

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Cognitive swarming: an approach from the theoretical neuroscience of hippocampal function

Joseph D. Monaco, Grace M. Hwang, Kevin M. Schultz, Kechen Zhang

Joseph D. Monaco, Grace M. Hwang, Kevin M. Schultz, Kechen Zhang, "Cognitive swarming: an approach from the theoretical neuroscience of hippocampal function," Proc. SPIE 10982, Micro- and Nanotechnology Sensors, Systems, and Applications XI, 109822D (13 May 2019); doi: 10.1117/12.2518966

SPIE.

Event: SPIE Defense + Commercial Sensing, 2019, Baltimore, Maryland, United States

Cognitive swarming: An approach from the theoretical neuroscience of hippocampal function

Joseph D. Monaco^a, Grace M. Hwang^b, Kevin M. Schultz^b, and Kechen Zhang^a

^aBiomedical Engineering, Johns Hopkins University School of Medicine, Baltimore, MD, USA

^bApplied Physics Laboratory, Johns Hopkins University, Laurel, MD, USA

ABSTRACT

The rise of mobile multi-agent robotic platforms is outpacing control paradigms for tasks that require operating in complex, realistic environments. To leverage inertial, energetic, and cost benefits of small-scale robots, critical future applications may depend on coordinating large numbers of agents with minimal onboard sensing and communication resources. In this article, we present the perspective that adaptive and resilient autonomous control of swarms of minimal agents might follow from a direct analogy with the neural circuits of spatial cognition in rodents. We focus on spatial neurons such as place cells found in the hippocampus. Two major emergent hippocampal phenomena, self-stabilizing attractor maps and temporal organization by shared oscillations, reveal theoretical solutions for decentralized self-organization and distributed communication in the brain. We consider that autonomous swarms of minimal agents with low-bandwidth communication are analogous to brain circuits of oscillatory neurons with spike-based propagation of information. The resulting notion of ‘neural swarm control’ has the potential to be scalable, adaptive to dynamic environments, and resilient to communication failures and agent attrition. We illustrate a path toward extending this analogy into multi-agent systems applications and discuss implications for advances in decentralized swarm control.

Keywords: distributed control, self-organization, swarming, spatial navigation, neuroscience, synchrony, emergence, place cells

1. INTRODUCTION

The rise of mobile multi-agent robotic platforms may potentially exponentiate the breadth of applications of autonomous technologies, but current control paradigms lack resilience and versatility in real-world environments. We suggest that this gap relates to the need to dynamically replan distributed responses to rapidly changing surroundings or other situational and task-relevant variables. Environmental, perceptual, and state-based sources of uncertainty may be defining features of critical future applications that could most benefit from the agility of low-footprint vehicle platforms. Crucially, solutions to online spatial decision-making in uncertain conditions have evolved in animals such as rodents whose daily survival depended on navigating to find food, take shelter, and escape predators.

The hippocampus is the center of lifelong memory formation in mammals including humans,¹ but research interest in this brain structure intensified after hippocampal neurons, known as ‘place cells,’ were discovered to carry a clear signal representing the animal’s current spatial location.² This was the first link in modern neuroscience³ between high-order cognition (e.g., the sense of place) and the output of single neurons (i.e., the ‘place field’ in which a place cell is active). Two major hippocampal phenomena, self-stabilizing attractors of spatial activity and temporal coding during shared oscillations, arise from the circuit-wide connectivity and collective activity of hippocampal networks.^{4–6} We suggest that these two phenomena outline possible brain-like control strategies for artificial swarms. Neurobiological circuits contend with noise, communication failures, and multi-level/multi-modal sources of variability and uncertainty in both perceptual and motivational states. Similarly, a ‘neural swarm control’ framework may confer adaptive resilience to critical real-world applications of autonomous multi-agent systems.

Further author information: (Send correspondence to J.D.M. or G.M.H.)

J.D.M.: E-mail: jmonaco@jhu.edu

G.M.H.: E-mail: grace.hwang@jhuapl.edu

Micro- and Nanotechnology Sensors, Systems, and Applications XI, edited by Thomas George,
M. Saif Islam, Proc. of SPIE Vol. 10982, 109822D · © 2019 SPIE
CCC code: 0277-786X/19/\$18 · doi: 10.1117/12.2518966

2. MOTIVATION

The growing development of multi-agent platforms with lightweight robotic vehicles may greatly expand the application domain of autonomous technologies, but current control frameworks lack adaptiveness, resilience, and computational efficiency, especially in complex, realistic, and dynamically changing conditions.^{7–10}

2.1 Limitations of state-of-the-art swarm control

At the level of algorithmic design, swarm control can be characterized by an intrinsic three-way trade-off between computational efficiency, stability, and performance.¹⁰ However, a more pragmatic trade-off driven by agent scale (e.g., physical size, mass, energy footprint/capacity) has similarly wide-ranging ramifications for computational control. Specifically, agent scale trades off against the degree of baseline autonomous capabilities allowed by onboard hardware (Fig. 1). While robotic vehicles are being engineered at scales spanning orders of magnitude, lower-footprint agents provide many benefits. With lower mass, agents have higher maximum angular acceleration and lower kinetic energy, which increases both safety and agility due to faster turning and velocity changes.¹¹ However, lower vehicle and payload weights restrict the diversity, range, and capacity of the onboard sensor suite that is critical to control, perception, and state estimation in autonomous applications.




Figure 1. The trade-off between agent scale and autonomy. Orders of magnitude differences in the scale of engineered robotic vehicles entail a wide range of sensing and communication capacities. The benefits of increased agility, safety, and maneuverability attained by reducing agent scale come at the expense of lower baseline autonomous capabilities of those vehicles. If the solution is a numerical increase in group size, then scalable, decentralized, and resilient control presents the crucial challenge. Inset (top left): This challenge might be approached by directly considering large swarms of small, limited agents that act as neurons in the spatial cognitive circuits of foraging animals such as rats.

The question becomes how to achieve high autonomy with small and functionally limited robots (Fig. 1, inset). Since smaller vehicles tend to have lower marginal costs, a promising approach is to scale up the number of robotic agents in an autonomous group as the size and computational complexity, etc., of the agents are scaled down. Current state-of-the-art approaches to autonomous control of large robotic swarms utilize fast onboard processors to continuously estimate state and compute optimal motion plans via exhaustive search over the space of possible trajectories with respect to neighboring agents and the environment at rates on the order of 1 kHz [e.g., 11]. If this represents a ceiling of computational power, then the resilience and adaptiveness of swarms in this regime (i.e., characterized by larger numbers of smaller, simpler agents) may benefit from alternative methodological approaches.

2.2 Place cells and self-organizing attractor maps

We conjecture that this inverse scaling of number and complexity may reflect, in the limit, biological solutions in evolved living systems that address the problem of collective goal-directed dynamics with strictly local information. For instance, the distributed neuronal networks in the brains of rodents seeking out food rewards (Fig. 1, inset) must integrate diverse and uncertain sources of information in order to produce a unified spatial plan for foraging and other spatial tasks. These neural computations in a rodent brain are limited to the interactions of $\sim 750,000$ hippocampal neurons¹² (cf. $\sim 1e9$ transistors in the Qualcomm[®] Snapdragon[™] 801 processor used in some quadrotor vehicles), each of which gradually build up sufficient capacitive charge across their cellular lipid bilayer membranes to generate an action potential (or ‘spike’) at a typical frequency (rate of occurrence, or ‘firing rate’) on the order of 1 Hz.^{2,13} This qualitatively different form of computation, without high sample rates or global optimization procedures, nonetheless allows a rat to make the spatial decisions that enable survival up to several years in the wild, during which it is always exploring farther out from home base.¹⁴

Hippocampal place cells, as described above, each fire within a restricted spatial region of the environment.² For example, a single place cell in a rat running laps on a circular track (Fig. 2A, top row) forms a new place field after an investigatory behavior in which the rat paused to visually examine the environment (Fig. 1, inset). This example shows a recently characterized behavioral-attentional mechanism of place field formation.¹⁵ Typically, it is not possible to discern which factors, external or internal, determine the location of a place field; regardless, the place fields of a population of place cells collectively form a map of the animal’s current environment³ (Fig. 2A, bottom row). A prominent theory of the stability of these spatial maps posits that attractor dynamics, including fixed points or continuous manifolds, drive place cell networks to represent low-dimensional spatial contexts.⁴ The idea is rather parsimonious, because models have shown that recurrent connectivity between place cells that perform nonlinear integration of their inputs, e.g.,

$$\frac{dr_i}{dt} = -r_i + g \left(\sum_j J_{ij} r_j + I_i \right), \quad (1)$$

where r_i is the firing rate of place cell i , I_i is its external spatial input, and g is a sigmoidal nonlinearity, is nearly sufficient to achieve a stable attractor map.^{16,17} The only additional requirement is that the recurrent weights J_{ij} represent the learned spatial associations between place cells (e.g., the degree of overlap between place fields) as a decremting function of distance between place field centers $\mathbf{x}$, such as

$$J_{ij} := F(\mathbf{x}_i - \mathbf{x}_j) = W_{\text{exc}} e^{-\frac{|\mathbf{x}_i - \mathbf{x}_j|}{\sigma}} - W_{\text{inh}}, \quad (2)$$

where σ is the spatial scale of an exponential kernel and W are the maximum excitatory and inhibitory synaptic strengths (note that this violates Dale’s law, which states that neurons only make connections of one type, and only serves to illustrate the expressiveness of the notion of attractor maps).¹⁷ A network constructed following Eqs. (1) and (2) supports self-organization of its activity into a singular, contiguous ‘bump’ that is anchored to the spatial attractor. Following random initialization, the activity bump will emerge as the network relaxes¹⁸ (Fig. 2B); the bump can then stably track the animal’s current position by responding to changes in external input as it moves around the environment.

2.3 Phaser cells and oscillatory phase synchronization

In addition to spatial modulation, the neuronal activity of the hippocampus and other spatial circuits is rhythmically modulated in time according to a continuous electric-field oscillation in the 5–12 Hz band called theta.⁵ Place cells exhibit a striking unidirectional shift, from late to early in the theta cycle, of spike theta-phase as a rat crosses the place field.²⁰ This ‘phase precession’ phenomenon is clearest on linear tracks, on which place fields arrange in an unambiguous sequential order; indeed, spike phase contributes additional spatial information (beyond firing rate alone) that allows more precise decoding of the rat’s location from neural recordings.²¹ However, in open two-dimensional environments, the relationship between the animal’s position and the theta phase of hippocampal neurons has remained unclear.




Figure 2. Place cells and phaser cells exemplify two neural dynamics of spatial computation. (A) Place cells were recorded on a 76-cm circular track as rats ran laps in the clockwise direction. A single place cell (top) that did not initially have a strong place field (left, circular firing-rate map; center text, maximum rate in spikes per s) eventually formed a place field (red, spikes) on laps 3 through 5 following activity during a visually attentive behavior (blue, spikes during a lateral head scan movement) that persisted throughout the recording (right, rate map for remainder of recording). (bottom) Place cells collectively map out the environment by forming fields at different locations. Adapted from Monaco et al. (2014) [15]. (B) A prominent theory to explain the spatial stability of place cell responses is the self-organization of attractor maps, in which the environment is represented by a continuum of locations in which overlapping place fields create self-reinforcing activity patterns in the recurrent network of the hippocampus. Adapted from Zhang (1996) [18]. (C) Phaser cells were found in the lateral septum, a subcortical brain area that receives dense input from the hippocampus. These cells convey spatial information in the timing of their spikes relative to the shared septal-hippocampal theta oscillation. (D) A simple model of rhythmic inhibition (sinusoid) and slowly-varying external input (triangle wave) accounts for the strong coupling of firing rate and theta phase observed in phaser cells. Adapted (C and D) from Monaco et al. (2019) [19].

To investigate this issue, we recently analyzed recordings from various subcortical brain areas connected to the hippocampus and discovered neurons, which we termed ‘phaser cells,’ that have a clear and unambiguous theta-phase code for spatial position in open environments.¹⁹ Phaser cells, which we found in the lateral septum, trade off positional precision for allocentric reliability; that is, instead of the unidirectional shift of hippocampal phase precession, phaser cells shift the phase of spiking in either direction according to the strength of external input at any given location in space. The signature of this symmetric phase modulation is that maps of average firing rate and average theta phase are highly correlated across space (see Fig. 2C for example phaser cell recordings showing the rate (top row) and phase (bottom row) maps of the 80-cm diameter recording arena). This strong rate-phase coupling can be accounted for with a simple model of excitatory external input interacting with inhibitory theta-rhythmic input (Fig. 2D). We took advantage of that mechanistic simplicity to demonstrate through simulations of oscillator networks that the diverse spatial patterns of phaser cells (Fig. 2C) may constitute an efficient basis set for creating arbitrary spatial patterns of phase synchrony.¹⁹

Location-dependent synchrony could serve various possible functions. A main component of spatial navigation, including in artificial autonomous systems, is path integration, which refers to updating internal estimates of position based on self motion. Based on the discovery of grid cells in the rodent brain,²² a theory was developed in which the relative phase amongst a set of velocity-controlled oscillators (VCOs) computes a path integration vector from a known reference point.^{23,24} It is unclear whether mammalian navigation utilizes the VCO (or any other) oscillatory mechanism,²⁵ but spatial synchrony driven by phaser cells could provide exactly the spatial phase-code feedback signal needed to counteract the accumulation of path integration errors over time.^{19,26,27} Regardless of experimental corroboration in the brain, in the consideration of engineered systems

these models from theoretical neuroscience can point to useful approaches. Certain properties of oscillatory phase-based mechanisms may benefit the communication design^{28–30} of large-scale autonomous systems: (1) the shared baseline oscillation may emerge dynamically via weakly-coupled rhythmic interactions; (2) oscillatory signals are necessarily repeated and thus robust to noise and communication failures; (3) population phase codes may synchronize within traveling waves that connect the disparate ends of a distributed system; (4) energy efficiency derives from the basis of pulsatile timing signals instead of rate, intensity, or amplitude modulation; (5) decentralized communication may arise from local phase organization that drives emergent higher-order states that reciprocally inform local behaviors.

3. THE BRAIN/SWARM ANALOGY

Swarms and the concept of ‘swarm intelligence’ have been analogized to the complex collective behaviors of groups of animals.³¹ Agent-local rule-based systems based on herding, flocking, and schooling have demonstrated complex emergent patterns in computer simulations and multi-vehicle platforms.^{32,33} However, these group-based behaviors typically reflect aggregative milling or dispersive responses in reaction to environmental changes. Achieving an active adaptiveness and resilience in numerically large multi-agent systems will require dynamics beyond close-to-equilibrium relaxation and reaction. Thus, we suggest that analogizing swarming behavior to the neural dynamics of a single intelligent animal’s brain (Fig. 3) may provide a path toward active, goal-directed spatial cognition in artificial swarms.




Figure 3. A swarming multi-agent group depicted as a neural network. (bottom) In-range peer-to-peer communication establishes a connection like the reciprocal synapses between two neurons. The connection graph can be considered as a weight matrix and reciprocal communication in local clusters of agents as the reverberating activity within a recurrent network. (middle) Instead of place fields (or phaser cell responses), agents construct an internal ‘task field’ representing preferences for certain environmental features or elements of the spatial task based on external inputs and each agent’s internal state, which may be represented as a phase variable. Images adapted for purposes of illustration from Ivancevic & Reid (2016) [34]. (top) Collective communication between agents could then drive processes of synchronization and self-organization as observed at the population level in neural circuits.

Targeting agent mobility to the level of a single neuron (Fig. 3) permits a distributed and decentralized paradigm of swarm control. The functional analogy between swarming and vertebrate neural architectures was discussed over a decade ago in the context of collective behaviors in animals, including honey bees³⁵ and ants,³⁶ but its impact across research disciplines has been unclear. By extending the notion of neural swarming to autonomous technologies, the neuroscientific principles of spatial cognition,³ decision making,^{37,38} and planning^{5,39,40} can be leveraged to provide the higher-order dynamics that allow individual animals to explore and exploit the real world. To satisfy the analogy, one crucial addition is required. As described above, neural communication is based on the propagation of spikes between neurons. Biophysically realistic spike generation may be formalized within phase models [e.g., 41], thus we posit that the internal state of neural swarming agents should be a phase variable (or collection of phase variables on a manifold).

4. NEURAL SWARM CONTROL

4.1 Swarming with mobile oscillators

The addition of an internal phase state to each agent implies that they may be understood as oscillators. Recently, O’Keeffe, Hong, & Strogatz (2017) [42] presented an expressive generalization of a class of biological systems that behave as populations of mobile oscillators whose phase state θ is coupled to spatial dynamics $\dot{\mathbf{x}}$, e.g.,

$$\dot{\mathbf{x}}_i = \frac{1}{N} \left[\sum_{j \neq i}^N \frac{\mathbf{x}_j - \mathbf{x}_i}{|\mathbf{x}_j - \mathbf{x}_i|} (A + J \cos(\theta_j - \theta_i)) - B \frac{\mathbf{x}_j - \mathbf{x}_i}{|\mathbf{x}_j - \mathbf{x}_i|^2} \right], \quad (3)$$

for scaling parameters A and B , and phase coupling J , and undergoes local phase synchronization, e.g.,

$$\dot{\theta}_i = \omega_i + \frac{K}{N} \sum_{j \neq i}^N \frac{\sin(\theta_j - \theta_i)}{|\mathbf{x}_j - \mathbf{x}_i|}, \quad (4)$$

for baseline frequency ω and synchronization parameter K . O’Keeffe et al. [42] demonstrated three static and two dynamic (or ‘active’) equilibria depending on coupling J and synchronization K . In the static states, the oscillator units spread outwards as the dynamics relax and converge on a regular hexagonal tiling of the unit disc (for $A = B = 1$) or an annulus that is angularly sorted according to phase (i.e., a ‘phase wave’). These states represent minimum-energy solutions given the mutual forces of attraction and repulsion, respectively governed by A and B in Eq. (3). The active states, however, frustrate the spatial and phase dynamics by pairing phase-tuned attraction with local desynchronization (e.g., $J \sim 1$ and $K \ll 0$) and produce constant milling or patrolling behaviors. Thus, this phase-based swarming formalism, termed ‘swarmalators,’ is a potentially useful framework for constructing more complex, higher-order dynamics.

4.2 Internal place fields for neural control

Beyond the addition of phase state, neural control for swarming agents must decouple the agent’s physical location from internal self-localization for two reasons. First, the multiplicity of agents is a qualitative difference between swarms and neural circuits; every hippocampal neuron corresponds to a single agent (e.g., the rat) and has particular connections and weights to other neurons with intrinsic recurrent and/or external sensory inputs.⁴³ That is, the individual rat has many place fields but the individual swarm agent should have only one, which cannot be identical to the agent’s objective location which depends on the complexities of a physical entity operating in the external world. Second, a recent line of experimental results have compellingly demonstrated that spatial path planning in hippocampal networks relies on activating sequences of remote locations represented by place cells.^{39,44,45} That is, dynamic replanning in mammalian spatial cognition may depend on internal representations that are separable from the animal’s (or agent’s) current objective position.

The notion of the internal place field (or phaser cell response) raises the question of its relationship to the agent and the agent’s motion dynamics. A straightforward approach might be to operate the swarming interaction (e.g., Eqs. (3) and (4)) on the internal position while the physical agent tracks the vector difference between its physical and internal positions in order to continuously approach the internal position. This method allows simulations to apply physics, energy constraints, or low-level controls (e.g., geometry-based obstacle avoidance) to each agent entity independently of the internal dynamics of the swarm. Given the distance-weight relationship described for attractor map connectivity in Eq. (2), this separation crucially provides a degree-of-freedom that may allow swarming to be recast as learning.

4.3 Neural swarming as mobile Hebbian learning

Inspection of the swarmalator dynamics in Eqs. (3) and (4) reveals a formal similarity to learning rules based on Hebbian associative synaptic modification.^{46–48} Hebbian learning in neural network models typically increments or decrements a synaptic weight according to a learning rate and a measure of the activity correlation between the pre-synaptic (input) and the post-synaptic (output) neurons. In this case, the phase coupling J acts as a learning rate and the cosine phase similarity provides the input/output correlation. Positive values of this

Hebbian-like measurement multiply the approaching unit vectors to bring the oscillator units closer together. To take this further, following the exponential attractor map kernel from Eq. (2), we could explicitly relate synaptic weight and agent distance, $W_{ij} = \exp(-D_{ij}/\alpha)$, for inter-agent distances $\mathbf{D}$ and spatial constant α . At this point, various approaches to updating the weight matrix $\mathbf{W}$ are possible.

We could consider that each agent's internal place field results from the triangulation of environmental cues. Modeling the field as a σ -bandwidth Gaussian function, we could define a time-continuous cue input

$$\tau \dot{c}_i = \exp \left(-\frac{1}{2\sigma^2} \left[\sum_{V_i^c} |D_i^c - D_i^{c*}| \right]^2 \right) - c_i, \quad (5)$$

from the cues V_i^c that are currently visible to unit i , which is distance D_i^c from the cue and whose internal place field center is distance D_i^{c*} from the cue. Similarly, recurrent input from peers in the swarm could follow

$$\tau \dot{q}_{ij} = V_{ij} \omega_j^h - q_{ij}, \quad (6)$$

where $\mathbf{V} \in \{0, 1\}^{N \times N}$ provides inter-agent visibility and ω^h is neural activation expressed in units of the baseline angular frequency ω of the phase state. With these feedforward ($\mathbf{c}$, cue) and recurrent ($\mathbf{q}$, swarm) inputs, we could define total 'conductances' to the neuron (agent) as

$$I_c = g_c \mathbf{c} \quad I_s = g_s \sum_{j=1}^N W_{\cdot j} \circ q_{\cdot j}, \quad (7)$$

where g_c and g_s indicate maximum conductance values and $\circ$ is the entrywise vector product. Similar to a model neuron, these inputs could drive a nonlinear activation function, e.g.,

$$\omega^h = k \tanh \left[\frac{1}{k} (I_c + I_s) \right], \quad (8)$$

where k sets the maximum activation relative to the baseline frequency. Thus, replacing Eq. (4), phase synchronization in this neural variation of the swarmalator formalism is expressed through the phase evolution $\dot{\theta} = 2\pi(1 + \omega^h) \circ \omega$. To construct a Hebbian updating rule, we lastly require a time-continuous post-synaptic signal, e.g., $\tau \dot{\mathbf{p}} = \omega^h - \mathbf{p}$. A naïve rule, such as $d\mathbf{W} = \eta \mathbf{q}^\top [\mathbf{p}, \dots, \mathbf{p}]$, would have weights grow unbounded; this causes pathological states in neural networks, and would cause a swarm to collapse into a singularity. Instead, if we assume that the sum of pre-synaptic weights is normalized to unity after each update, then the updated weights may follow

$$\mathbf{W}' = \mathbf{W} + \eta (\mathbf{q} - \mathbf{W}[\mathbf{p}, \dots, \mathbf{p}])^\top [\mathbf{p}, \dots, \mathbf{p}], \quad (9)$$

with learning rate η , according to Oja's derivation.⁴⁹ The effect of normalization on Hebbian learning in Eq. (9) is the addition of a subtractive term, quadratic in the post-synaptic signal $\mathbf{p}$, that depresses the growth of overly active synapses. In network models of place cell learning, such synaptic depression is critical to spreading out place fields to efficiently map an environment.^{50,51} Strikingly, this synaptic depression is functionally analogous to the quadratic repulsion term governed by the parameter B in Eq. (3); thus, a minimal form of stable Hebbian learning recapitulates generalized swarming dynamics.

To fully integrate neural network learning with swarming, the updated weights $\mathbf{W}'$ must drive agent movements. Given the exponential kernel that produced the baseline weight matrix $\mathbf{W}$, we can convert the updated weights into the desired distance between agents i and j ,

$$D'_{ij} = -\alpha \log W'_{ij}, \quad (10)$$

and perform a unit-wise averaging of the implied position shifts between mutually visible pairs of agents, e.g.,

$$d\mathbf{x}_i = \frac{1}{2 \sum_j V_{ij}} \sum_{j=1}^N V_{ij} (D'_{ij} - D_{ij}) \frac{\mathbf{x}_j - \mathbf{x}_i}{|\mathbf{x}_j - \mathbf{x}_i|}, \quad (11)$$

to compute the effective motion resulting from modified synaptic weights. The foregoing description illustrates one possible reformulation of phase-organized swarming as a dynamical neural network problem.

5. SUMMARY & DISCUSSION

We introduced the idea that neurobiological solutions in foraging animals may contribute to the development of many-vehicle autonomous technologies. We outlined a critical problem for artificial swarms (Fig. 1), which will require novel control schemes as robotic vehicles are scaled down (decreasing agent resources) and the numerical size of swarms is scaled up (increasing communication and coordination requirements). A similar distributed scaling problem may have been solved by the neural architecture of mammalian brains. Particularly, the neural circuits of the hippocampus and related brain areas support spatial cognition, decision making, and planning through all of the complex external conditions that, e.g., a rat will encounter throughout its life.

Neuronal networks such as hippocampal place cells and lateral septal phaser cells have produced theories of brain function derived from two main concepts: (1) that sparse, recurrent connectivity can produce attractor maps with self-organized, self-stabilizing activity patterns that reflect spatial learning and experience (Fig. 2A+B); (2) that phase coding based on shared oscillations can temporally organize spatial information as emergent global states (Fig. 2C+D). Oscillatory and attractor dynamics both derive from synaptic connections between neurons and the spikes that propagate across those connections. Rethinking swarm communication as spike propagation may reduce the density and complexity of peer-to-peer data messaging commensurate with the reduced communication capabilities of small-scale vehicles.

Compared to signal comprehension, signal production errors may be particularly deleterious to large-scale distributed inference and consensus.⁵² Thus, potential onboard suites for ‘cognitive swarming’ platforms should emphasize reliable, real-time transmission of pulsatile low-bandwidth (even single bit) packets. Low fidelity inputs from, e.g., cameras or receivers are more easily compensated by distributed processing within the swarm; i.e., sensor designs should emphasize energy and cost to maximize deployment duration and swarm size.

We made the explicit analogy from swarms and agents to neural circuits and neurons. This analogy allows the tools of theoretical neuroscience to be leveraged in developing artificial autonomous systems. We suggested two components that facilitate the implementation of the neural swarming analogy: (1) a phase-based internal state; (2) decoupling of the agent’s physical location from its internal self-localization. The phase state naturally drives spike generation and leverages the expressive complexity of mobile oscillator formalisms such as swarmalators (Eqs. (3) and (4)). The separation of physical/internal self-localization allows swarm motion dynamics to be interpreted as Hebbian-like learning in a place cell network. On this basis, emergent hippocampal phenomena such as oscillatory phase coding and attractor bump formation provide a framework for advances in decentralized swarm control for tasks that may require dynamic replanning within complex and changing environments.

ACKNOWLEDGMENTS

This work was supported by NSF NCS/FO award 1835279, “Spatial Intelligence for Swarms Based on Hippocampal Dynamics,” to K. Zhang (PI, JHU/School of Medicine) and G. Hwang & K. Schultz (PIs, JHU/Applied Physics Laboratory), who also received JHU/APL internal research and development support. The authors thank Robert W. Chalmers, Dwight Carr, Elise Buckley, Bryanna Yeh, and Clare Lau at the JHU/APL for insightful discussions and explorations of these ideas.

REFERENCES

- [1] Scoville, W. B. and Milner, B., “Loss of recent memory after bilateral hippocampal lesions,” *J Neurol Neurosurg Psychiatry* **20**(1), 11–21 (1957).
- [2] O’Keefe, J. and Dostrovsky, J., “The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat,” *Brain Res* **34**(1), 171–175 (1971).
- [3] O’Keefe, J. and Nadel, L., [*The Hippocampus as a Cognitive Map*], Clarendon Press, Oxford, UK (1978).
- [4] Knierim, J. J. and Zhang, K., “Attractor dynamics of spatially correlated neural activity in the limbic system,” *Annu Rev Neurosci* **35**, 267–85 (2012).
- [5] Buzsáki, G. and Moser, E. I., “Memory, navigation and theta rhythm in the hippocampal-entorhinal system,” *Nat Neurosci* **16**(2), 130 (2013).
- [6] Sosa, M., Gillespie, A. K., and Frank, L. M., [*Neural Activity Patterns Underlying Spatial Coding in the Hippocampus*], 43–100, Springer International Publishing, Cham (2018).

- [7] Murray, R. M., "Recent research in cooperative control of multivehicle systems," *Journal of Dynamic Systems, Measurement, and Control* **129**(5), 571–583 (2007).
- [8] Yang, G.-Z., Bellingham, J., Dupont, P. E., Fischer, P., Floridi, L., Full, R., Jacobstein, N., Kumar, V., McNutt, M., Merrifield, R., et al., "The grand challenges of Science Robotics," *Science Robotics* **3**(14), eaar7650 (2018).
- [9] Hamann, H., Khaluf, Y., Botev, J., Divband Soorati, M., Ferrante, E., Kosak, O., Montanier, J.-M., Mostaghim, S., Redpath, R., Timmis, J., et al., "Hybrid societies: challenges and perspectives in the design of collective behavior in self-organizing systems," *Frontiers in Robotics and AI* **3**, 14 (2016).
- [10] Chung, S.-J., Paranjape, A. A., Dames, P., Shen, S., and Kumar, V., "A survey on aerial swarm robotics," *IEEE Transactions on Robotics* **34**(4), 837–855 (2018).
- [11] Loianno, G., Brunner, C., McGrath, G., and Kumar, V., "Estimation, control, and planning for aggressive flight with a small quadrotor with a single camera and IMU," *IEEE Robotics and Automation Letters* **2**(2), 404–411 (2017).
- [12] Amaral, D. G., Ishizuka, N., and Claiborne, B., "Neurons, numbers and the hippocampal network," *Prog Brain Res* **83**, 1–11 (1990).
- [13] Sengupta, B., Laughlin, S. B., and Niven, J. E., "Consequences of converting graded to action potentials upon neural information coding and energy efficiency," *PLoS Comput Biol* **10**(1), e1003439 (2014).
- [14] Tchernichovski, O., Benjamini, Y., and Golani, I., "The dynamics of long-term exploration in the rat. Part I. A phase-plane analysis of the relationship between location and velocity," *Biol Cybern* **78**(6), 423–32 (1998).
- [15] Monaco, J. D., Rao, G., Roth, E. D., and Knierim, J. J., "Attentive scanning behavior drives one-trial potentiation of hippocampal place fields," *Nat Neurosci* **17**(5), 725–731 (2014).
- [16] Samsonovich, A. and McNaughton, B. L., "Path integration and cognitive mapping in a continuous attractor neural network model," *J Neurosci* **17**(15), 5900–5920 (1997).
- [17] Tsodyks, M., "Attractor neural network models of spatial maps in hippocampus," *Hippocampus* **9**(4), 481–9 (1999).
- [18] Zhang, K., "Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory," *J Neurosci* **16**(6), 2112–26 (1996).
- [19] Monaco, J. D., De Guzman, R. M., Blair, H. T., and Zhang, K., "Spatial synchronization codes from coupled rate-phase neurons," *PLoS Comput Biol* **15**(1), e1006741 (2019).
- [20] O'Keefe, J. and Recce, M. L., "Phase relationship between hippocampal place units and the EEG theta rhythm," *Hippocampus* **3**(3), 317–30 (1993).
- [21] Jensen, O. and Lisman, J. E., "Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding," *J Neurophysiol* **83**(5), 2602–9 (2000).
- [22] Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E. I., "Microstructure of a spatial map in the entorhinal cortex," *Nature* **436**(7052), 801–806 (2005).
- [23] Burgess, N., "Grid cells and theta as oscillatory interference: theory and predictions," *Hippocampus* **18**(12), 1157–74 (2008).
- [24] Blair, H. T., Gupta, K., and Zhang, K., "Conversion of a phase- to a rate-coded position signal by a three-stage model of theta cells, grid cells, and place cells," *Hippocampus* **18**(12), 1239–55 (2008).
- [25] Savelli, F. and Knierim, J. J., "Origin and role of path integration in the cognitive representations of the hippocampus: computational insights into open questions," *J Exp Biol* **222**(Pt Suppl 1) (2019).
- [26] Monaco, J. D., Knierim, J. J., and Zhang, K., "Sensory feedback, error correction, and remapping in a multiple oscillator model of place-cell activity," *Front Comput Neurosci* **5**, 39 (2011).
- [27] Issa, J. B. and Zhang, K., "Universal conditions for exact path integration in neural systems," *Proc Natl Acad Sci U S A* **109**(17), 6716–20 (2012).
- [28] Kirst, C., Timme, M., and Battaglia, D., "Dynamic information routing in complex networks," *Nature Communications* **7**, 11061 EP – (2016).
- [29] Varughese, J. C., Hornischer, H., Thenius, R., Zahadat, P., Wotawa, F., and Schmickl, T., "Introduction to WOSPP: wave oriented swarm programming paradigm," *CoRR* **abs/1804.04202** (2018).

- [30] Avena-Koenigsberger, A., Yan, X., Kolchinsky, A., van den Heuvel, M., Hagmann, P., and Sporns, O., “A spectrum of routing strategies for brain networks,” *PLOS Computational Biology* **15**(3), 1–24 (2019).
- [31] Bonabeau, E., Marco, D. d. R. D. F., Dorigo, M., Théraulaz, G., et al., [*Swarm intelligence: from natural to artificial systems*], no. 1, Oxford university press (1999).
- [32] Reynolds, C. W., “Flocks, herds and schools: A distributed behavioral model,” *SIGGRAPH Comput. Graph.* **21**(4), 25–34 (1987).
- [33] Saska, M., Vakula, J., and Přeucil, L., “Swarms of micro aerial vehicles stabilized under a visual relative localization,” in [*2014 IEEE International Conference on Robotics and Automation (ICRA)*], 3570–3575, IEEE (2014).
- [34] Ivancevic, V. G. and Reid, D. J., “Cognitive supervisor for an autonomous swarm of robots,” *Intelligent Control and Automation* **8**(01), 44 (2016).
- [35] Passino, K. M., Seeley, T. D., and Visscher, P. K., “Swarm cognition in honey bees,” *Behavioral Ecology and Sociobiology* **62**(3), 401–414 (2008).
- [36] Couzin, I. D., “Collective cognition in animal groups,” *Trends in Cognitive Sciences* **13**(1), 36–43 (2009).
- [37] Dehaene, S. and Changeux, J.-P., “Reward-dependent learning in neuronal networks for planning and decision making,” in [*Cognition, emotion and autonomic responses: The integrative role of the prefrontal cortex and limbic structures*], *Progress in Brain Research* **126**, 217 – 229, Elsevier (2000).
- [38] Papale, A. E., Zielinski, M. C., Frank, L. M., Jadhav, S. P., and Redish, A. D., “Interplay between hippocampal sharp-wave-ripple events and vicarious trial and error behaviors in decision making,” *Neuron* **92**(5), 975–982 (2016).
- [39] Momennejad, I., Otto, A. R., Daw, N. D., and Norman, K. A., “Offline replay supports planning in human reinforcement learning,” *eLife* **7**, e32548 (2018).
- [40] Vikbladh, O. M., Meager, M. R., King, J., Blackmon, K., Devinsky, O., Shohamy, D., Burgess, N., and Daw, N. D., “Hippocampal contributions to model-based planning and spatial memory,” *Neuron* [**Epub ahead of print**] (2019).
- [41] Gutkin, B. S. and Ermentrout, G. B., “Dynamics of membrane excitability determine interspike interval variability: a link between spike generation mechanisms and cortical spike train statistics,” *Neural Comput* **10**(5), 1047–65 (1998).
- [42] O’Keefe, K. P., Hong, H., and Strogatz, S. H., “Oscillators that sync and swarm,” *Nature communications* **8**(1), 1504 (2017).
- [43] Wang, Y., Roth, Z., and Pastalkova, E., “Synchronized excitability in a network enables generation of internal neuronal sequences,” *Elife* **5** (2016).
- [44] Pfeiffer, B. E. and Foster, D. J., “Hippocampal place-cell sequences depict future paths to remembered goals,” *Nature* **497**(7447), 74–9 (2013).
- [45] Ólafsdóttir, H. F., Bush, D., and Barry, C., “The role of hippocampal replay in memory and planning,” *Curr Biol* **28**(1), R37–R50 (2018).
- [46] Hebb, D. O., [*The Organization of Behavior: A Neuropsychological Theory*], Wiley and Sons, New York (1949).
- [47] Levy, W. B. and Steward, O., “Synapses as associative memory elements in the hippocampal formation,” *Brain Res* **175**(2), 233–45 (1979).
- [48] Eichenbaum, H., “Barlow versus Hebb: When is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition?,” *Neuroscience Letters* **680**, 88–93 (2018). New Perspectives on the Hippocampus and Memory.
- [49] Oja, E., “Simplified neuron model as a principal component analyzer,” *Journal of Mathematical Biology* **15**(3), 267–273 (1982).
- [50] Monaco, J. D. and Levy, W. B., “T-maze training of a recurrent CA3 model reveals the necessity of novelty-based modulation of LTP in hippocampal region CA3,” in [*Proc Int Joint Conf Neural Netw*], 1655–1660, IEEE (2003).
- [51] Monaco, J. D. and Abbott, L. F., “Modular realignment of entorhinal grid cell activity as a basis for hippocampal remapping,” *J Neurosci* **31**(25), 9414–25 (2011).
- [52] Salahshour, M., Rouhani, S., and Roudi, Y., “Phase transitions and asymmetry between signal comprehension and production in biological communication,” *Scientific Reports* **9**(1), 3428 (2019).