

# Who will control the swarm?

July 25 2017, by Andrew Myers

---



Credit: Asad Photo Maldives from Pexels

The world is already well on its way to a day when innumerable autonomous cars and drones buzz about, shuffling commuters to work and packages to doorsteps. In fact, there is new term for it floating around the circles of engineers and venture capitalists who hope to see the day arrive sooner rather than later: They call it "The Swarm."

But this dramatic-sounding reality raises a critical question that has yet to be answered: Who will control the swarm? Some say control will be distributed. Each car and every drone will be its own self-sustaining unit – individually aware of its surroundings, individually directed where to go and individually outfitted with all the computational power to make it through the world efficiently and without accident.

A collection of faculty at Stanford have a different view. They believe that [device](#) swarms will be managed centrally, using applications running in large datacenters, much the way the cloud centralized big data. The faculty members have formed a new laboratory at Stanford, called the Platform Lab, to develop infrastructure for these new "Big Control" applications.

"We think all these self-driving cars and drones will be controlled not individually, but centrally, in a coordinated fashion," says John Ousterhout, faculty director of the lab. "This has the potential to change how society functions on a daily basis."

While most current research into autonomous vehicles assumes a distributed model – relatively autonomous devices, controlled in a peer-to-peer fashion with each machine doing its own calculations – the concentrated model has its advantages, says Ousterhout.

First is the ease of creating applications. Writing applications for the distributed model is very difficult, since each device has limited information about the state of the world. With the centralized approach, data from all the devices is collected in one place. This provides a big-picture view of the world that allows better control of higher-level tasks like system-wide situational perception, decision-making and large-scale traffic planning.

Second, control applications running in datacenters have many more

resources available, such as computing horsepower and large back-end datasets. This allows them to implement more sophisticated collaborative behaviors for the device swarm. In addition, the centralized applications can take advantage of powerful machine learning algorithms, which allow the control system to learn and improve its behavior.

An added challenge of the distributed model that centralization solves is that computational limitations in the devices themselves hinder overall system sophistication, an issue that becomes ever-more pronounced as scale increases. As new technologies arise, they either must be retrofitted into each device, or, when the old devices become obsolete, the devices must be replaced – an expensive proposition.

In the centralized model, the vehicle is merely a tool – a relatively dumb device fitted with equipment by which to see the road and the skies ahead, to detect obstacles and other vehicles in the way, to provide geolocation and so forth. The gathered data gets transferred back to the cloud and processed en masse by much faster computers able to handle the mathematical demand of keeping track of those millions of vehicles and to plan ways around bottlenecks and hazards to efficiently and safely guide passengers and packages to their many destinations.

"From a technology standpoint, it is attractive and easiest to centralize control – to amass data, plan and then disseminate a singular view to all devices," says Ousterhout.

Not all functions are suited to the centralized model, however. The Platform Lab foresees that devices will retain local control for things like device stability and near-term collision avoidance. Such control needs microsecond or sub-millisecond response time and must happen on the device.

The implications of a centralized system reach beyond easier commutes.

One prime example is in disaster recovery. In a community devastated by earthquake, fire or flood, it often proves too dangerous for human first responders. In such cases, a flock of autonomous drones could be dispatched to assess the situation, allowing emergency management to triage from afar.

Another example cited would be a massive warehouse where 10,000 or more drones operate indoors, in a closed environment all watched over by cameras and sensors to monitor, organize and move millions of packages each day.

There remains one looming problem in all this talk of centralization: Much of the infrastructure does not yet exist. The range of computational and communication capabilities necessary to pull it off is staggering – GPS, mapping, wireless communications, situational awareness and traffic coordination are only the most obvious components.

Another challenge is to provide a massive amount of computing with extremely low and predictable latency. That means leveraging new machine learning and artificial intelligence techniques to ensure planning and control happen fast without data inference.

None of this deters Ousterhout or Guru Parulkar, the lab's executive director and a consulting professor of electrical engineering. With pieces of the puzzle still in flux, the mission is to imagine how this centralized future might function and to determine what pieces exist, which need to be improved and what others are yet to be created to make it all function seamlessly. To this end, they are laying out the roadmap of the platform architecture, from the applications that will be needed to track packages, manage commutes, coordinate disaster relief and provide mapping to deep learning, adaptive scheduling and data-gathering tools.

Deeper down in the platform, the lab foresees the need for things like new hardware accelerators, better ways for computers to manage the many computing threads occurring simultaneously, rapid data storage and retrieval, and improved cluster scheduling necessary to execute the massive number of computations centralized [control](#) will demand. And, of course, security will be a preeminent concern. Most of these things must still be created, but simply knowing the need is a first step to realizing the vision.

If it works, the economic and social consequences will be profound. Ousterhout and Parulkar point to separate studies by the Federal Aviation Administration and PricewaterhouseCoopers that foresee 7 million drones driving \$127 billion in economic activity, respectively, all by the year 2020.

The centralized vision is not without its controversy, but settling those debates is precisely why the founders created the lab. For industry leaders who are well along the path to creating individually autonomous cars and drones under the assumption of a distributed model, Parulkar says there is common ground and a great need for that sort of experience and opinion at the lab. Big Control represents a new class of applications that is both exciting and not so futuristic as to be impractical. "It's big, but it's achievable," Parulkar says.

Provided by Stanford University

Citation: Who will control the swarm? (2017, July 25) retrieved 11 May 2024 from <https://phys.org/news/2017-07-swarm.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.