

Ergodic dynamics for large-scale distributed robot systems

Dylan A. Shell and Maja J. Matarić

Department of Computer Science
University of Southern California
Los Angeles, CA USA 90089-0781
{shell | mataric}@usc.edu

Abstract. Intelligent autonomous robotics is a promising area with many potential applications that could benefit from non-traditional models of computation. Information processing systems interfaced with the real world must deal with a continuous and uncertain environment, and must cope with interactions across a range of time-scales. Robotics problems resist existing tools and, consequently, new perspectives are needed to address these challenges. Toward that end, we describe a dynamics-based model for computing in large-scale distributed robot systems. The proposed method employs a compositional approach, constructing robot controllers from ergodic processes. We describe application of the method to two multi-robot tasks: decentralised task allocation, and collective strategy selection.

1 Introduction

Intelligent autonomous robotics is the study of autonomous agents coupled with the physical world. Robots are equipped with sensors to perceive their environment and actuators that allow aspects of the agent, and often parts of the environment, to be controlled. Even robots with limited information processing capabilities can exhibit complex, nontrivial behaviour due to the feedback introduced between the robot and its environment. Systems like autonomous robots, that interface with the world, introduce several unique challenges. We highlight those challenges that may be best addressed by non-traditional computing ideas, illustrating that robotics is a potentially rich application area for broad, alternative notions of computation.

This paper also considers the specific problem of producing predictable task-oriented collective behaviours in systems of many simple robots, or so-called *swarm* robot systems. The word “swarm” is a reference to those natural systems which inspire the research. Several biological systems are capable of operating in distributed and decentralised ways, exploiting synergism of simple individuals, and achieving robustness through massive redundancy. We wish to construct artificial systems with similar features, which is the subject of ongoing research.

We describe a method for programming swarm robot systems applicable to a range of tasks domains. The method explores a dynamics-based formalism primarily using the composition of elementary processes, each possessing the ergodic

property. Statistical mechanics techniques allow a link to be established between individual (microscopic) interactions and the group (macroscopic) behaviour. The result is an approach for structuring collective interactions in loosely-coupled distributed systems, which emphasises equilibrium states rather than the steps necessary to achieve particular operations.

In order to demonstrate our method, we consider two complementary variations on an entomologically-inspired multi-robot foraging task. We show that our method is sufficient to enable group coordination in both tasks. This also serves to demonstrate the types of capabilities of interest within the multi-robot community. We present data from large-scale simulations.

2 Robotics and its challenges

Researchers from several disciplines have shown an interest in synthesising, analysing and studying autonomous system behaviour. The cybernetics community, with its basis in control theoretic ideas, produced early pioneering systems using analogue electronic methods (e.g., [29]). Later, some within the AI community assumed the goal of constructing artificial beings capable of exhibiting intelligent behaviour. Bekey [4] recently produced a comprehensive history and discussion of the field.

2.1 Intelligent robotics

For several years the AI community attempted to apply search and planning techniques to robotics problems. In the late 1980's arguments against the traditional AI methods were put forward by Brooks and his collaborators (see for example Brooks [7]). The arguments highlighted assumptions which ignored key aspects of robot systems, such as the implicit belief in the physical symbol hypothesis [23], and an assumption that symbolic search could become tractable.

Robots inhabit a continuous world. Traditional planning methods however, require discretization. Any robot relying on a discrete representation for its successful functioning may become brittle, with small errors causing the robot's beliefs about the world to diverge from reality. The result is execution of inappropriate actions. Some viewed this as a fundamental shortcoming of symbolic representations. Subsequent researchers explored alternatives, like connectionist methods [22], while others have introduced randomised algorithms for planning over continuous spaces [18], or highly unorthodox methods [1]. Currently, probabilistic techniques for explicitly representing and reasoning about the world are the most popular [25]. However, the world is generally a dynamic place and few explicit representations are able to provide machinery for updating models over time.

The physical world is fraught with uncertainty. In addition to the noisy sensors and inaccurate actuators, no robot can reasonably expect to have complete information about the environment in which it finds itself. Uncertainty presents a challenge to the robot as well as the researcher intending to model the world.

Many of the methods for designing industrial robots are inappropriate when robots are expected to operate in unstructured environments. For example, an assumption of known friction constants becomes tenuous.

Robotic systems operate across disparate time-scales. The physical dynamics, hardware interfaces, and high-level computation typically each execute with frequencies that are orders of magnitudes apart. The task specification itself may have non-trivial timing requirements. In all but the most benign environments, robots will have real-time constraints in order to ensure survival and maintain safety through obstacle avoidance. Often the robot has multiple high-level goals (typically several will be time-extended) that must be reconciled with the lower-level constraints.

The choice of computing model dictates the way one approaches any particular problem. Brooks [8] has argued that a dominant computing model has effected thinking about intelligence itself. A re-examination in this context has changed the perception of robots and their role in robot-environment interactions. However, beyond a few biological inspirations and proposals for novel controller architectures, a comprehensive re-examination of alternative computing paradigms is yet to come to robotics.

2.2 Multi-robot systems

There are numerous task domains in which multiple robots operating concurrently offer advantages over a single robot. A task that is impossible for an individual may be feasible for a group. Also, additional robots may improve performance—although typically only up to a point. After that, additional robots aggravate resource conflicts, thus making sophisticated allocation and management policies necessary. Unlike traditional distributed systems, communication cannot always be assumed. Explicit wireless communication may be ephemeral. Further, the notion of implicit communication, through a shared environment, rather than through a dedicated communication channel, is particularly relevant in multi-robot systems.

Swarm systems follow biological inspirations. They are usually homogeneous groups of simple robots, which are individually *minimalist* [10] in that they are an attempt to find the smallest set of capabilities necessary to achieve a particular task. Frequently, implicit communication is the only form of interaction among the robots. Examples of implicitly coordinated systems include instances of puck clustering and sorting. Simulations have been used to hypothesise about the sensing and computational requirements on individual ants [9]; experiments with real robots showed that fewer sensing capabilities are required from individuals because physical dynamics can aid the clustering process [3]. Further robot experiments showed that overall performance depends critically on several physical features of the experiment and the robots themselves [14]. Minimalist cooperative box-pushing [17], again inspired by social ants, has been demonstrated on physical robots. In those results, positive feedback has led to self-organised task-achievement [6] of the group as a whole. Specific mechanisms from nature

have also been generalised, such as the use of pheromone trails, which have been employed in robotics [21, 26].

In stark contrast to swarms are multi-robot systems employing explicit coordination methods [16]. These systems use explicit communicative acts generated by a distributed algorithm executing as a layer within an otherwise standard single-robot controller. Such designs may target heterogeneous robot systems, often with fewer robots (≤ 10) than contained in swarms. The physical properties of the robots are not directly considered, whereas in swarm systems, robot physics may be critical for task-achievement. Programming explicitly coordinated systems is algorithmic, while swarms are often better treated as dynamic systems.

When considering a large swarm of robots, in addition to distinct time-scales, the system can be characterised at multiple spatial scales. Local microscopic interactions occur between individual robots. Since the individual robots act autonomously, programming must ultimately be grounded at the microscopic scale. Descriptions of the properties of the complete system (or large parts, relative to the radius of communication) are considered the macroscopic level-of-detail. Collective properties are exposed at this level by considering average system behaviour. Structured local dynamics can produce complex global phenomena. Simulated distributed systems of simple interacting agents have exhibited global behaviour ranging from point and periodic attractors to chaos [13] and, of course, there are classic examples of those capable of universal computation [5].

Challenges in dealing with robot systems, and other systems that interface computation with the physical world, have a large number of potential applications. Visions of a future with ubiquitous computing [28] are far more likely to be realised if we have an appropriate computational model. We believe that such a model would tolerate random failures, allow multiple levels-of-description, and be capable of dealing with approximate and incomplete information. The key question is whether such a model exists, and if so, whether the model is sufficiently expressive. A model which satisfies the requirements, but which sacrifices the ability to perform task-oriented computation, is of little use.

3 Large-scale multi-robot systems

Next, we consider the problem of prescribing a control and communication policy for homogeneous large-scale systems with hundreds of robots, i.e., robot swarms. Such systems are already conceivable within a research setting, but several engineering issues must be overcome in order for such systems to become common. Hundreds of robots would be demanding for most existing methods, as few researchers evaluate scaling properties for more than twenty robots. It is at these numbers of robots for which the method we propose here begins to become feasible.

The method achieves distributed computation through the interactions of coupled processes with the ergodic property. This places a condition on the tem-

poral structure of the process dynamics so that all accessible regions of the processes phase-space can be explored (see pg. 6 for the formal definition). The ergodic property permits programming to occur at an elevated level-of-description. The underlying philosophy is that microscopic details are not always necessary for controller construction.

3.1 Definitions

We take a *process* P to be a tuple comprising a state-space definition S and a dynamics function Φ , writing $P = \langle S, \Phi \rangle$. The set S need not be finite (nor countable), and represents all possible states a process could occur in. Also Φ is a generally non-deterministic dynamics that produces the process trajectory $s(t) \in S$ for all non-negative $t \in \mathbb{Z}^+$. Suppose such a trajectory is produced by each of the n robots executing process P . A global snap-shot is given by the state vector $\mathbf{s}(t) = [s_1(t), s_2(t), \dots, s_n(t)]$. The *microstate* vector $\mathbf{s}(t)$ captures the complete microscopic details of process P across all of the n robots for time t . The system-wide evolution can be interpreted in terms of vectorised dynamics function Φ constructed by n copies of Φ operating on each component in $\mathbf{s}(t)$ to yield a $\mathbf{s}(t + \delta t)$.

We consider a restriction to the general $\Phi : \mathcal{S} \rightarrow \mathcal{S}$ so that the i^{th} element may depend only on nearby robots, that is, robots within a given communication disk. Each robot requires only local information in order to update its internal state for each process, and so all the dynamics functions we consider will have this form.

An arbitrary function $\mathbf{G} : \mathcal{S} \rightarrow \mathbb{R}$ produces an associated equivalence relation $\mathbf{s} \underset{\mathbf{G}}{\sim} \mathbf{r} \iff \mathbf{G}(\mathbf{s}) = \mathbf{G}(\mathbf{r})$ and hence partitions \mathcal{S} into equi- \mathbf{G} -valued equivalence classes. Each class is called a *macrostate*. Any suitable \mathbf{G} gives a macroscopic view of the system in which any two microstates which are equivalent with respect to $\underset{\mathbf{G}}{\sim}$ are identified. A convenient interpretation is that any two such states appear indistinguishable to an observer capable only of measuring \mathbf{G} . One question frequently asked is the relative number of states with $\mathbf{G}(\cdot) = C_1$ and $\mathbf{G}(\cdot) = C_2$. The set of microstates can be compared through the *entropy*, calculated—following Boltzmann—as a log function of the set cardinality $\mathcal{S}(X) \triangleq \ln ||X||$.

3.2 Ergodicity

In a many degree-of-freedom dynamic system two factors direct the overall system behaviour. First, the phase-space \mathcal{S} , that is, the conceivable states available to the system. Second, the spatio-temporal dynamics Φ through which the system explores those states. In designing a system expected to perform information processing, the vast majority of approaches (and thinking) is directed toward constructing sophisticated spatio-temporal dynamics. As already mentioned, even simple interaction rules can produce complex collective dynamics.

These complexities make the selection of rules in order to perform a given task extremely difficult. Prediction of the system may require exact initial conditions, perfect models, or may even be undecidable.

We propose that the dynamics Φ be chosen so as to enable prediction. Task-oriented computation is performed through changes to the structure of \mathcal{S} . Next, we define what is meant by “enable prediction.”

The average macrostate \mathbf{G} likely to be externally observed from time 0 to T is given by

$$\overline{\mathbf{G}}_T \triangleq \frac{1}{T} \int_0^T \mathbf{G}(\Phi^t(\mathbf{s}(0))) dt, \quad (1)$$

where $\Phi^t(\mathbf{s}(0)) = \overbrace{\Phi(\Phi(\dots\Phi(\mathbf{s}(0))))}^{t \text{ times}}$ generates a single trajectory from initial condition $\mathbf{s}(0)$.

Dynamics possessing the ergodicity property induce a probability measure on the phase-space. More formally a system that exhibits *ergodic dynamics* has the properties that:

1. A measure $\mathbf{P} : \mathcal{S} \rightarrow (0, 1]$ exists with $\int_{\mathcal{S}} \mathbf{P}(\mathbf{s}) d\mathbf{s} = 1$
2. With any initial conditions the system will evolve exploring \mathcal{S} entirely given sufficient time. The probability of finding the system in states $\mathcal{B} \subseteq \mathcal{S}$ is given by the probability mass of \mathbf{P} that lies within the sub-volume \mathcal{B} of the phase-space. That is $\Pr(\mathbf{s}(t) \in \mathcal{B}) = \int_{\mathcal{B}} \mathbf{P}(\mathbf{r}) d\mathbf{r}$.

The key here is that the current microstate places very “loose” restrictions on future states, the dynamics is free to explore the state space. Long-term history and initial conditions are not important in predicting the system’s trajectory. Dynamics functions typically produce short-term temporal regularity. In such cases, analysis must consider durations of sufficient length.

Measure \mathbf{P} allows for the phase-space average of \mathbf{G} defined as

$$\langle \mathbf{G} \rangle \triangleq \int_{\mathcal{S}} \mathbf{P}(\mathbf{s}) \mathbf{G}(\mathbf{s}) d\mathbf{s}. \quad (2)$$

For ergodic processes long time-averages equal phase-space averages,

$$\lim_{T \rightarrow \infty} \overline{\mathbf{G}}_T = \langle \mathbf{G} \rangle. \quad (3)$$

In other words the macroscopic behaviour of such processes can be described by modelling the phase-space rather than resorting to simulation of the dynamics. The value of the two means identified in Eqn. 3 is the called equilibrium value of \mathbf{G} .

3.3 Coupling processes

The definitions thus far describe a method for predicting the mean macrostate of a process executing on an homogeneous system of robots. They also show the probabilistic nature of the law of increasing entropy. But in order to perform any useful processing, multiple connected processes must be considered. First, we must generalise the notion of a process by permitting parametrisation. Let $P(m) = \langle S(m), \Phi(m) \rangle$ denote one such process in which the phase-space and dynamics are variable and each depend on m . As already mentioned, it is the phase-space which we are most interested in.

Suppose the robots are executing two parametrised ergodic processes: $P^0(m) = \langle S^0(m), \Phi^0 \rangle$, $P^1(n) = \langle S^1(n), \Phi^1 \rangle$. Two state vectors $\mathbf{s}^1(t)$ and $\mathbf{s}^2(t)$, one for each of the processes, gives the complete state of the system. Eqn. 2 can be applied to each independent parametrisation (e.g., $n = n_0, m = m_0$) to calculate expected behaviour.

The parametrised processes are coupled by defining a relation on the parameters. The previous processes can be constrained so that $m + n = C$. A known C implies that a single degree-of-freedom describes the parametrisation. We call this a *macroscopic degree-of-freedom*. A dynamics can be defined that operates on parameters m and n that respects the conservation constraint. If this coupling dynamics is slow compared with P^0 and P^1 then the two processes can be suitably modelled as a single composite process. This composite process can be predicted as before, provided behaviour is analysed on long time-scales compared to the coupling dynamics. The constraint relation structures a composite phase-space from P^0 and P^1 .

Clearly this procedure be applied repeatedly and recursively, while time-scales remain suitable.

4 Example problem domains

Foraging is a canonical and one of the most widely studied tasks in distributed robotics [2, 12, 20]. It is entomologically-inspired and requires robots to locate items (called pucks) scattered throughout an environment, and transport them back to a central location (called the home region).

4.1 Domain 1: task allocation

We consider a variation of foraging in which two varieties of puck exist, call them type A and type B . Like Jones and Matarić [15], we consider the case in which each robot may forage only one variety of puck at a time. The task allocation problem involves switching robots from one variety to another in order to emulate the fractional distribution of pucks within the environment. We measure the effectiveness of the allocation by comparing the distribution of pucks within the environment with the proportion of the robots foraging each type.

The robot controllers include traditional behaviours like obstacle avoidance, basic searching, homing, etc. Two ergodic processes are layered above these

behaviours in order to provide each robot with a local strategy for choosing the variety of puck to forage.

For n robots, define $P_a(n, m_a) = \langle \{0, \dots, m_a\}, \Phi_a \rangle$ where

$$\begin{aligned} \Phi_a [s_k(t), s^{I \in A_k}(t)] &= s_k(t) - \sum_{i \in A_k} \omega_{k,i}(s_k(t), s_i(t)) \\ &\quad + \sum_{j \in A_k} \omega_{j,k}(s_j(t), s_k(t)) \end{aligned}$$

where $\omega_{k,i}(s_k(t), s_i(t))$ for $i \in A_k$ are uniformly drawn integer random variables constrained so that $\sum_{i \in A_k} \omega_{k,i}(s_k(t), s_i(t)) \leq \alpha s_k(t)$, ($\alpha = 10^{-2}$). The set A_k gives the indices of robots within communication distance of robot k (which technically depend on t). We require symmetrical neighbourhood sets so that $i \in A_k \implies k \in A_i$. The dynamics rule says that robot i distributes random portions of its current state value among its neighbours. Despite using only local communication, this interaction rule conserves the global property, i.e., $\forall t, \sum_i^n s_i(t) = K$. We set initial conditions so that $K = m_a$.

This process has two defining parameters: n and m_a . In the case where $n \ll m_a$, the total number of states is a simple combinatorial exercise,

$$\|\mathcal{S}(n, m_a)\| = \binom{m_a + n - 1}{n - 1}.$$

Define $\mathbf{G}_j([s^1, \dots, s^n]) = s^j$. To get the average state for robot j , we apply Eqn. 2. Formally proving that Φ_a is ergodic is beyond the scope of this paper, but observe that the dynamics is symmetrical with respect to robots: no single robot is favoured over another. The dynamics randomly explores the configuration available to it. The full density function is unnecessary to calculate $\langle \mathbf{G}_j \rangle = \frac{m_a}{n}$.

A second process $P_b(n, m_b)$ is defined identically. Dynamics Φ_a and Φ_b explore their respective state spaces with expected state calculable in terms of n , m_a and m_b . The system is initialised with $m_a = 10^5$, $m_b = 10^5$. These m_a and m_b values are adjusted on each puck observation when the robot alters the local state of the processes $s^a(t)$ and $s^b(t)$. Observation of an A puck causes the following transition: $s^a(t + \delta t) = (1 - \gamma)s^a(t)$ and $s^b(t + \delta t) = \gamma s^a(t) + s^b(t)$ where γ is a tunable parameter. The converse happens on observing a B puck. Each robot independently decides which type of puck to forage using $s^a(t)$ and $s^b(t)$, the local states of the two processes. Pucks of variety A are chosen with probability $pr_{a-puck} = s^a(t)/(s^a(t) + s^b(t))$. Thus, the transitions simply skew the probability by a factor of γ .

Robots randomly encounter either type A or B pucks, making observations of each type in proportion with the puck distribution. These observations are smoothed by the dynamics of the two processes. Low probability observations (e.g., observing ten types of minority puck) are averaged out over the entire group of robots.

We simulated 100 robots within a 64m \times 64m arena. Initially, 3000 pucks were randomly scattered throughout the arena. Pucks started in an initial 50%/50%




Fig. 1. Performance of task allocation processes. The vertical axis gives the proportion of tasks (for the broken line), and the division of robots among tasks (the solid lines). Plots show mean and standard deviation for 5 runs.

distribution. Robots explored from random initial locations. After stumbling on an appropriate puck, the robot would transport it to the home region. For each puck foraged, a new one, of the same type, was introduced at a random location. Thus the puck density was maintained throughout.

The puck distribution was altered at three stages, at $t = 2000$ it was changed to 95%/5%, at $t = 8000$ to 5%/95%, and at $t = 13000$ to 75%/25%. In Figure 1, the dotted line shows the puck distribution. The plot shows experimental runs with three different settings for the γ parameter. The system shows hysteresis and a response time dependant on γ . In all cases, however, the system adapts so as to find a distribution applicable for the environmental conditions.

4.2 Domain 2: collective strategy selection

Interference can adversely affect task performance in a multi-robot system. It is a particularly important factor in large-scale robot systems. Robot foraging is one of few problem domains for which interference has been well-studied (e.g., Goldberg [12]) and is suitable for large-scale systems. Efficient strategy selection can ameliorate the negative effects inter-robot interference. More generally, however, this form of distributed decision-making underlies many tasks.

In this domain we consider two distinct strategies for foraging (with a single puck type). The first, *homogeneous method*, involves each robot searching for a puck, locating a puck, and delivering that puck home [12]. In the alternative, *bucket brigading method* [19], each robot deals with only those pucks in a particular sub-region of the environment (also termed the caste method [12]). Figure 2 shows performance data for simulations with 250 robots within a $25\text{m} \times 25\text{m}$ arena. The left graphs show data for 500 pucks. With that puck density the homogeneous strategy out performs bucket brigading. But with higher puck density (right graph is for 2000 pucks) the homogeneous strategy is inferior to bucket brigading. The homogeneous strategy results in considerable interference around the home region with high puck densities because a large number of pucks results in a short mean time between puck discoveries. Bucket brigading becomes more effective with increasing puck density because the robots have no perception of pucks at a distance. On the other hand, with too few pucks each robots takes a long time to find pucks dropped by its peers.




Fig. 2. Plot of the number of pucks within home region versus time for both foraging strategies (with 250 robots). Left figure has low puck density, right has high puck density. Plots show mean and standard deviation for 5 independent simulation runs.

Each robot’s controller consists of an implementation of both foraging strategies, with a binary variable to control the one currently in use. This variable provides an interface to the two coupled ergodic processes. The first is the process defined in the preceding section; call it $P_a(n, m_a)$. The second is defined as $P_c(n, m_c, e_c) = \langle \{-1, +1\}, \Phi_c \rangle$. Again that is for n robots. Here $m_c = \sum_{i=1}^n s_k(t)$. Less obviously, $e_c = \sum_{j,k, s.t. j \in A_k} -s_j(t)s_k(t)$ measures the number of neighbours that have the same state. The value of m_c gives a measure of “agreement” among robots, while e_c a measure of “frustration.” A range of m_c values are applicable for a given e_c .

The dynamics Φ_c is constructed so that e_c remains constant. The process has only two states and the transition from one state to another (e.g., $+1$ to -1) is called a flip. Two robots within communication range, say i and j , calculate δe_i^c and δe_j^c , these are the changes in e_c that would result if robot i and j flip states. Both can be calculated using local information. If $\delta e_i^c = -\delta e_j^c$ then the two robots carry out this flip operation.

Both P_c and P_a execute a coupling dynamics that operates on the e_c and m_a parameters. The process on robot i may flip with a resulting δe_i^c , provided that this can be supplanted by subtracting an equivalent value from P_a . Thus a large m_a value can result in an increased e_c and hence effect the robot’s “agreement” and, similarly, small m_a results in decreased e_c . The manner by which this change occurs is crucial for strategy selection.

Analysis of P_c ’s macroscopic behaviour is less obvious than P_a . The process has the same structure as the Ising ferromagnetic model [11], which has been well studied in the limit of infinite system size, and under controlled temperature conditions. In the limit the model exhibits a symmetry breaking phase-transition from mixed spin values (with $m_c \sim 0$) to a state with alignment ($m_c = 1$ or $m_c = -1$).

During the foraging task, local sensing of task progress enables the m_a value to be tuned appropriately. While bucket brigading, each robot kept track of the time between puck discoveries, a noisy local measurement of puck density.

For homogeneous foraging, interference was estimated by measuring distance travelled for a short time. Both of these decreased the local value of $s^a(t)$ on each robot. Each time a puck was dropped over the home region $s^a(t)$ was decreased. The system achieves a steady-state between m_a creation (through interference, obstacle avoidance) and m_a deletion (from pucks homed). Different puck densities have different steady-states, and a mismatch of puck density and strategy is sufficient to drive the phase transition of P_c .

5 Discussion

Statistical mechanics methods enable predictions of behaviour by characterising macroscopically identical systems. These *ensemble* predictions are feasible in systems with hundreds of robots, and with increasing system size predictions become progressively accurate. We believe that, for large systems, existing metaphors — e.g., message-passing — break down and must be superseded. We thus explore and exploit the benefits afforded by the large number of degrees-of-freedom, in which we include both aspects of the physical robot and controller state-space.

The two foraging domains were carefully chosen. The “decentralised task allocation” domain required estimation of a continuous quantity. It is typical of the calculations and optimisation that might be performed by an implicitly coordinated swarm system. Small local estimates can be used to make a decision, and that can be shared with other agents easily. This notion seems similar to the Downhill Principle [27], but in a distributed sense.

On the other hand, explicitly coordinated systems often solve discrete problems with hard constraints. For such problems, gradient techniques are not useful. The “collective strategy selection” was intended to demonstrate that discrete notions can also be feasibly tackled with the ergodic process approach. Since communication times are not zero, the system does take time to transition between states. The transition is all-or-nothing, with infinite size.

An unconventional aspect of the proposed methodology is the focus on equilibrium solutions rather than dynamics in computing. In a sense this is not unlike the electronic analogue computers of the past, in which the initial transients were ignored, with attention paid to steady-state solutions [24]. In robot systems, it is envisioned that changes within environmental factors would trigger adaptation within the ergodic processes, as they reach a new equilibrium.

6 Conclusion

This paper has argued that non-conventional computing models may be a way to elegantly address the challenges raised by physically situated robots. We defined our own compositional method for synthesising controllers for large-scale multi-robot systems and proposed the use of ergodic processes as elemental distributed building blocks. This contrasts with current methods that produce task-oriented

behaviour through dynamics with rich temporal structure. Two complementary coordinated foraging problems were used to demonstrate the proposed method.

Acknowledgements Support for this work was provided by the National Science Foundation Grant No. IIS-0413321. We thank G. Sibley for proofreading.

Bibliography

- [1] A. Adamatzky, B. De Lacy Costello, C. Melhuish, and N. Ratcliffe. Experimental reaction–diffusion chemical processors for robot path planning. *Journal of Intelligent Robotic Systems*, 37:233–249, 2003.
- [2] R.C. Arkin, T. Balch, and E. Nitz. Communication of behavioral state in multi-agent retrieval tasks. In *Proceedings IEEE Conference on Robotics and Automation*, pages 588–594, Atlanta, GA, 1993.
- [3] R. Beckers, O. E. Holland, and J.-L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In *Artificial Life IV.*, pages 181–189, Cambridge, MA, Jul 1994.
- [4] G. A. Bekey. *Autonomous Robots: From Biological Inspiration to Implementation and Control*. MIT Press, Cambridge, MA, 2005.
- [5] E. R. Berlekamp, J. H. Conway, and R. K. Guy. *Winning ways for your mathematical plays*. Academic Press, New York, 1982.
- [6] E. Bonabeau, M. Dorigo, and G. Thraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NY, 1999.
- [7] R. A. Brooks. *Cambrian Intelligence: The Early History of the New AI*. MIT Press, Cambridge, MA, 2001.
- [8] R. A. Brooks. Intelligence without reason. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 569–595, Sydney, Australia, 1991. Morgan Kaufmann publishers Inc.: San Mateo, CA. Reprinted in [7].
- [9] J.-L. Deneubourg, S. Goss, N. R. Franks, A. B. Sendova-Franks, C. Detrain, and L. Chrétien. The dynamics of collective sorting robot-like ants and ant-like robots. In *Proc. First International Conference on Simulation of Adaptive Behavior*, pages 356–363, Paris, France, September 1990.
- [10] B. Donald, J. Jennings, and D. Rus. Minimalism + Distribution = Supermodularity. *Journal of Experimental and Theoretical AI*, 9(20):293–321, 1997.
- [11] M. E. Fisher. The theory of equilibrium critical phenomena. *Reports on Progress in Physics*, 30:615–730, 1967.
- [12] D. Goldberg. *Evaluating the Dynamics of Agent-Environment Interaction*. PhD thesis, Department of Computer Science, University of Southern California, May 2001.
- [13] T. Hogg and B. A. Huberman. Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics (Special Section on DAI)*, 21(6):1325–1332, 1991.
- [14] O. E. Holland and C. Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5(2):173–202, 1999.
- [15] C. V. Jones and M. J. Matarić. Adaptive division of labor in large-scale minimalist multi-robot systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1969–1974, Las Vegas, NV, Oct 2003.

- [16] C. V. Jones, D. A. Shell, M. J. Matarić, and B. P. Gerkey. Principled Approaches to the Design of Multi-Robot Systems. In *Invited contribution to Workshop on Networked Robotics, International Conference on Intelligent Robots and Systems*, pages 71–80, Sendai, Japan, September 2004.
- [17] C. R. Kube and H. Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–219, 1993.
- [18] J. C. Latombe. *Robot motion planning*. Kluwer Academic Publishers, Norwell, MA, 1991.
- [19] E. Østergaard, G. S. Sukhatme, and M. J. Matarić. Emergent bucket brigading — a simple mechanism for improving performance in multi-robot constrained-space foraging tasks. In *International Conference on Autonomous Agents*, pages 29–30, Montreal, Canada, May 2001.
- [20] L. E. Parker. ALLIANCE: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, Apr 1998.
- [21] D. Payton, M. Daily, R. Estkowsky, M. Howard, and C. Lee. Pheromone robotics. *Autonomous Robots*, 11(3):319–324, November 2001.
- [22] D. Pomerleau. *Neural network perception for mobile robot guidance*. Kluwer Academic Publishers, Norwell, MA, 1993.
- [23] H. A. Simon. *The Sciences of the Artificial*. MIT Press, Cambridge, Massachusetts, third edition, 1996.
- [24] J. S. Small. *The Analogue Alternative: The Electric Analogue Computer in Britain and the USA, 1930-1975*. Routledge, London and New York, 2001.
- [25] S. Thrun, W. Burgard, and Dieter Fox. *Probabilistic Robotics*. MIT Press, Cambridge, MA, 2005.
- [26] R. T. Vaughan, K. Støy, G. S. Sukhatme, and M. J. Matarić. Blazing a trail: insect-inspired resource transportation by a robot team. In *Proceedings of the International Symposium on Distributed Autonomous Robotic Systems*, pages 111–120, Knoxville, TN, Oct 2000.
- [27] A. Vergis, K. Steiglitz, and B. Dickinson. The complexity of analog computation. *Mathematics and Computers in Simulation*, 28:91–113, 1986.
- [28] M. Weiser. Some Computer Science issues in Ubiquitous Computing. *Communications of the ACM*, 36(7):75–84, 1993.
- [29] N. Wiener. *Cybernetics, or Control and Communication in the animal and the machine*. MIT Press, Cambridge MA, second edition, 1962.