Hierarchal Control of Intelligent Agents for Cyber Physical Systems

Aaron Becker

***Abstract*— Technological innovations continue to decrease the cost, size, and energy demands of computation. Traditional centralized control techniques are being transferred to distributed control, where individual agents sense their environment, exchange information with neighboring agents, and act rationally according to their model of the world. While intelligent agents in large populations can exhibit collective wisdom, unfortunately they are often prone to mob mentality failures: bandwagon errors, cascade failures, and panic. Our proposed research is to investigate supervisory control for large populations of intelligent, distributed agents. We will investigate hierarchal control in the context of multi-robot manipulation, with potential applications in transportation, supply-management, medical treatment, and defense.**

I. Background

Large populations of computing devices that (1) interact with the physical world and (2) are joined in communication networks are in use today. The number of these cyber physical systems (CPS) is increasing, and the large numbers of devices strains traditional control architectures.

As in human societies, where communal decision making gave way to a single leader, which in turn was succeeded by hierarchal government, so too will CPS require levels of abstraction to efficiently robustly respond and sustainably exploit resources.

We desire complex CPS to demonstrate *wisdom of the crowd*. That is, we expect to CPS to have superior judgement over individual decision-makers. Unfortunately, crowds are often susceptible to poor judgment, as illustrated in Fig. 2a. On January 27, thousands of intelligent motorists were stranded on interstate highways because they all tried to escape a coming snow storm at roughly the same time. Evacuations are exemplary of this phenomenon, because

A. Becker is with the Department of Cardiovascular Surgery, Boston Children's Hospital and Harvard Medical School, Boston, MA, 02115 USA aaron.becker@childrens.harvard.edu.

(a) r-one robot platform

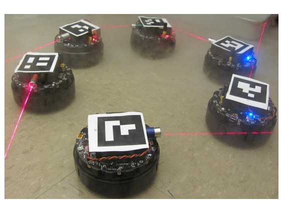

(b) Centralized controllers

Fig. 1. Centralized controllers for manipulation avoid local minima, but ignore the decision-making capabilities of individual robots and due to communication constraints, miss much of the sensory inputs of the individual robots [1], [2]. These robots, which cost $250 USD, represent how the component cyber-physical agents in a CPS can *each* be sophisticated platforms with many sensors, actuators, and substantial processing power.

(a) crowd wisdom fails

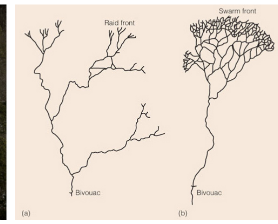

(b) army ant forage patterns

Fig. 2. **(b)** Three inches of snow turned areas of Atlanta's interstate into parking lots, with cars stalled for 36 hours. Though each car operator was a rational agent, the catastrophe stemmed from schools, government, and businesses dismissing students and workers simultaneously. (Photo credit: http://www.wbur.org/npr). **(b)** Army ants are some of the best foragers in the natural world. This figure shows different foraging patterns from different species of ants, illustrating how a small genomic change leads to a broad array of behaviors [3].

often many agents choose the same option from a set of choices, and *none* are able to exploit it.

Similar problems occur in allocating communication bandwidth [4]. Other population-related problems include *cascading failure*, such as when the failure of one financial institution may cause other financial institutions to fail, and at the cellular level in stroke victims, when a small ischemic attack releases toxins that kill off many more cells than the initial attack.

CPS research can benefit from work in many areas that combine distributed decision making with a centralized controller.

a) Entomology: army ants are efficient foragers that rely strongly on distributed decision making. By comparing different species it is apparent that small genomic changes strongly influence foraging patterns. These patterns can be simulated with high fidelity, as shown in Fig. 2b. A centralized controller could efficiently command the hive to switch behaviors to better exploit resources if a timely abstraction of the world state is available.

b) Cognitive Radio: There are two leading models for allocating bandwidth in cognitive radio, a centralized, iterative water-filling approach that is rooted in communication theory; and a distributed, no-regret algorithm rooted in machine learning [4]. The centralized approach is faster, but the distributed approach is more robust. A hybrid approach could combine these.

c) Defense: centralized controllers waste resources if they deny the decision-making capabilities of the agents they control. A famous example is the charge of the light brigade, where a lightly armed calvary, instead of being sent to harass a retreating artillery battery, was mistakenly sent against a dug-in artillery battery. As Tennyson immortalized them, *Theirs not to make reply, Theirs not to reason why,* *Theirs but to do and die.* To avoid this fate, communication between agents and decision-makers is necessary. Individual CPS agents should inform decision-makers about the current world state, but due to the sheer numbers of CPS agents, this information must be presented in an *abstracted*, i.e. organized and averaged, form.

II. Proposed Research

The key research questions are "Why do mobs make worse decisions than individuals?" and "what information and control structure will avoid this?"

As cyber physical systems increase in number and responsibilities, we need control structures that guarantee that groups of agents behave no worse than individuals.

The key idea is that each component system is intelligent, and the price for computation and the requisite energy is dropping. Because the systems are increasingly intelligent, control that is centralized and dictatorial is unreasonable. A centralized controller must tradeoff between micro-managing and generic control. Choosing to be efficient in communication by sending generic commands to all the agents inefficiently utilizes the agents, but micro-managing the agents to maximize each agents' contribution requires large amounts of communication.

Compared to centralized control, distributed systems of homogenous agents require longer to reach consensus. Communication is the bottleneck. Similarly in political states, once the group grows beyond a certain size, pure democracies become inefficient. Instead specialized representatives are appointed. These representatives contribute a greater amount of time and every to decision making and communication, and naturally form a hierarchy.

My background is multi-robotic systems, so I propose to study these issues in the context of *massive manipulation*—manipulation with very large populations—from an algorithmic and a control-theoretic perspective. Robotic manipulation at the micro- and nanoscale can fundamentally transform how we build and assemble objects. One of my goals is to assemble structures from the bottom-up; to fabricate evermore complex assemblies by pushing the minimum size of a component to a few microns, or even a few dozen molecules. This kind of precision manipulation must be coupled with a large population of manipulators in order to enable rapid progress. The potential impact is broad: large populations of micro-manipulators would enable surgeons to eliminate cancer at the cellular level, engineers develop complex MEMS assemblies, and biologists to simultaneously sort all the cells on a Petri dish. Moving to the nanoscale, manipulation offers almost limitless possibilities, enabling construction at the molecular level.

The central research questions are how to: (1) minimize communication between hierarchal controllers (2) minimize the minimum number of informed agents while maximizing the number of actively working agents (3) provide performance guarantees.

To make progress, we are investigating human control of large populations of robots. Though online experiments we can efficiently run large-scale experiments with 1000s of human users and 100s of CPS robots (Fig. 3). These approaches can then be refined with hardware simulations involving 100s of communicating hardware robots (Fig. 4b).

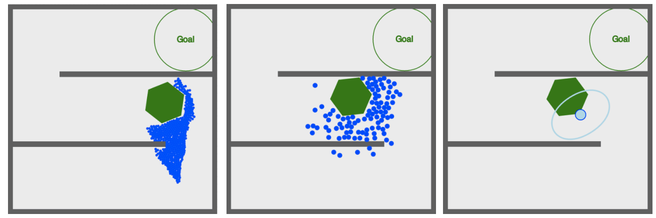

(a) Vary Number (b) Vary Visual Feedback

Fig. 3. Screenshots from our online experiments swarmcontrol.net controlling multi-robot systems. These experiments allow us to run large-scale experiments (1000s of trials) where human users control 100s of CPSs (simulated robots) in real time [5].

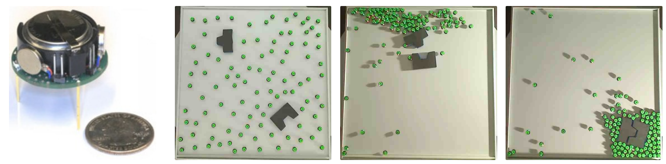

(a) 1 Kilobot (b) Directing a CPS swarm of kilobots to assemble a part

III. Potential Impact to CPS

The potential impacts of efficient hierarchal control of CPS are broad. The optimal division of labor between laborers that sense and actuate, and decision-makers that form abstractions and deliberate is a difficult question, but is tractable in the context of multi-robot manipulation. Advances on this problem can be applied to other arenas with large populations. These are diverse and include optimizing transportation both in rush hours and evacuations; systemic control of medical micro robots treating a patient; and warehouse and supply-chain management.

Both economists and political scientists study how large populations of intelligent agents react and make decisions. Models from these domains can be leveraged for large populations of CPSs. CPS have the advantage of being logical, but interacting CPS will have competing goals. Efficiently using resources requires communication.

References

[1] A. Becker, G. Habibi, J. Werfel, M. Rubenstein, and J. McLurkin, "Massive uniform manipulation: Controlling large populations of simple robots with a common input signal," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov. 2013, pp. 520–527.

[2] A. Becker, E. Demaine, S. Fekete, G. Habibi, and J. McLurkin, "Reconfiguring massive particle swarms with limited, global control," in *International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics (ALGOSENSORS)*, Sep. 2013.

[3] I. Couzin and N. Franks, "Self-organized lane formation and optimized traffic flow in army ants," *Proceedings of the Royal Society of London. Series B: Biological Sciences*, vol. 270, no. 1511, pp. 139–146, 2003.

[4] S. Haykin, "Cognitive radio: brain-empowered wireless communications," *Selected Areas in Communications, IEEE Journal on*, vol. 23, no. 2, pp. 201–220, 2005.

[5] A. Becker, C. Ertel, and J. McLurkin, "Online massive manipulation: A user study on controlling large swarms of simple robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.