Verifying Emergence of Bounded Time Properties in Probabilistic Swarm Systems

Alessio Lomuscio, Edoardo Pirovano

Imperial College London

Introduction

- Robot swarms give a number of advantages over single robots, such as scalability and resistance to faults.
- Robots in a swarm ("agents") follow simple protcols but these can lead to overall global behaviours such as moving in a formation.
- We refer to a behaviour that is displayed when a sufficient number of agents is present in the swarm as an *emergent property* and the number of agents that need to be present for this to be the case as an *emergence threshold*.

A swarm of kilobots

Models

- We model our swarms by an agent template, which captures the possible states and behaviours of one of the agents, and an environment which captures the rest of the state of the system.
- A *probabilistic agent template* is a tuple $T = \langle S, \iota, Act, P, t\rangle$ where:
  - The set S represents a set of agent local states.
  - $\iota \in S$ is a distinguished initial state.
  - The set $Act \neq \emptyset$ is a finite set of possible agent actions.
  - The agent's protocol function $P : S \times Act \to [0, 1]$ is such that for each $s \in S$ we have $\sum_{a \in Act} P(s, a) = 1$ giving the probability distribution for the next action of an agent given its current state.
  - The agent's transition function $t : S \times S_E \times \mathcal{P}(Act) \times Act \to S$ returns the agent's next state given its current state, the environment's current state, the set of actions performed by all the agents and the action performed by this agent at this time step.
- We similarly define an environment. Unlike the agents, the environment is deterministic. So, the environment evolves uniquely based on the actions of the agents.

Specifications

Specifications are expressed in a restricted form of PCTL given by the grammar, for $a \in AP$ and $k, i \in \mathbb{Z}$:

$$\chi ::= P_{\geq x}[\psi], \text{ for } x \in [0,1]$$
$$\psi ::= X^k\phi \mid \phi\ U^{\leq k}\phi \qquad \text{(path formulas)}$$
$$\phi ::= \top \mid (a, i) \mid \neg\phi \mid \phi \wedge \phi \qquad \text{(state formulas)}$$

For example, the property

$$P_{\geq 0.9}[X^{10}(connected, 1)]$$

says that, with probability at least 90%, agent 1 is conncected after 10 time steps.

EPI and ETI Decision Problems

- The emergent property identification (EPI) problem is concerned with establishing whether a given property $P_{\geq x}[\psi]$ is an emergent property of a given swarm system.
- We solve this by constructing an abstract model of the swarm system.
- The abstract model assumes that at each time step every action that is possible with some probability > 0 occurs.
- This assumption can be made to hold with a probability close to 1 by having a sufficiently large number of agents in the system.
- Thus, the model satisfies precisely those properties that are emergent properties of the system, and we can solve the EPI problem by simply model checking this abstract model.
- The related emergence threshold identification (ETI) problem involves finding a sufficiently large number of agents for the property to be guaranteed to hold.
- We solve this by constructing and solving a numerical equation which ensures that, with a sufficiently high probability, the assumption made in the abstract model holds.

Implementation and Evaluation

- We implemented our techniques in an open-source tool called PSV-BD (**P**robabilistic **S**warm **V**erifier for **B**ounde**D** time properties) based on the probabilistic model checker PRISM.
- We used this model an *ant coverage algorithm*.
- In this algorithm, a number of agents move in an $n \times n$ grid (that loops around so cell $(1, 1)$ is left of cell $(1, n)$).
- They store a value u (initialised at 0) on each cell and at each time step move to the neighbour with the lowest value of u (breaking ties randomly with equal probability) and increment this cell's u value.
- We considered the property $P_{\geq p}[F^{\leq k}\text{allVisited}]$ which says "with probability at least p, all cells are visited within k time steps".
- For different values of n and p, we used EPI to find the least k for which the property held, and ETI to find a corresponding emergence threshold.

Results

n	p	k	threshold	EPI time (s)	ETI time (s)
3	0.25	7	16	2.1	2.3
3	0.50	7	20	2.1	2.3
3	0.75	7	26	2.1	2.3
3	0.95	7	38	2.1	2.3
3	0.99	7	51	2.1	2.3
5	0.25	13	70	2.2	81
5	0.50	13	85	2.2	81
5	0.75	13	105	2.2	81
5	0.95	13	145	2.2	81
5	0.99	13	183	2.2	81
10	0.25	31	-	3.9	*timeout*

Conclusion

- We have introduced a semantics for formally verifying emergent properties of probabilistic swarm systems, and solved the corresponding decision problems.
- In future work, we intend to apply this to more swarm algorithms and investigate richer semantics (for example, specifications expressed in a larger fragment of PCTL).