

ERGODIC DYNAMICS BY DESIGN: A ROUTE TO PREDICTABLE MULTI- ROBOT SYSTEMS

Dylan A. Shell, Chris V. Jones, Maja J. Matarić
Computer Science Department, University of Southern California
Los Angeles, California 90089-0781, USA
dshell@cs.usc.edu, cvjones@cs.usc.edu, mataric@cs.usc.edu

Abstract We define and discuss a class of multi-robot systems possessing ergodic dynamics and show that they are realizable on physical hardware and useful for a variety of tasks while being amenable to analysis. We describe robot controllers synthesized to possess these dynamics and also physics-based methodologies that allow macroscopic structures to be uncovered and exploited for task execution in systems with large numbers of robots.

Keywords: Multi-robot systems, Ergodicity, Formal methods.

1. Introduction

Multi-robot systems can both enhance and expand the capabilities of single robots, but robots must act in a coordinated manner. So far, examples of coordinated robot systems have comprised of largely domain-specific solutions, with few notable exceptions. In this paper we describe our ongoing work on the development of formal methodologies for synthesis of multi-robot systems that address these issues in a principled fashion.

We focus here on inter-robot dynamics, the roles played by those dynamics toward task achievement, and their implications in feasible formal methods for synthesis and analysis. We describe automated synthesis of controllers that capitalize on so-called *ergodic* dynamics, which enable mathematical arguments about system behavior to be simplified considerably. Sensor-based simulations and physical robot implementations show that these controllers to be feasible for real systems. We further suggest that this approach will scale to systems with large numbers of robots.

2. Related formal methodologies

Formal methodologies for synthesis and analysis of multi-robot systems differ based on the type of systems they aim to address. One successful method




Figure 1 The Markovian *world* as states with action transitions. The sequential *task* specification is a particular sequence of world transitions. The task is shown here shaded within the world specification.

for analysis of swarm systems is based on the theory of stochastic processes: for example, in the phenomenological modeling and analysis of multi-robot cooperative stick-pulling, a macroscopic difference equation for the rates of change of each type of robot state is derived from the stochastic master equation and sensor-based simulations are used to estimate parameter values (Martinoli et al., 2004). An extension to the same theory (but using continuous differential equations instead) allows adaptive systems to be modeled (Lerman and Galstyan, 2003). When applied to foraging, the analysis enabled system design improvements. Explicitly coordinated systems are typically addressed at the algorithmic level, such as in the Computation and Control Language (Klavins, 2003) and formal studies of multi-robot task allocation methodologies (Gerkey and Matarić, 2004). Also related is Donald’s (1995) Information Invariants Theory, and Erdmann’s (1989) studies of the advantages of probabilistic controllers.

3. Behavioral configuration space and ergodicity

The physical configuration space, common in robotics for representing physical arrangements, can be augmented to include additional dimensions for each of the robot’s internal control variables that observable behavior. We call this the *behavioral configuration space* (BCS). It is a useful mental representation for a multi-robot system and for reasoning about the overall system dynamics. For practical applications, we will only consider particular subspaces, never the full configuration space.

The BCS of a single robot consists of dimensions for the physical configuration (e.g., the pose variables, and velocities if necessary) and dimensions for internal state variables (continuous or discrete values within memory). The range of each dimension is determined by constraints on state variables. The BCS of an ensemble of robots is constructed from essentially a Cartesian product of individuals spaces and the spaces of movable obstacles, etc. The constraints (e.g., two robots simultaneously occupying the same location) subtract components from this product. Couplings between the robots via communication channels, mutual observation, etc., further restrict this space.

The global state of the multi-robot system at any specific time can be represented by a point in BCS and likewise, the time-evolution of the system, as a trajectory. A system that exhibits *ergodic dynamics* completely visits all parts of the configuration space with probability that is dependent only on the volume of that part of the space. Long term history is unimportant in predict-

ing the dynamical behavior because the system “forgets” previous trajectories. Time averages of some system property (over a duration longer than the underlying dynamics timescale), are equal to (configuration) spatial averages. Few useful robotics systems are entirely ergodic, but various sub-parts of the BCS may be ergodic. The next section describes one such system.

4. Automated synthesis for sequential tasks

Jones and Matarić (2004a, 2004b) have developed a framework for automatic and systematic synthesis of minimalist multi-robot controllers for sequential tasks. The framework consists of a suite of algorithms that take as input a formal specification of the environmental effects, the task requirements, and the capabilities of the robots. The algorithms produce either provably correct robot controllers, or point to the exact scenarios and task segments which make (algorithmically) guaranteed task completion impossible. The type of controller and prospect of successful task execution depend on the capabilities of the individual robots. Current options include the possibly of broadcast inter-robot communications (Jones and Matarić, 2004a), and a local memory on each of the robots permitting non-reactive controllers (Jones and Matarić, 2004b). Two complementary analysis techniques allow various statistical performance claims to be made without the cost of a full implementation and exhaustive experimentation.

We do not provide full details of the framework here, but instead focus on the (non-obvious) role of ergodic dynamics in the work. The framework uses a set of states S to denote the possible states that the (assumed to be Markovian) world can be in. The set A contains actions which act upon the world state, producing state transitions defined in some probabilistic manner (see Figure 1). A particular sequence of states, say T , ($T \subset S$) makes up the task. In actuality the robots are only interested in producing the single task sequence, and thus only those transitions need to be stored. Thus, S is never stored or calculated, only T need be kept, and $|T| \ll |S|$. Robots then move around the environment making observations, perhaps consulting internal memory or listening to the broadcast communication channel if suitably equipped. If a robot has sufficient information to ensure that the performance of a particular action (from A) can only result a world transition that is part of the task (i.e., result in a state in T) then it may perform that action.

Returning to the notion of behavioral configuration space, each of the world states in S represents entire subspaces of the overall system’s space. Figure 2 shows that the entire behavioral configuration space as it fits into this formalism. A hypothetical projection of this entire (huge) configuration space separates the configurations into subspaces, each subspace representing a single state in T . Actions (from A) evolve the world state and hence transition the system from one subspace to another. *We design the system so that within each*




Figure 2 Three different views of the system dynamics: (a) A trajectory within the entire behavioral configuration space; (b) A hypothetical projection of the configuration space showing planar subspaces for behavior not part of S , and actions from A connecting these spaces; (c) The abstracted S representation as used by the described formalism.

subspace the dynamics are ergodic. Work that has used this formal framework ensured this property by having the robots perform randomized exploration policies. The randomized strategy needs to have sufficient effect to overpower other systematic biases in the system that could produce large scale effects and ignore some part of the configuration space.

Both controllers with memory (Jones and Matarić, 2004b) and ones endowed with communication capabilities (Jones and Matarić, 2004a) were demonstrated in simulation and on physical hardware in a multi-robot construction domain. The task involves the sequential placement of colored cubes into a planar arrangement. The sequence T contains simply the required evolution of the structure, actions A being the placement of an individual cube. Referring back to Figure 2; in the construction domain the motions within each subspace are random walks by the robots, and the transitions between spaces are cube placement actions.

Analytical techniques developed in order to predict task execution are aided by the ergodic components of the robots behavior in this domain. One example is in the macroscopic model (Jones and Matarić, 2004b, pp. 4–5) applied to this formal framework. This model calculates the probability of successful task completion by calculating a large multiplication of all possible memory states that get set, in each possible world state, after each possible observation, calculating the probability that only the correct action will result and includes terms for when actions may result in other, or null, world transitions. A fundamental assumption for that calculation is that no “structure” in the world results in the observation and action sequences that correlate. When endowed with navigational controllers that have ergodic dynamics, we know that this is true because the observation of an ergodic system at N arbitrary instants in time is statistically the same as N arbitrary points within the behavioral space (McQuarrie, 1976, pp. 554).

This section has demonstrated that dynamics with a high degree of ergodicity are achievable on physical robot systems. They can play a role in systems for which analytical methods exist, and as a very simple form of dynamics they can aid in simplifying particular aspects of system design.

5. Large-scale multi-robot systems

We consider large-scale multi-robot systems those with robots on the order of thousands. In spite of the fact that manufacturing and tractable simulation remain open challenges, a variety of tasks have been proposed for systems of this type. Increasing the number of robots increases the total number of degrees-of-freedom in a system, and results in a highly dimensional BCS. Coordination approaches that couple robot interactions as loosely as possible are most likely to scale to large sizes.

Mathematical techniques employed in statistical mechanics are useful for establishing the relationship between microscopic behavior and macroscopic structures (McQuarrie, 1976). Typical system sizes for classical work are significantly larger ($\sim 10^{23}$) than the numbers currently conceivable for robots. In the case of large (or infinite) systems, interesting macroscopic structures can result even from ergodic local dynamics. global structures like equilibrium phases, phase transitions, coexistence lines, and critical points are widely studied in thermodynamics. Recent work attempts to reformulate many of these classical notions for systems with fewer entities (Gross, 2001).

We are pursuing a methodology for coordination of large-scale systems through the study of a small set of mechanisms for producing general macroscopic phenomena. One candidate mechanism is a protocol for achieving consensus. The Potts (1952) model is illustrative; it is an archetypal magnetic spin system that models interactions between particles at a number of fixed locations within a graph or lattice. The Ising model (a specific Potts model) has also been used to model gas flow. Neither model is a perfect fit for robots, but illustrates macroscopic structure from simple interactions.

Mapping the spin interactions at spin sites to robots allows for the development of a communication algorithm that possesses ergodic dynamics (and an energy conservation constraint) that permits the definition of a partition function \mathcal{Z} that can be solved using a numerical method for pseudo-dynamics simulation (or in trivial cases analytically). This admits a prediction of global behavior because exhaustive parameter variations enable construction of a phase diagram. In the case of the Potts and Ising models this phase diagram is well known. Particular regions of the phase space in the Ising model represent regions of maximal order. For robots this means unanimity; consensus is reached through a second-order phase transition.

The ability to prescribe ergodic dynamics for large-scale robot systems makes those analytical approaches that focus only on constraint space topology feasible for predictions of global structure. This means that task directed actions can be tackled directly from a macroscopic perspective.

6. Summary and Conclusion

We have taken a dynamics-centric approach to describing multi-robot behavior. This view has suggested that the notion of ergodicity may be useful within a robotics context, something that we have demonstrated throughout the paper. After defining a behavioral configuration space, we demonstrated that subspaces in which the robot dynamics are essentially ergodic can be used to produce meaningful behavior, and allow automated synthesis techniques to focus on a small set of task-oriented states, rather than the entire ensemble configuration space. Also, in at least one case, ergodicity simplifies analysis of system behavior. Implementations on physical and simulated robots show that ergodicity is indeed achievable in the real world. Future promise of this general approach is suggested in a discussion of large-scale multi-robot systems.

Acknowledgments

This research was conducted at the Interaction Lab, part of the Robotics Research Lab at USC and of the Center for Robotics and Embedded Systems. It was supported by the Office of Naval Research MURI Grant N00014-01-1-0890.

References

- Donald, B. R. (1995). On information invariants in robotics. *AI*, 72(1-2):217–304.
- Erdmann, M. A. (1989). *On Probabilistic Strategies for Robot Tasks*. PhD thesis, M.I.T.
- Gerkey, B. P. and Mataric, M. J. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research*, 23(9):939–954.
- Gross, D. H. E. (2001). *Microcanonical Thermodynamics: Phase Transitions in “Small” Systems*. World Scientific Lecture Notes in Physics - Volume 66. World Scientific, Singapore.
- Jones, C. V. and Mataric, M. J. (2004a). Automatic Synthesis of Communication-Based Coordinated Multi-Robot Systems. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-04)*, Sendai, Japan.
- Jones, C. V. and Mataric, M. J. (2004b). Synthesis and Analysis of Non-Reactive Controllers for Multi-Robot Sequential Task Domains. In *Proc. of the International Symposium on Experimental Robotics*, Singapore.
- Klavins, E. (2003). Communication Complexity of Multi-Robot Systems. In Boissonnat, J.-D., Burdick, J., Goldberg, K., and Hutchinson, S., editors, *Algorithmic Foundations of Robotics V*, volume 7 of *Springer Tracts in Advanced Robotics*, pages 275–292. Springer.
- Lerman, K. and Galstyan, A. (2003). Macroscopic Analysis of Adaptive Task Allocation in Robots. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-03)*, pages 1951–1956, Las Vegas, NV., USA.
- Martinoli, A., Easton, K., and Agassounon, W. (2004). Modeling Swarm Robotic Systems: A Case Study in Collaborative Distributed Manipulation. *IJRR*, 23(4):415–436.
- McQuarrie, D. A. (1976). *Statistical Mechanics*. Harper and Row. reprinted by University Science Books, Sausalito, CA., USA in 2000.
- Potts, R. B. (1952). Some generalized order-disorder transformations. In *Proc. of the Cambridge Philosophical Society*, volume 48, pages 106–109.