A design pattern for best-of-n collective decisions

Andreagiovanni Reina,[†] Gabriele Valentini,[†] Cristian Fernández-Oto,[‡]
Marco Dorigo,[†] and Vito Trianni[⋆]
[†]IRIDIA and [‡]Faculté de Sciences, Université Libre de Bruxelles, Belgium
[⋆]ISTC, Consiglio Nazionale delle Ricerche, Italy

The engineering of large-scale decentralised systems requires sound methodologies to guarantee the attainment of the desired macroscopic system-level behaviour given the microscopic individual-level implementation. While a general-purpose methodology is currently out of reach, specific solutions can be given to broad classes of problems by means of well-conceived *design patterns*, which recall the well-known concept exploited in software engineering [1]. Design patterns provide formal guidelines to deal with recurring problems in a specific domain. For the particular case of distributed systems, design patterns should prescribe the individual-level microscopic behaviour required to obtain desired system-level macroscopic properties [2, 3].

Here, we present a design pattern for decentralised decision-making—a fundamental ability in several contexts and application domains [4, 5, 6]. The design pattern is based on the nest-site selection behaviour of honeybee swarms [7, 8, 9]. Previous experimental and theoretical studies have demonstrated near-optimal speed-accuracy tradeoffs in the selection of the most profitable option among a set of alternative nesting sites by honeybees [7, 8]. Most importantly, inhibitory signals among bees provide an adaptive mechanism to quickly break deadlocks and tune the decision dynamics according to the perceived quality of the discovered options [8, 9]. The above properties of the nest-site selection process are relevant for many practical decision-making scenarios in decentralised systems.

Starting from the macroscopic description of the nest-site selection dynamics [8, 9], we derive the exact relationship between microscopic and macroscopic models—also including finite-size effects—for the general case of a best-of-n decision problem. The inter-related models represent the core of the design pattern, which is completed by formal guidelines for the implementation of collective decisions in multiagent systems. We provide guidelines for implementation by means of either homogenous or heterogeneous agents, as well as guidelines for the inclusion of spatial and topological factors that have a bearing in determining the microscopic interaction patterns. We report here a case study that illustrates the application of the design pattern, and we briefly discuss the results obtained in other case studies, as well as the relevance of the obtained results for better understanding the behaviour of natural systems.

Models We consider a best-of-n decision problem, that is, the choice of the best option, or any of the equal-best options, among n different alternatives. Each option $i \in \{1, \ldots, n\}$ is characterised by its quality $v_i \in [v_m, v_M]$. We study decision making for a population of N agents, where each agent is either committed to one of the available options i (sub-population size N_i and fraction $\Psi_i = N_i/N$) or is uncommitted (sub-population size N_U and fraction Ψ_U). At the macroscopic level, a decision is taken as soon as the entire population or a large fraction Ψ_q (hereafter referred to as *quorum*) is committed to a single option.

According to the model proposed for honeybee nest-site selection [8], four concurrent processes determine the distribution of agents across populations: (i) uncommitted agents spontaneously *discover* option i at rate γ_i (ii) agents committed to option i spontaneously *abandon* commitment at rate α_i; (iii) agents committed to option i *recruit* uncommitted agents at rate ρ_i; and (iv) agents committed to option $j \neq i$ *inhibit* agents committed to option i at rate σ_j. The mean-field macroscopic dynamics are well described by an n-dimensional ODE system, which extends the binary version discussed in [8, 9]:

$$\begin{cases} \dot{\Psi}_i &= \gamma_i \Psi_U - \alpha_i \Psi_i + \rho_i \Psi_i \Psi_U - \sum_{j \neq i} \sigma_j \Psi_i \Psi_j \\ \Psi_U &= 1 - \sum_i \Psi_i \end{cases}, \qquad i \in \{1, \ldots, n\} \tag{1}$$

The transition rates (γ_i, α_i, ρ_i and σ_i) are functions of the quality v_i:

$$\alpha_i = f_\alpha(v_i),\ \gamma_i = f_\gamma(v_i),\ \rho_i = f_\rho(v_i),\ \sigma_i = f_\sigma(v_i). \tag{2}$$

The relations between option quality and transition rates determine the macroscopic dynamics [9].

The microscopic behaviour of the average agent is represented by the probabilistic finite state machine (PFSM) shown in Fig. 1(a). It describes the *commitment dynamics* of a single agent in interaction with agents belonging to different sub-populations. An agent can be either uncommitted (state C_U) or committed to option i (state C_i), and changes state every τ seconds according to two types of transitions: spontaneous and interactive. Spontaneous transitions model discovery of the option i with probability $\mathcal{P}_\gamma(v_i)$ and abandonment of commitment to option i with probability $\mathcal{P}_\alpha(v_i)$. Interactive transitions model the recruitment and cross-inhibition processes resulting from the interaction between agents belonging to different populations. We refer to the probability of any agent interacting with an agent committed to option i as $\mathcal{P}_{\Psi_i}$. We assume a well-mixed system, so that $\mathcal{P}_{\Psi_i} = N_i/N$. Recruitment for option i is modelled by a transition from C_U to C_i with overall probability $\mathcal{P}_{\Psi_i}\mathcal{P}_\rho(v_i)$. Cross-inhibition of an agent committed to option i is instead modelled as the cumulative effect of the interaction with agents committed to a different option, with overall probability $\sum_{j\neq i}\mathcal{P}_{\Psi_j}\mathcal{P}_\sigma(v_j)$.

Implementation guidelines The actual implementation of the agent behaviour requires choosing the way in which transitions are executed in relation to the limited information available to the individual agent. For instance, the estimation of the population-size dependent probability $\mathcal{P}_{\Psi_i}$ by individual agents requires some sampling of the current population size. The other transition probabilities ($\mathcal{P}_\lambda(v_i), \lambda \in \{\gamma,\alpha,\rho,\sigma\}$) should instead provide an exact relationship with the macroscopic transition rates to obtain the desired dynamics determined by Eq. (2).

We propose two strategies based either on a *homogeneous* or on a *heterogeneous* implementation. In the homogeneous case, all agents compute their transition probabilities in the same way as a function of the estimated quality $\hat{v}_i$. In this case, it is possible to establish a direct correspondence between micro and macro parameters:

$$P_{\lambda,g}(\hat{v}_i) = \lambda_i\tau = f_\lambda(\hat{v}_i)\tau, \qquad \begin{array}{c} i \in \{1,\ldots,n\} \\ \lambda \in \{\gamma,\alpha,\rho,\sigma\} \end{array} \tag{3}$$

where $P_{\lambda,g}$ represents the actual probability for the agent a_g to undergo the transition λ. In the heterogeneous case, agents compute their own transition probabilities differently from each other. We propose a simple response threshold scheme, so that agent a_g follows a transition with a fixed probability if the (estimated) option quality $\hat{v}_i$ exceeds a given threshold δ_g:

$$P_{\lambda,g}(\hat{v}_i) = \left\{ \begin{array}{ll} \mathcal{P}_{\lambda\uparrow} & \hat{v}_i > \delta_g \\ \mathcal{P}_{\lambda\downarrow} & \hat{v}_i \leq \delta_g \end{array} \right., \qquad \begin{array}{c} i \in \{1,\ldots,n\} \\ \lambda \in \{\gamma,\alpha,\rho,\sigma\} \end{array} \tag{4}$$

where $\mathcal{P}_{\lambda\uparrow}$ and $\mathcal{P}_{\lambda\downarrow}$ are tuneable parameters, and the value δ_g is drawn for each agent a_g from a probability distribution $\mathcal{D}_\lambda$ over the range $[v_m, v_M]$. With this implementation, it is possible to establish a relationship between microscopic and macroscopic parameters through the cumulative distribution function $F_{\mathcal{D}_\lambda}$ of $\mathcal{D}_\lambda$:

$$F_{\mathcal{D}_\lambda} = \frac{\lambda\tau - \mathcal{P}_{\lambda\downarrow}}{\mathcal{P}_{\lambda\uparrow} - \mathcal{P}_{\lambda\downarrow}}, \qquad \lambda \in \{\gamma,\alpha,\rho,\sigma\} \tag{5}$$

For both homogeneous and heterogeneous strategies, the derivation of the relationship between microscopic and macroscopic description levels passes through the introduction of the master equation of the finite-size macroscopic model.

In a practical application scenario, agents might not be able to interact with neighbours every τ seconds. For instance, an agent might be busy estimating the quality of a discovered option, or spatial/topological factors might prevent frequent interactions. Agents unable to interact are latent, as opposed to interactive ones. We model changes in this activity state (i.e., the *activity dynamics*) by considering that an agent becomes latent with probability $\mathcal{P}_L$, and returns interactive with probability $\mathcal{P}_I$. In these conditions, a fraction of $\eta_I = \mathcal{P}_I/(\mathcal{P}_I + \mathcal{P}_L)$ agents can be found asymptotically in the interactive state. Depending on the microscopic implementation, it is possible to derive the correspondence between micro and macro parameters by dividing the macroscopic transition rates by η_I or η_L.

Case study We illustrate the implementation of decentralised decision-making for a multiagent system in which each agent can potentially interact with any other agent. We focus on the general case of value-sensitive decision-making as described in [9]. We consider a quality range $v \in [1, 10]$, we fix the quorum for the collective decision to $\Psi_q = 0.8$ and we limit the total execution time to $T = 40$ s. Following [9], discovery

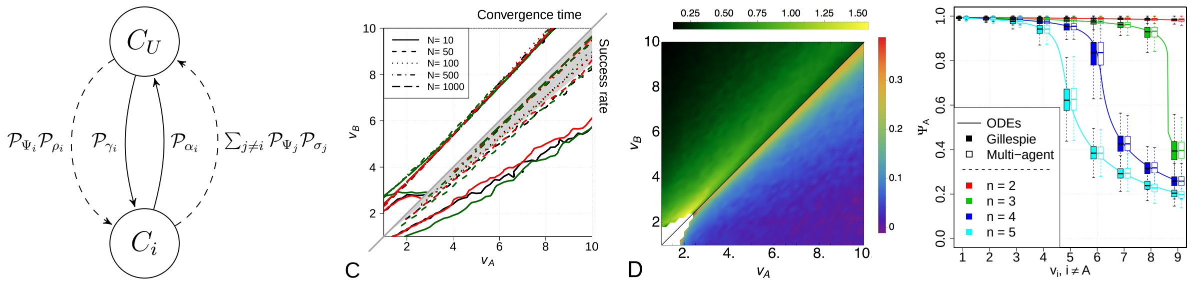

(a) (b) (c) (d)

Figure 1: (a) Probabilistic Finite State Machine (PFSM) describing the microscopic behaviour of an agent in average. Spontaneous transitions are represented by bold lines, interactive transitions by dashed lines. (b) Comparison between the stochastic finite-size macroscopic model (black lines) and the multiagent implementation with both the homogeneous strategy (red lines) and the heterogenous strategy (green lines). Results are displayed for varying system size N. For each possible configuration (v_A, v_B), 500 independent runs are performed. In the bottom-right half of the plot, we show the isolines for the success rate at the value $\mathcal{S} = 0.9$, and the gray triangle indicates quality value pairs below the target resolution $R = |v_A - v_B|/\max(v_A, v_B) = 0.15$. In the top-left half we show for symmetric quality pairs (v_B, v_A) the isolines for convergence time at the value $\mathcal{C} = 1\,\mathrm{s}$. (c) Scaling of the convergence time $\mathcal{C}$ with the system size N. For each configuration (v_A, v_B)—and the symmetric case (v_B, v_A)—we fitted the curve $\mathcal{C} = b\,N^a$ and we show the heat-map for the fitted coefficient a (bottom-right) and b (top-left) across the decision space. (d) Micro-macro link with varying number of options. We compare the macroscopic dynamics predicted by the mean-filed model (1), the finite-size macroscopic dynamics approximated by the Gillespie algorithm and the microscopic dynamics resulting from homogeneous multiagent simulations ($N = 500$ agents). The plot shows the fraction of the population committed to option A at the end of the simulation, plotted against the lower option quality v_i. Boxes represent the inter-quartile range of the data (2000 runs), while the horizontal lines inside the boxes mark the median values. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range. Outliers are not shown.

and recruitment rates are assumed to be linearly proportional to the option quality v_i (i.e., $\gamma_i = \rho_i = v_i$), the abandonment rate is inversely proportional to quality (i.e., $\alpha_i = 1/v_i$), while the cross-inhibition rate is constant ($\sigma_i = \bar{\sigma}$), which we fix to $\bar{\sigma} = 10$. Given the macroscopic parameterisation, we analyse the finite-size effects produced by the system size N by approximating the macroscopic dynamics using the Gillespie algorithm [10]. Then, we design the multiagent system following both the homogeneous and the heterogeneous strategies mentioned above.

We first focus on a binary decision problem, in which the available options are referred to as A and B, and their quality as v_A and v_B. Figure 1(b) shows the match between the macroscopic Gillespie simulations and the multiagent implementation with both the homogeneous and the heterogeneous strategy, for varying system size N. The correspondence between macroscopic model and microscopic implementation is remarkable. The results show that the studied parameterisation allows to reliably take decisions for above-resolution decision problems already with $N = 100$, as indicated by the success rate $\mathcal{S}$ in the bottom-right part of Fig. 1(b). Conversely, the convergence time $\mathcal{C}$ is very similar across different system sizes. We analysed the scaling behaviour of the convergence time and found adherence with a power law behaviour $\mathcal{C} = b\,N^a$ (see Fig. 1(c)). With the proposed parameterisation, $\mathcal{C}$ becomes nearly independent of the system size N in large parts of the problem space. Finally, we studied the micro-macro link in a best-of-n scenario. We fix the best option (A) to the maximum quality $v_A = 1$, and all other options to the same, lower quality v_i. The results presented in Fig. 1(d) reveal a very good correspondence between multiagent and Gillespie simulations, therefore validating the methodology beyond the binary decision problems presented above.

Discussion The design pattern methodology we propose provides a complete framework that allows to move from the choice of the macroscopic parameterisation down to the implementation of the individual behaviour. Each step is supported by the principled understanding of the causal relationship between microscopic choices and macroscopic effects. We have substantiated the methodology with other case studies beyond the one presented here. In particular, we studied the micro-macro link for a problem of resource discovery and exploitation. In this case study, mobile agents have to recognise and collectively select the best option among several available. The spatial factors have a bearing on the interaction between agents, so that particular care must be given to the implementation. The design pattern provides guidelines also for such case, thanks to the inclusion of latent states for individual agents that allows to preserve the micro-macro link also when interactions are sporadic or when spatiality interferes with the well-mixed assumption.

Besides engineering, our results can be relevant for better understanding the behaviour of natural systems, by providing testable hypotheses to be verified by field experiments. The heterogeneous implementation strategy represents one such case. The choice of response thresholds is supported by the large literature on inter-individual variability in social insects [11]. Recent studies have recognised the importance of including individual differences in behaviour to better understand the dynamics of collective behaviours [12, 13]. Here, we have highlighted the relationship between the distribution of individual thresholds and the collective response function, so that macroscopic predictions could be matched against estimates of the real threshold distribution [14]. Response thresholds well adhere with adaptive mechanisms for threshold adaptation, allowing to finely tune the macroscopic response to match the statistical regularities that characterise the task. This adaptivity can result from evolutionary factors as well as from development and learning [11]. Integrating adaptive mechanisms in the microscopic implementation could lead to improved performance, and represents a natural extension for the proposed design pattern.

References

[1] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional, Boston, MA, 1995.

[2] G. Babaoğlu, G. Canright, A. Deutsch, G. A. Di Caro, F. Ducatelle, L. M. Gambardella, N. Ganguly, M. Jelasity, R. Montemanni, and A. Montresor. Design patterns from biology for distributed computing. *ACM Transactions on Autonomous and Adaptive Systems*, 1(1):26–66, 2006.

[3] A. Reina, R. Miletitch, M. Dorigo, and V. Trianni. A quantitative micro-macro link for collective decision: The shortest path discovery/selection example. *Swarm Intelligence*, (in press) http://dx.doi.org/10.1007/s11721-015-0105-y, 2015.

[4] I.F. Akyldiz, B.F. Lo, and R. Balakrishnan. Cooperative spectrum sensing in cognitive radio networks: A survey. *Physical Communication*, 4(1):40–62, 2011.

[5] V. Srivastava and N.E. Leonard. Collective decision-making in ideal networks: The speed-accuracy tradeoff. *IEEE Transactions on Control of Network Systems*, 1(1):121–132, 2014.

[6] M. Vigelius, B. Meyer, and G. Pascoe. Multiscale modelling and analysis of collective decision making in swarm robotics. *PLoS ONE*, 9(11):e111542–19, 2014.

[7] J.A.R. Marshall, R. Bogacz, A. Dornhaus, R. Planqué, T. Kovacs, and N.R. Franks. On optimal decision-making in brains and social insect colonies. *J R Soc Interface*, 6(40):1065–74, 2009.

[8] 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.

[9] D. Pais, P.M. Hogan, T. Schlegel, N.R. Franks, N.E. Leonard, and J.A.R. Marshall. A mechanism for value-sensitive decision-making. *PLoS ONE*, 8(9):e73216, 2013.

[10] D.T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J Phys Chem A*, 81(25):2340–2361, 1977.

[11] R. Jeanson and A. Weidenmüller. Interindividual variability in social insects - proximate causes and ultimate consequences. *Biol Rev Camb Philos Soc*, 89(3):671–687, December 2013.

[12] M.K. Wray and T.D. Seeley. Consistent personality differences in house-hunting behavior but not decision speed in swarms of honey bees (*Apis mellifera*). *Behav Ecol Sociobiol*, 65(11):2061–2070, 2011.

[13] I. Planas-Sitjà, J.-L. Deneubourg, C. Gibon, and G. Sempo. Group personality during collective decision-making: a multi-level approach. *Proc Biol Sci*, 282(1802):20142515, 2015.

[14] A. Weidenmüller. The control of nest climate in bumblebee (bombus terrestris) colonies: interindividual variability and self reinforcement in fanning response. *Behav Ecol*, 15(1):120–128, 2004.