Collective Decision Making in Distributed Systems Inspired by Honeybees Behaviour

(Extended Abstract)

Andreagiovanni Reina,[1] Marco Dorigo[1] and Vito Trianni[2]

[1]IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium
[2]ISTC, Italian National Research Council, Rome, Italy
{areina@ulb.ac.be, mdorigo@ulb.ac.be, vito.trianni@istc.cnr.it}

We propose *cognitive design patterns* (CDPs) as a general methodology to design distributed artificial systems capable of cognitive processing. CDPs are reusable solutions that leverage the current understanding of cognitive processing in natural distributed systems, and that put this knowledge in use for the design of artificial ones [2]. In this paper, we derive a CDP for collective decision making inspired by the nest-site selection behaviour of honeybee swarms [1]. During this process, honeybees leave the swarm in search of a new nest site. A bee that discovers a potential nest gets committed to it and returns to the swarm to recruit other uncommitted scouts through the waggle dance. Bees committed to different alternatives deliver to each-other *stop signals*, which make them return to a uncommitted state. Additionally, scouts have a certain probability of spontaneously abandoning commitment. This collective process is based on peer-to-peer interactions among bees, and leads to a consensus decision for the best available alternative. Additionally, thanks to the cross-inhibition mechanism implemented through stop signals, the colony is able to break decision deadlocks in case of same-quality alternatives, and to randomly selects any of the two.

An analytical model of the nest-site selection process in a binary-choice scenario has been developed and confronted with empirical results [1]. The model is a ODE system that describes the dynamics of the fraction $\Psi_i = N_i/N$ of individuals belonging to population $i \in \{U, A, B\}$, where U is the population of uncommitted scouts, and A and B are the populations of bees committed to option A or B. The model assumes that agents switch between populations through four types of transitions with constant rate: discovery (γ_i), abandonment (α_i), recruitment (ρ_i) and cross-inhibition (σ_i):

$$\begin{cases} \dot{\Psi}_A = & \Psi_U(\gamma_A + \rho_A \Psi_A) - \Psi_A(\alpha_A + \sigma_B \Psi_B) \\ \dot{\Psi}_B = & \Psi_U(\gamma_B + \rho_B \Psi_B) - \Psi_B(\alpha_B + \sigma_A \Psi_A) \end{cases} \quad (1)$$

where $\Psi_U = 1 - \Psi_A - \Psi_B$.

It is worth noting that this model does not require any explicit comparison of the alternatives' quality by the single individuals. The quality value of the two alternatives is instead encoded in the transition rates: different-quality alternatives correspond to biased transition rates, while same-

quality alternatives to unbiased ones. An extensive analysis of the model showed that cross-inhibition rates σ_i determine a phase transition of the system that allows to break the deadlock in problems with equal quality alternatives. For low rates of cross-inhibition, the system remains deadlocked at indecision with equal number of individuals committed to either alternative ($\Psi_A = \Psi_B$), while for cross-inhibition values greater than the critical value $\sigma^* = \frac{4\alpha\gamma\rho}{(\rho-\alpha)^2}$ the system converges to a collective choice [1]. Therefore, through a suitable parameterisation it is possible to control the global system behaviour and the resulting collective dynamics.

The mechanistic description above clarifies the working regimes and suggests how the transition rates should be chosen to obtain the desired macroscopic behaviour. However, to guide the implementation of a distributed multi-agent system it is also necessary to define the main features of (i) the individual agent behaviour, (ii) the agent-to-agent interactions (e.g., information exchange and integration), and (iii) spatial and topological factors (e.g., connection topology among agents). These features determine the CDP, and incorporate the knowledge gained from the theoretical models to provide the minimal requirements for obtaining the desired system behaviour.

Following this line, we define the collective decision making CDP, and we demonstrate its application through a simple, spatial multi-agent scenario. The case study is the collective choice of the shortest path between two alternatives in a 1D space: agents move on a circle and need to collectively select and exploit the shortest path between two *target areas*. Two alternatives are possible: the upper and the lower path, respectively labelled as A and B. The angle θ between the target areas defines the decision problem: the best alternative is A for $\theta < \pi$, B for $\theta > \pi$ and any of the two for $\theta = \pi$. To identify and exploit a path, agents need to navigate back and forth between target areas. We assume that agents have local sensing and communication, move at constant speed, track the angular distance of the areas through dead reckoning and cumulate error in the position estimates. All agents start uncommitted with no knowledge about the target areas. Starting from the theoretical model, we have developed the CDP that we implement as follow:

(i) *Discovery*: an agent explores the environment through correlated random walk, and gets committed to a path as soon as it stores the position of the two target areas. In this way, we obtain $\gamma_i \propto 1/\theta$ because shorter paths are easier to discover through random walk;

(ii) *Abandonment*: an agent abandons its commitment and

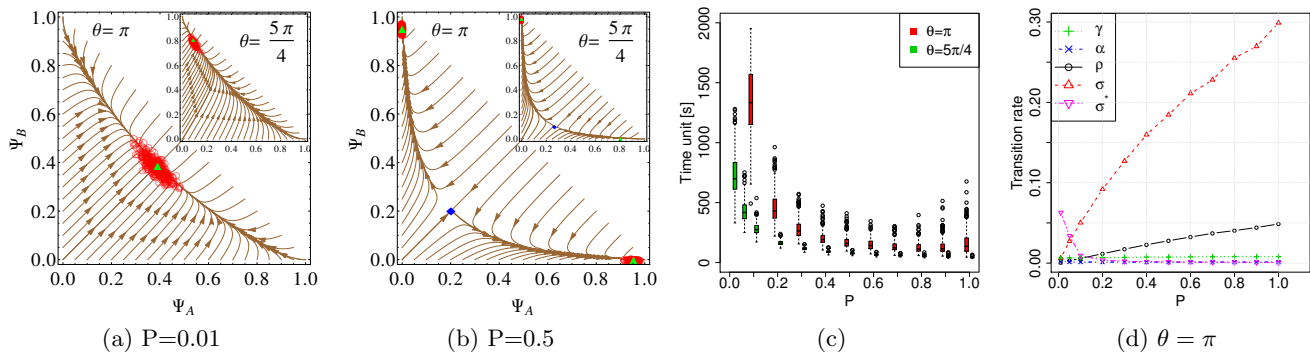(a) P=0.01   (b) P=0.5   (c)   (d) $\theta = \pi$

Figure 1: (a)-(b) Comparison between macroscopic dynamics (trajectories, stable points as triangles, and saddle points as rhombus) and 300 multi-agent simulations (final repartition of $N = 400$ agents between the two populations, shown as red empty dots) for the symmetric case $\theta = \pi$ (main) and the asymmetric case $\theta = 5\pi/4$ (insets). (c) Convergence time for symmetric and asymmetric cases as a function of P, with decision threshold at $\Psi_i = 0.7$ and $t_{max} = 2000s$. (d) Transition rates as a function of P for the symmetric case.

resumes exploration if it fails to attain a target area due to errors in the position estimates. Here, we obtain $\alpha_i \propto \theta$ because lower abandonment rates result from smaller cumulative error on shorter paths.

(iii) *Recruitment*: an uncommitted agent that interacts with an agent committed to alternative i gets recruited with fixed probability P_ρ.

(iv) *Cross-inhibition*: an agent committed to alternative i that interacts with an agent committed to alternative $j \neq i$ becomes uncommitted with fixed probability P_σ: it erases the stored locations and resumes exploration.

(v) *Well-mix*: to provide equal probability of interaction with agents exploiting different paths, interactions are possible only when agents are within the same target area. Each agent has a maximum of one interaction per time unit.

Note that we have not specified a direct way to control the transition rates for discovery and abandonment, while recruitment and cross-inhibition are determined by the control probabilities P_ρ and P_σ. We choose fixed probabilities independently of the possible differences in the path lengths. As discussed above, this should be sufficient to produce a collective choice, provided that the discovery rates are biased toward the best option. To simplify the system analysis, we fix $P_\rho = P_\sigma = P$ which we refer to as the *interaction probability*. We study the system behaviour varying P and θ, while the other parameters are kept constant.

To verify the correctness of the design pattern and to study how the collective behaviour changes as a function of P and θ, we check the adherence of the multi-agent system with the macroscopic model (Figure 1(a)-(b)). To this purpose, we statistically estimate the transition rates directly from the simulations through survival analysis. The agreement between multi-agent system and macroscopic dynamics is remarkable: the predictions of the model with estimated parameters perfectly match the final agent distribution. In the symmetric case with $\theta = \pi$, the alternatives have same quality and this—potentially—corresponds to a decision deadlock. This is actually the case for very low values of the interaction probability (e.g., $P = 0.01$ shown in the main plot of Figure 1(a)). For increasing interaction probability, we observe a phase transition which lets the system

break the symmetry: two stable solutions appear indicating a collective choice for either A or B (Figure 1(b)-main).

In case $\theta = 5\pi/4$, the decision problem should lead to the systematic choice of the alternative B. We note that for low values of the interaction probability P there exists a single stable fixed point for $\Psi_B \approx 1$ (inset of Figure 1(a)); while, for higher values of P, the system undergoes a bifurcation, and a second stable fixed point appears for the inferior option (inset of Figure 1(b)). The bifurcation appears when cross-inhibition is sufficiently strong compared to the other transition rates. This may lead to errors in the decision making if the system happens to be in the basin of attraction of the inferior choice. Conversely, larger cross-inhibition rates lead to increased decision speed (see Figure 1(c)).

Figure 1(d) shows how the estimated transition rates vary with respect to P for the symmetric case. While discovery and abandonment remain roughly constant, both ρ and σ increase quasi-linearly with P, indicating that a higher probability of interaction among agents directly translates in increased recruitment and cross-inhibition rates. Note that the estimated cross-inhibition rate is initially below the critical value ($\sigma < \sigma^*$) for small interaction probabilities ($P < 0.07$), which are actually the values at which the multi-agent simulations remain deadlocked at indecision. For larger P, cross-inhibition is sufficiently high and the collective decision is efficiently performed.

Acknowledgements

REFERENCES

[1] T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–11, 2012.

[2] V. Trianni, E. Tuci, K. M. Passino, and J. A. R. Marshall. Swarm Cognition: an interdisciplinary approach to the study of self-organising biological collectives. *Swarm Intelligence*, 5(1):3–18, 2010.