Towards an Engineering Science of Robot Foraging

Alan FT Winfield

Abstract Foraging is a benchmark problem in robotics - especially for distributed autonomous robotic systems. The systematic study of robot foraging is important for several reasons: firstly, because foraging is a metaphor for the broad class of problems integrating robotic exploration, navigation and object identification, manipulation and transport; secondly, in multi-robot systems foraging is a canonical problem for the study of robot-robot cooperation; and thirdly, many and diverse actual or potential real-world applications for robotics are instances of foraging robots, for example, for cleaning, harvesting, search and rescue, landmine clearance or planetary astrobiology. This paper sets out a theoretical framework, structured upon an abstract model and taxonomy of robot foraging. A framework which, it is hoped, might provide the basis of a principled approach to the engineering of future real-world robot foraging systems.

1 Introduction

Foraging is a benchmark problem for robotics, especially for multi-robot systems. It is a powerful benchmark problem for several reasons: (1) sophisticated foraging observed in social insects, recently becoming well understood, provides both inspiration and system level models for artificial systems. (2) Foraging is a complex task involving the coordination of several - each also difficult - tasks including efficient exploration (searching) for objects, food or prey; physical collection (harvesting) of objects almost certainly requiring physical manipulation; homing or navigation whilst transporting those objects to collection point(s), and deposition of the objects before returning to foraging. (3) Effective multi-robot foraging requires cooperation between

---

Bristol Robotics Laboratory, University of the West of England, Bristol, UK, e-mail: Alan.Winfield@uwe.ac.uk

individuals involving either communication to signal to others where objects may be found (e.g. pheromone trails, or direction giving) and/or cooperative transport of objects too large for a single individual to transport.

There are, at the time of writing, very few types of foraging robots successfully employed in real-world applications. Most foraging robots are to be found in research laboratories or, if they are aimed at real-world applications, are at the stage of prototype or proof-of-concept. The reason for this is that foraging is a complex task which requires a range of competencies to be tightly integrated within the physical robot and, although the principles of robot foraging are now becoming established, many of the sub-system technologies required for foraging robots remain very challenging. In particular, sensing and situational awareness; power and energy autonomy; actuation, locomotion and safe navigation in unknown physical environments and proof of safety and dependability all remain difficult problems in robotics.

This paper proceeds as follows. Section 2 describes an abstract model of robot foraging, using a finite state machine (FSM) description to define the discrete sub-tasks, or states, that constitute foraging. Section 3 develops a taxonomy of robot foraging, together with some generalised performance metrics and a multi-robot case study. Section 4 then concludes the paper.

2 An Abstract model of Robot Foraging

Foraging robots are mobile robots capable of searching for and, when found, transporting objects to one or more collection points. Foraging robots may be single robots operating individually, or multiple robots operating collectively. Single foraging robots may be remotely tele-operated or semi-autonomous; multiple foraging robots are more likely to be fully autonomous systems [12].

Fig. 1 Finite State Machine for Basic Foraging

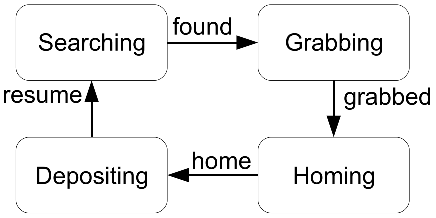

Figure 1 shows a Finite State Machine (FSM) representation of a foraging robot (or robots). In the model the robot is in always in one of four states: *searching*, *grabbing*, *homing* or *depositing*. Implied in this model is, firstly, that the environment or search space contains more than one of the target objects; secondly, that there is a single collection point (hence this model is sometimes referred to as central-place foraging), and thirdly, that the process continues indefinitely. The four states are defined as follows.

1. Searching. In this state the robot is physically moving through the search space using its sensors to locate and recognise the target items. At this level of abstraction we do not need to state how the robot searches: it could, for instance, wander at random, or it could employ a systematic strategy such as moving alternately left and right in a search pattern. The fact that the robot has to search at all follows from the pragmatic real-world assumptions that either the robot's sensors are of short range and/or the items are hidden (behind occluding obstacles for instance); in either event we must assume that the robot cannot find items simply by staying in one place and scanning the whole environment with its sensors. Object identification or recognition could require one of a wide range of sensors and techniques. When the robot finds an item it changes state from searching to grabbing. If the robot fails to find the target item then it remains in the searching state forever; searching is therefore the 'default' state.
2. Grabbing. In this state the robot physically captures and grabs the item ready to transport it back to the home region. Here we assume that the item is capable of being grabbed and conveyed by a single robot. As soon as the item has been grabbed the robot will change state to homing.
3. Homing. In this state the robot must move, with its collected object, to a home or nest region. Homing clearly requires a number of stages, firstly, determination of the position of the home region relative to where the robot is now, secondly, orientation toward that position and, thirdly, navigation to the home region. Again there are a number of strategies for homing: one would be to re-trace the robot's path back to the home region using, for instance, odometry or by following a marker trail; another would be to home in on a beacon with a long range beacon sensor. When the robot has successfully reached the home region it will change state to depositing.
4. Depositing. In this state the robot deposits or delivers the item in the home region, and then immediately changes state to searching and hence resumes its search.

There are clearly numerous variations on this basic foraging model. Some are simplifications: for instance if a robot is searching for one or a known fixed number of objects then the process will not loop indefinitely. Real robots do not have infinite energy and so a model of practical foraging would need to take account of energy management. However, many variations entail either complexity within one or more of the four basic states (consider, for instance, objects that actively evade capture - a predator-prey model of foraging), or complexity in the interaction or cooperation between robots in multi-robot foraging. Thus the basic model stands as a powerful top-level abstraction and a useful basis for extension to more complex foraging systems.

3 A Taxonomy of Robot Foraging

In robotics several taxonomies have been proposed for multi-robot systems. Dudek *et al* [3] define seven taxonomic axes: collective size; communications [range, topology and bandwidth]; collective reconfigurability; processing ability and collective composition. In contrast to Dudek's taxonomy which is based upon the characteristics of the robot(s), Balch [1] characterises tasks and rewards. Balch's task taxonomy is particularly relevant to robot foraging because it leads naturally to the definition of performance metrics. Balch articulates six task axes: time; criteria; subject of action; resource limits; group movement and platform capabilities. See also [4] for a formal analysis and taxonomy of task allocation. Østergaard *et al* [11] develop a simple taxonomy of foraging by defining eight characteristics each of which has two values: number of robots; number of sinks (collection points for foraged items); number of source areas (of objects to be collected); search space: unbounded or constrained; number of types of object to be collected; object placement: in fixed areas or randomly scattered; robots: homogeneous or heterogeneous and communication: none or with. This taxonomy does not capture task performance criteria, nor does it specify the strategy for either searching for, physically collecting or retrieving objects.

Tables 1 and 2 propose a more comprehensive taxonomy for robot foraging that incorporates the robot-centric and task/performance oriented features of Dudek *et al* and Balch, respectively, with the environmental features of Østergaard *et al.* The four major axes are Environment, Robot(s), Performance and Strategy. Each major axis has several minor axes and each of these can take the values enumerated in the third column of tables 1 and 2. The majority of the values are self-explanatory, those that are not are annotated.

3.1 Performance metrics

Following Balch [1], we can formalise successful object collection and retrieval as follows:

$$F(O_i, t) = \begin{cases} 1 \text{ if object } O_i \text{ is in a sink at time } t \\ 0 \text{ otherwise} \end{cases} \tag{1}$$

If the foraging task is performance time limited (Performance time = fixed) and the objective is to maximise the number of objects foraged within fixed time T, then we may define a performance metric for the number of objects collected in time T,

$$P = \sum_{i=1}^{N_o} F(O_i, t_0 + T) \tag{2}$$

Table 1 A taxonomy of robot foraging, part A

Major Axis	Minor Axis	Value	Notes
Environment	search space	unbounded	
		constrained	
	source areas	single limited	fixed number of objects
		single unlimited	e.g. objects 're-grow'
		multiple	
	sinks	single	home, nest or collection point
		multiple	multiple collection points
	object types	single static	unmoving object, food or 'prey'
		multiple static	multiple types of static object
		single active	e.g. prey which evades capture
	object placement	fixed known locations	
		uniform distribution	
		clustered	
Robot(s)	number	single	one robot
		multiple	multi-robot system
	type	homogeneous	one type of robot
		heterogeneous	multiple robot types
	object sensing	limited	short-range sensing
		unlimited	unlimited-range sensing
	localisation	none	
		relative	robots know relative position
		absolute	robots know absolute position
	communications	none	
		near	limited range robot-robot comms
		infinite	robots have infinite comms range
	power	limited	robots can run out of energy
		forage	robots forage for own energy
		unlimited	robots have unlimited energy

where N_o is the number of objects available for collection and t_0 is the start time. A metric for the number of objects foraged per second is clearly, $P_t = P/T$. P as defined here is independent of the number of robots.

In order to measure the performance improvement of multi-robot foraging, for example the benefit gained by search or homing with trail following, recruitment or coordination (assuming the task can be completed by a single robot, grabbing = single and transport = single), then we may define the performance of a single robot $P_s = P$ as defined in equation 2 and use this a baseline for the normalised performance P_m of a multi-robot system,

$$P_m = \frac{P}{N_r} \tag{3}$$

where N_r is the total number of robots. The efficiency of multi-robot foraging is then the ratio P_m/P_s.

Consider now that we wish instead to minimise the energy cost of foraging (Performance energy = minimum). If the energy cost of foraging object i is E_i, then we may define a performance metric for the number of objects

Table 2 A taxonomy of robot foraging, part B

Major Axis	Minor Axis	Value	Notes
Performance	time	fixed	metric: objects foraged per second
		minimum	minimise time to forage
		unlimited	
	energy	fixed	metric: objects foraged per Joule
		minimum	minimise energy used
		unlimited	
Strategy	search	random wander	
		geometrical pattern	
		trail following	e.g. follow ‘pheromone’ trail
		follow other robots	
		in teams	robots organise into search gangs
	grabbing	single	one robot can grab the object
		cooperative	e.g. stick pulling
	transport	single	one robot can transport the object
		cooperative	several robots needed for transport
	homing	self-navigation	e.g. using odometry
		home on beacon	
		follow trail	e.g. pheromone trail to home
	recruitment	none	
		direct	a robot that finds objects...
		indirect	...recruits others to the area
	coordination	none	
		self-organised	emergent coordination
		distributed	controlled coordination
		master slave	one robot acts as master to others
		central control	all robots under central command

foraged per Joule of energy,

$$P_e = \frac{N_o}{\sum_{i=1}^{N_o} E_i} \tag{4}$$

then seek the foraging strategy that achieves the highest value for P_e.

3.2 Multi-robot foraging with division of labour

As a taxonomic case study consider multi-robot foraging with division of labour. Division of labour in ant colonies has been well studied and in particular a response threshold model is described in [2]; in essence a threshold model means that an individual will engage in a task when the level of some task-associated stimulus exceeds its threshold.

For threshold-based multi-robot foraging with division of labour Figure 2 shows a generalised FSM for each robot. In this foraging model a robot will not search endlessly. If a robot fails to find a food-item because, for

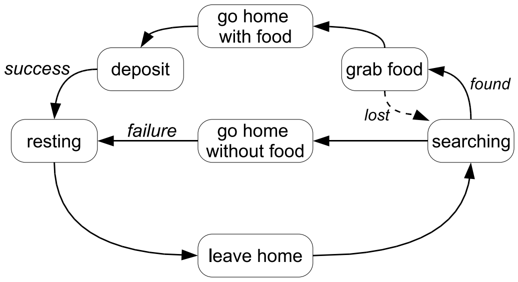

Fig. 2 Finite State Machine for Foraging with Division of Labour, adapted from [9]. *Environment search space=constrained; source areas=single unlimited; Robot number=multiple; power=forage; Performance energy=minimum; Strategy coordination=self-organised.*

instance, its searching time exceeds a maximum search time threshold T_s, or its energy level falls below a minimum energy threshold, then it will abandon its search and return home without food, shown as *failure*. Conversely *success* means food was found, grabbed and deposited. Note, however, that a robot might see a food-item but fail to grab it because, for instance, of competition with another robot for the same food-item. The robot now also has a *resting* state during which time it remains in the nest conserving energy. The robot will stop resting and leave home which might be according to some threshold criterion, such as its resting time exceeding the maximum rest time threshold T_r as in [9], or the overall nest energy falling below a given threshold [5].

4 Conclusion: Towards an Engineering Science

An engineering science requires both a theoretical framework and a set of tools for design and analysis. The abstract model, taxonomy and performance metrics set out above provide such a theoretical framework. Within this brief paper a comprehensive review of tools for design and analysis is not possible. However, because of its importance to a rigourous approach to robot foraging we shall here briefly review mathematical modelling.

A multi-robot system of foraging robots is typically a stochastic non-linear dynamical system and therefore challenging to mathematically model, but without such models any claims about the correctness of foraging algorithms are weak. Experiments in computer simulation or with real-robots (which provide in effect an ‘embodied’ simulation) allow limited exploration of the parameter space and can at best only provide weak inductive proof of correctness. Mathematical models, on the other hand, allow analysis of the whole parameter space and discovery of optimal parameters. In real-world applications, validation of a foraging robot system for safety and dependability will require a range of formal approaches including mathematical modelling.

Lerman, Martinoli and co-workers have developed the *macroscopic* approach to directly describe the collective behaviour of the robotic swarm. A class of macroscopic models have been used to study the effect of interference in a swarm of foraging robots [6] and collaborative stick-pulling [10]. Lerman *et al* [7] successfully expanded the macroscopic probabilistic model to study dynamic task allocation in a group of robots engaged in a puck collecting task. More recently Liu *et al* [8] have applied the macroscopic approach to develop a mathematical model for foraging with division of labour.

Although the principles of robot foraging are well understood, the engineering realisation of those principles remains a research problem. Consider multi-robot cooperative robot foraging. Although separate aspects have been thoroughly researched and demonstrated there has, to date, been no demonstration which fully integrates self-organised cooperative search, object manipulation and transport in unknown or unstructured real-world environments. Such a demonstration would be a precursor to a number of compelling real-world applications including search and rescue, toxic waste cleanup or foraging for recycling of materials.

References

1. Balch, T.: Taxonomies of multirobot task and reward. In: T. Balch, L. Parker (eds.) Robot Teams, pp. 23–35. A K Peters (2002)
2. Bonabeau, E., Theraulaz, G., Deneubourg, J.L.: Fixed response thresholds and the regulation of division of labour in insect societies. Bulletin of Mathematical Biology **60**, 753–807 (1998)
3. Dudek, G., Jenkin, M., Milios, E., Wilkes, D.: A taxonomy for multi-agent robotics. Autonomous Robots **3**, 375–397 (1996)
4. Gerkey, B.P., Matarić, M.J.: A formal analysis and taxonomy of task allocation in multi-robot systems. International Journal of Robotics Research **23**(9), 939–954 (2004)
5. Krieger, M., Billeter, J.B.: The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. Jour. of Robotics & Autonomous Systems **30**, 65–84 (2000)
6. Lerman, K.: Mathematical model of foraging in a group of robots: Effect of interference. Autonomous Robots, **13**(2), 127–141 (2002)
7. Lerman, K., Jones, C., Galstyan, A., Matarić, M.J.: Analysis of dynamic task allocation in multi-robot systems. Int. Journal of Robotics Research **25**(3), 225–242 (2006)
8. Liu, W., Winfield, A.F.T., Sa, J.: Modelling swarm robotic systems: A case study in collective foraging. In: Towards Autonomous Robotic Systems (TAROS 07), pp. 25–32. Aberystwyth (2007)
9. Liu, W., Winfield, A.F.T., Sa, J., Chen, J., Dou, L.: Towards energy optimisation: Emergent task allocation in a swarm of foraging robots. Adaptive Behaviour **15**(3), 289–305 (2007)
10. Martinoli, A., Easton, K., Agassounon, W.: Modeling swarm robotic systems: A case study in collaborative distributed manipulation. Int. Journal of Robotics Research, Special Issue on Experimental Robotics **23**(4), 415–436 (2004)
11. Østergaard, E.H., Sukhatme, G.S., Matarić, M.J.: Emergent bucket brigading: A simple mechanism for improving performance in multi-robot constrained-space foraging tasks. In: Proc. Int. Conf. on Autonomous Agents. Montreal, Canada (2001)
12. Winfield, A.F.T.: Foraging robots. In: R. Meyers (ed.) Encyclopedia of Complexity and System Science. Springer (2009, in press)