Stigmergy” was first introduced by *Pierre Paul Grassé*, a French entomologist, in 1959. He used the term to describe the coordination of activities of ants in carrying out complex activities, such as nest building or foraging, without direct communication amongst themselves. It is evident that stigmergy describes a form of asynchronous interaction and information interchange between entities mediated by an “active” environment.

which it will have to rejoin the queue. The agent incurs a cost once it has locked a resource. A resource diminishes by a certain value while the process/task is being executed. The j th process assigned to the i th agent is p_{ij} , and costs it C_{ij} . Individual resource cost is $C_{ij}^{r_n}$, for the resource r_n . Cumulative cost associated with j th process is,

$$C_{ij} = \sum_{n=1}^N C_{ij}^{r_n} \quad (1)$$

Agents can cooperate and form groups to collaboratively execute the process/task or choose to forage alone. The throughput of the system is calculated as a function of successfully completed jobs in the minimum time and with minimal costs. Agents probabilistically (ρ) cooperate with other agents, and decide to share information through the HotSpot or not. If the agent incurs a cost which is higher than the cost on its previous task $C_{ij+1} > C_{ij}$, it then either chooses to collaborate with other agent(s) by forming a group and/or change its degree of cooperation. This acts as a simple adaptive learning mechanism and some form of reciprocity. An agent's cooperative strategy (probability of publishing/sharing information) changes after every process or during successive simulation runs. This is more or less an equivalent NASH equilibrium³ for the agent.

We have considered the resource handling time as negligible and the process execution time as a random time factor. Other agents looking for the same resource can access the HotSpot and search through the advertised resources/services. The HotSpot contains the information about resources and their location. Each resource published at the HotSpot has a reinforcement value (similar to pheromone deposit) associated with it, which signifies the demand (Δ) of the resource.

Every time an agent accesses resource information at the HotSpot, it reinforces the pheromone deposit so that the resource path continues to exist, whereas if the reinforcement value goes below a certain value, it gets over written by the first new resource that appears in the system. Hence, the table is constantly updated with the latest information about resource paths. Agents attempt to optimise costs locally and globally in accordance with the dynamics of their interactions.

³ *Nash Equilibrium* is a combination of strategies for the players of a game, such that each player's strategy is a *best response* to the other players' strategies. A *best response* is a strategy, which maximises a player's expected payoff against a fixed combination of strategies played by the others.

3.1 Results and Analysis

We analyse some aspects of artificial and biological social systems, such as, optimal number of agents in the system [11], throughput of the system, degree of cooperation (which can depend on an implicit factor of relatedness). Demonstration of the use of Nash equilibrium, to show the “tragedy of the commons” for certain situations both in the simulations and in real life, e.g. Slash Dot effect. How a certain resource gets over exploited because of it being over publicised and may lead to its exhaustion/starvation. Similarities with Caraco’s food calling game [1], [15], when agents individually look for resources and on finding it, decide to publish it or not. According to Caraco’s model if they decide against publishing the information, then they are more susceptible to predation.

3.1.1 Optimal Group Size?

In general, we observe a peaked fitness function [6] when we analyse the system as a collection of agents trying to maximise the throughput and minimise the delay in acquiring information. The peaked function we see in Fig. 2 illustrates the existence of only one optimal agent population size for which, the throughput of the system is maximum, given that certain other parameters in the simulation remain fixed, like the number of resources.

This suggests that initially an increase in the agent population is beneficial in obtaining a good throughput, but the throughput peaks at some point for a certain size of population implying that there are enough agents to process requests for resources any further increase will result in delays due to queuing for resources. The Increasing Fitness plot is an indication of abundance of resources.




Fig. 2. Optimal Number of Agents.

3.1.2 Throughput of the System

The time taken to find all the resources for a request can vary depending on the number of resources required. Therefore, we calculated the average time (τ_{avg}) taken for finding the various resources over a series of runs and accumulated the data for all the possible number of resources in the system. We were interested in finding out the trend that follows in terms of time/hops taken to locate all those resources. As seen from Fig. 3(a) out that there is an increasing trend with respect to the number of hops. As the number of required resources increases it takes more time to find them, but the trend shows that there could be a decrease later on in the system as the agents develop an optimum response for each request, as the number of resources increase. This also may lead to a drop in the number of cooperators, meaning that individual foraging can sometimes also be a useful strategy Fig. 3(b). Fig. 3(c) shows the average cost incurred by agents over successive simulation runs. The drop in average cost suggests an increase in information sharing and level of cooperation.

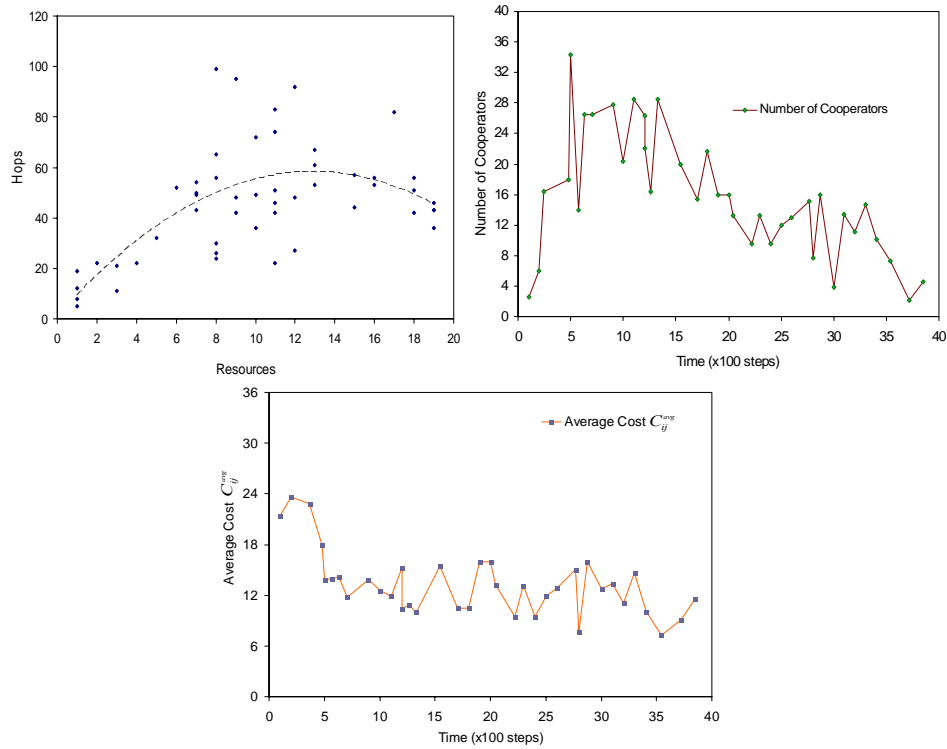


Fig. 3. (a) Time (τ_{avg}) to establish a resource path. (b) Number of Cooperators. (c) Average cost C_{ij}^{avg} over successive runs.

3.1.3 Slash dot effect/Kleptoparasitic Behaviour

Slash dot effect, whereby popular data becomes less accessible because of the load of the requests on a central server. The following Fig. 4 demonstrates the percentage increase in the number of agents in the queue for a resource e.g. resource r_{12} in this figure. The figure also displays the corresponding decline in the throughput for processes requiring the service r_{12} .

This implies that popular request for a service can lead to it being highly advertised or “vocalised”, resulting in the depletion and decreased performance of the service. Therefore, unless there is a way to adapt to this phenomenon, the services will continue to fail or perform at a sub-optimal behaviour. Current work is aimed at studying the possibility of introducing service replication in the locality of the current service. This will distribute the load of the service and help process more requests. Also, it will handle to a certain extent the dynamic nature of the system wherein the services can fail.




Fig. 4. Demonstration of Slash dot effect at a resource r_{12} and the corresponding drop in throughput for processes requiring that resource.

Kleptoparasitic behaviour [14] is observed when an agent frequently refers to the environment for information regarding resources instead of foraging itself. Also, there isn't a change observed in its cooperative strategy, if anything, there is evidence of decreasing cooperation. Implying that the agent is satisfied getting most of its information from other agents that have published/shared the information and itself does not gather information.

3.1.4 Vocalisation/Persistence of Resources:

There are various resources that appear and disappear in the system over the duration of the simulation. The requests and usage of resources helps reinforce their life in the simulation. The Fig. 5 below shows a graph indicating the appearance and persistence of resources during one run of the simulation.




Fig. 5. Vocalisation/Persistence of resources.

After the emergence of a resource, its life depends upon the reinforcement or failure. The figure gives an indication of the existence and use of resources at particular times, and helps hypothesize the prospects of their consumption, which is very helpful in evaluating the various strategies being used for sharing information. The perseverance of some particular strategies in the system gives an indication of the behaviour of agents to particular situations and also, determines if certain behaviours recur in agents over time, but without them having the benefit of hindsight.

4 Discussion/Conclusions

Our experiments explore various cooperative/competitive strategies that encompass most aspects of social behaviour. Mixed strategy models [8], [9] showing the possibility of freeloaders or lying. Ongoing implementations include scenarios like modelling trust in the system, altruism, and misinformation/malicious agents. To show how information sharing models can make novel, quantitative, and testable predictions concerning social foraging theory, within the application domain of distributed systems e.g. P2P networks.

The experiments reveal some interesting dynamics of the system with respect to the information dissemination algorithm. Our main objective has been to keep the agent imperceptible and its behaviour very simple, and to understand the local dynamics of interacting agents that lead to complex global behaviours. We draw our inspiration for this work from biological social networks, e.g. Ant colonies, Bee colonies, and other relevant theories in behavioural ecology. We are currently developing formalisations for the current algorithmic approach, so as to do a detailed mathematical analysis of the underlying theory. Our study hopefully gives insights into certain kinds of behaviour persistent in the system, which bear some resemblance to biological social systems. Especially to areas such as foraging, danger of predation, sharing information regarding food/nest sites etc, [17], [5], and [16]. Issue of trust and reputation once incorporated into the simulation should yield some more interesting dynamics. The simulation model discussed should eventually be able to help understand some of the contexts in which cooperation emerges, is beneficial or not, and to what extent.

References

1. Caraco and Giraldeau: *Social Foraging Theory*, Princeton University Press. (2000)
2. Luc-Alain Giraldeau, Thomas Valone et. al.: Potential disadvantages of using socially acquired information, *Phil. Trans. of the Royal Society, London*. (2002)
3. Mangel, M.: Dynamic Information in uncertain and changing worlds. *J. Theoretical Biology* (1990) 146: 317-332
4. Stephens, D.W.: Variance and the value of information. *American Naturalist* (1989) 134: 128-140
5. Seeley, T.: *The Wisdom of the Hive – The Social Physiology of Honey Bee Colonies*, Harvard University Press, Cambridge (1995)
6. Clark, C.W., and Mangel, M.: The Evolutionary Advantages of Group Foraging. *Theoretical Population Biology*, Vol. 30 (1986)
8. Axelrod: *The Evolution of Cooperation*, Basic Books Inc. (1984)
9. Axelrod: *The Complexity of Cooperation, Agent-Based Models of Competition and Collaboration*. Princeton University Press. (1997)
10. Dugatkin: *Game Theory and Animal Behaviour*, Oxford University Press (1998)
11. Pacala, S.W., Gordon, D.M., Godfray, H.C.J.: Effects of social group size on information transfer and task allocation. *Evolutionary Ecology*, (1996) 10, 127-165
12. Hamilton: Altruism and Related Phenomena, Mainly in Social Insects. *Annual Review of Ecology and Systemics* (1972) 3:193-232
13. Dall and Johnstone: Managing Uncertainty: Information and insurance under the risk of starvation. *Phil. Trans. Royal Society. London*. (2002)
14. Hamilton: Kleptoparasitism and the distribution of unequal competitors. *Behavioural Ecology*. Vol. 13, (2001) 2: 260-267
15. Real and Caraco: Risk and Foraging in Stochastic Environments. *Annual Review Ecol. Syst.* (1986) 17:371-90
16. Caraco, T., Uetz, G.W., Gillespie, R.G., Giraldeau, Luc-Alain.: Resource Consumption Variance within and among individual: On coloniality in Spiders. *Ecology*, (1966) 76(1), pp. 196-205
17. Franks, N.R., Pratt, S.C., Mallon, E.A., Britton, N.F., Sumpter, D.: Information flow, opinion polling and collective intelligence in house-hunting social insects. *Phil. Trans. Royal Society. London* (2002)
18. Beckers, R., Goss, S., Deneubourg, J.L. & Pasteels, J.M.. Colony size, communication and ant foraging strategy. *Psyche* 96, 239-256 (1989).
19. Jaffe, K. & Hebling-Beraldo, M. J. Respirometry and the evolution of order: negentropy criteria applied to the evolution of ants. In *Proceedings of the 11th Conference of the International Union for the Study of Social Insects* (Bangalore, India) (ed. G. K. Vereesh, B. Mallik and C. A. Viraktamath), p. 538. Oxford and IPH Publishing Co., New Delhi (1990).
20. Deneubourg, J. L., Pasteels, J.M. & Verhaeghe, J.C. Probabilistic behaviour in ants: a strategy of errors? *Journal of Theoretical Biology* 105, 259-271 (1983).
21. Deneubourg, J. L., Aron, S., Goss, S. & Pasteels, J.M. Error, communication and learning in ant societies. *European Journal of Operations Research* 30,168-172 (1987).
22. Weidenmuller, A. & Seeley, T. D. Imprecision in waggle dances of the honeybee (*Apis mellifera*) for nearby food sources: error or adaptation. *Behavioral Ecology and Sociobiology* 46, 190-199 (1999).
23. Seeley, T. D. Honey bee colonies are group-level adaptive units. *American Naturalist* 150, 22-41 (1997).
24. Bonabeau, E. Social insect colonies as complex adaptive systems. *Ecosystems* 1, 437-443 (1998).