

Analysis of ant colony behavior could yield better algorithms for network communication

13 July 2016, by Larry Hardesty



"It's intuitive that if a bunch of people are randomly walking around an area, the number of times they bump into each other will be a surrogate of the population density," says Cameron Musco, an MIT graduate student in electrical engineering and computer science. "What we're doing is giving a rigorous analysis behind that intuition, and also saying that the estimate is a very good estimate, rather than some coarse estimate."

Credit: Christine Daniloff/MIT

Ants, it turns out, are extremely good at estimating the concentration of other ants in their vicinity. This ability appears to play a role in several communal activities, particularly in the voting procedure whereby an ant colony selects a new nest.

Biologists have long suspected that ants base their population-density estimates on the frequency with which they—literally—bump into other ants while randomly exploring their environments.

That theory gets new support from a theoretical paper that researchers from MIT's Computer Science and Artificial Intelligence Laboratory will present at the Association for Computing

Machinery's Symposium on Principles of Distributed Computing conference later this month. The paper shows that observations from random exploration of the environment converge very quickly on an accurate estimate of population density. Indeed, they converge about as quickly as is theoretically possible.

Beyond offering support for biologists' suppositions, this theoretical framework also applies to the analysis of social networks, of collective decision making among robot swarms, and of communication in [ad hoc networks](#), such as networks of low-cost sensors scattered in forbidding environments.

"It's intuitive that if a bunch of people are randomly walking around an area, the number of times they bump into each other will be a surrogate of the population density," says Cameron Musco, an MIT graduate student in electrical engineering and computer science and a co-author on the new paper. "What we're doing is giving a rigorous analysis behind that intuition, and also saying that the estimate is a very good estimate, rather than some coarse estimate. As a function of time, it gets more and more accurate, and it goes nearly as fast as you would expect you could ever do."

Random walks

Musco and his coauthors—his advisor, NEC Professor of Software Science and Engineering Nancy Lynch, and Hsin-Hao Su, a postdoc in Lynch's group—characterize an ant's environment as a [grid](#), with some number of other ants scattered randomly across it. The ant of interest—call it the explorer—starts at some cell of the grid and, with equal probability, moves to one of the adjacent cells. Then, with equal probability, it moves to one of the cells adjacent to that one, and so on. In

statistics, this is referred to as a "random walk." The explorer counts the number of other ants inhabiting every cell it visits.

In their paper, the researchers compare the random walk to [random sampling](#), in which cells are selected from the grid at random and the number of ants counted. The accuracy of both approaches improves with each additional sample, but remarkably, the random walk converges on the true population density virtually as quickly as random sampling does.

That's important because in many practical cases, random sampling isn't an option. Suppose, for instance, that you want to write an algorithm to analyze an online social network—say, to estimate what fraction of the network self-describes as Republican. There's no publicly available list of the network's members; the only way to explore it is to pick an individual member and start tracing connections.

Similarly, in ad hoc networks, a given device knows only the locations of the devices in its immediate vicinity; it doesn't know the layout of the network as a whole. An algorithm that uses random walks to aggregate information from multiple devices would be much easier to implement than one that has to characterize the network as a whole.

Repeat encounters

The researchers' result is surprising because, at every step of a random walk, the explorer has a significant likelihood of returning to a cell that it has already visited. An estimate derived from random walks thus has a much higher chance of oversampling particular cells than one based on random sampling does.

Initially, Musco says, he and his colleagues assumed that this was a liability that an algorithm for estimating population density would have to overcome. But their attempts to filter out oversampled data seemed to worsen their algorithm's performance rather than improve it. Ultimately, they were able to explain why, theoretically.

"If you're randomly walking around a grid, you're not going to bump into everybody, because you're not going to cross the whole grid," Musco says. "So there's somebody on the far side of the grid that I have pretty much a zero percent chance of bumping into. But while I'll bump into those guys less, I'll bump into local guys more. I need to count all my interactions with the local guys to make up for the fact that there are these faraway guys that I'm never going to bump into. It sort of perfectly balances out. It's really easy to prove that, but it's not very intuitive, so it took us a while to realize this."

Generalizations

The grid that the researchers used to model the ants' environment is just a special instance of a data structure called a graph. A graph consists of nodes, typically represented by circles, and edges, typically represented as line segments connecting nodes. In the grid, each cell is a node, and it shares edges only with those cells immediately adjacent to it.

The researchers' analytic techniques, however, apply to any graph, such as one describing which members of a social network are connected, or which devices in an ad hoc network are within communication range of each other.

If the graph is not very well connected—if, for instance, it's just a chain of nodes, each connected only to the two nodes adjacent to it—then oversampling can become a problem. In a chain of, say, 100 nodes, an explorer taking a random walk could get stuck traversing the same five or six nodes over and over again.

But as long as two random walks starting from the same node are likely to branch out in different directions, as is often the case in graphs describing communication networks, random walks remain virtually as good as random sampling.

Moreover, in the new paper, the researchers analyze random walks executed by a single explorer. Pooling observations from many explorers would converge on an accurate estimate more quickly. "If they were robots instead of ants, they

could get gains by talking to each other and saying, 'Oh, this is my estimate,'" Musco says.

"Nancy's field is distributed computing, which has various strategies and methods that are pretty much unknown to biologists," say Anna Dornhaus, an associate professor of ecology and evolutionary biology at the University of Arizona, who studies social insects. "Nancy [Lynch] is at the forefront of realizing that these tools can actually be very useful to biologists. She's trying to do this interdisciplinary research and really enable us to perhaps make a leap forward in understanding biological systems."

"People always debate whether [ants](#) or bees can recognize other individuals," Dornhaus explains. "This paper shows that at least in this context, that's not necessary. For me, that's the main surprising result here. But of course, they can also prove mathematically how accurate that strategy is."

More information: Ant-Inspired Density Estimation via Random Walks.
arxiv.org/pdf/1603.02981v1.pdf

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

APA citation: Analysis of ant colony behavior could yield better algorithms for network communication (2016, July 13) retrieved 12 May 2021 from <https://phys.org/news/2016-07-analysis-ant-colony-behavior-yield.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.