Coordinating Robotic Bee Swarms

Karthik Dantu, Bryan Kate, Jason Waterman, Peter Bailis and Matt Welsh
School of Engineering and Applied Sciences, Harvard University
{kar,bkate,waterman,pbailis,mdw}@eecs.harvard.edu

Introduction

The RoboBees project[1] is building swarms of micro-aerial vehicles at unprecedented size and scale. Individually, each sub-inch MAV will be extremely resource limited in terms of energy, computation, communication, and sensing. By harnessing the massive redundancy of the swarm numbering in the thousands, we can allow efficient and robust swarm operation. Most applications for swarms of robobees are performing operations in space. They could range from mapping environments, search and rescue, surveillance, commercial crop pollination etc. The robobees are expected to either search for objects of interest in space (eg. search for a target of interest) and/or actuate on previously acquired information (eg., track a target).

Our first effort in this direction has been to treat the robobees as simple drones that do not communicate with each other and are tasked by a central controller (hive). We built a system (*Karma*) that provides simple yet flexible programming model to be able to write programs treating the swarm as a whole and a scheduling mechanism that efficiently allocates bees to tasks in a robust fashion. A major shortfall of this design was lack of communication between the bees. Correspondingly, there is no coordination between the bees in the field introducing long latencies to react to observations. Also, actuation is an order of magnitude more expensive in energy consumption in comparison to computing and communication in the robobees. It is fairly expensive for the bees to return to the hive to get re-tasked. A solution to this problem requires the bees to coordinate in the field and operate independently when they are disconnected from the hive. A parallel observation is that the most relevant information for coordinated decision making for a robobee is information that is geographically local. An example is in a search application, once a robobee finishes searching its allotted sector, it could talk to its neighbors to go complete some other unfinished sectors instead of returning to the hive to be tasked the same task.

Vision

The conjecture is that each robobee minimally needs a *foveal view* of the world with information of higher resolution from places closer geographically to where the robobee

[1] `http://robobees.seas.harvard.edu`

is and decreasing for other place geographically further away. Our envisioned solution is a distributed data store that is shared among the bees providing the property mentioned above. The data store is accompanied by a programming model to interact with this data store allowing for easy storage and querying of information collected. Finally, there need to be some mechanisms for coordination that use such a data store for efficient operation of the swarm.

In this view, the swarm could be thought of as a distributed system with many concurrent operations being executed on the world to accomplish a given task. We view our solution analogous to being MapReduce for the physical world. But unlike traditional distributed systems, it is hard to provide a consistent view of the data that these operations are being performed on. Instead, an approximation is provided by providing a level of inconsistency that is acceptable (like the foveal view). We envision this as a useful property for other coordinated cyber-physical systems as well.

Also, unlike MapReduce on an array of servers where processors execute precisely, robobees could fail/have high variance in performance due to various reasons including environmental dynamics, individual failure and sensing/actuation error. Disconnected operation makes a central allocation more inefficient requiring coordination among sub-groups of processors to better execute parts of the task at hand.

Finally, prior sensor network algorithms like Trickle allow for reliable dissemination. We envision the need for the networking substrate to allow for reliability to be an application-defined value between 0 and 1 instead of being perfectly reliable all the time along with a user-defined aggregation function that can be used to disseminate information for such applications.

Background

I am a postdoctoral fellow with the School of Engineering and Applied Sciences at Harvard University affiliated with the RoboBees project working with Prof. Matt Welsh. I have worked on problems in load balancing in tiered sensor networks and connectivity/coverage problems in robot networks. I have also worked on low-power embedded systems working on algorithms for application-specific voltage scaling and power-aware networking. Most recently, as a postdoc on the robobees project, I have been working on *Karma*, a resource management system to make programming easy for a swarm of robotic bees while performing robustly and efficiently towards a given application goal. My current interests are in coordinated sensing systems with mobility/actuation.