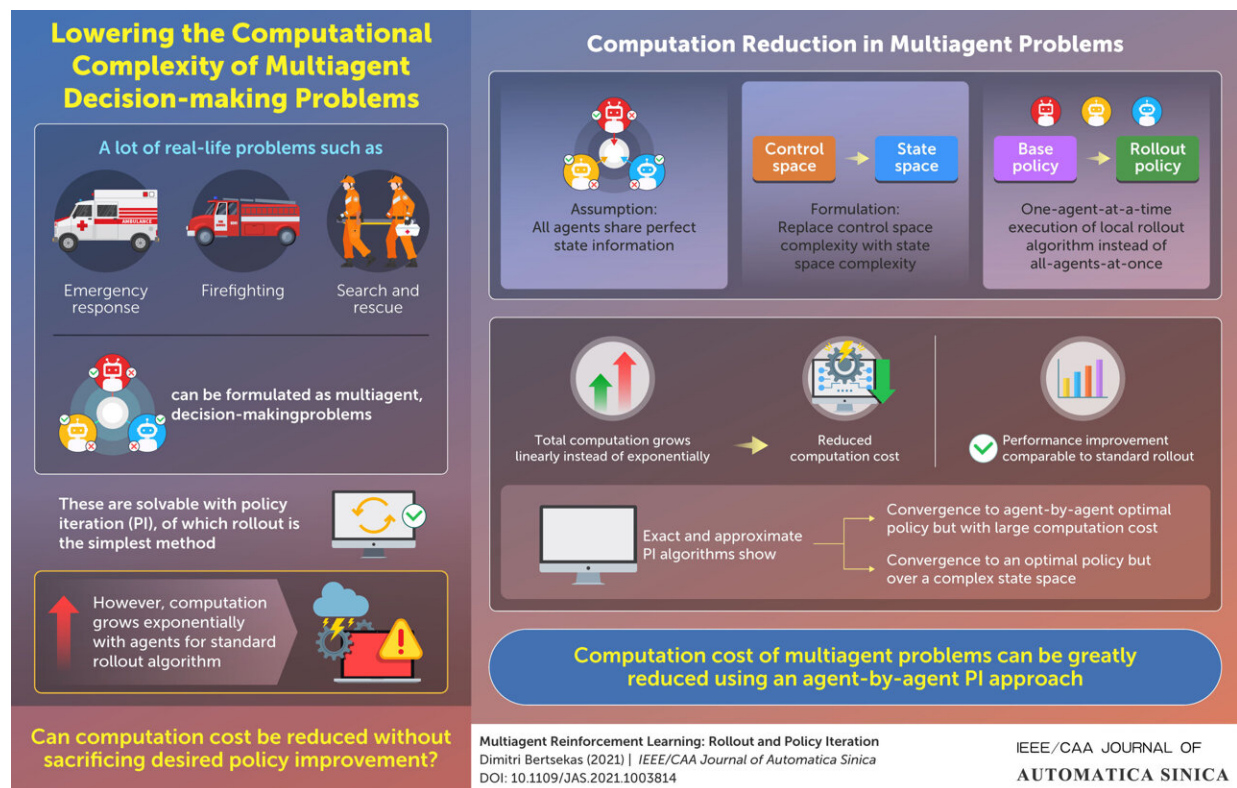


New algorithm makes it easier for computers to solve decision making problems

April 28 2021



Computation cost of multiagent problems can be greatly reduced using an agent-by-agent PI approach. Credit: IEEE/CAA Journal of Automatica Sinica


Computer scientists often encounter problems relevant to real-life scenarios. For instance, "multiagent problems," a category characterized by multi-stage decision-making by multiple decision makers or "agents,"

has relevant applications in search-and-rescue missions, firefighting, and emergency response.

Multiagent problems are often solved using a machine learning technique known as reinforcement learning (RL), which concerns itself with how intelligent agents make decisions in an environment unfamiliar to them. An approach usually adopted in such an endeavor is policy iteration (PI), which starts off with a 'base policy' and then improves on it to generate a 'rollout policy' (with the process of generation called a rollout). Rollout is simple, reliable, and well-suited for an on-line, model-free implementation.

There is, however, a serious issue. "In a standard rollout algorithm, the amount of total computation grows exponentially with the number of agents. This can make the computations prohibitively expensive even for a modest number of agents," explains Prof. Dimitri Bertsekas from Massachusetts Institute of Technology and Arizona State University, USA, who studies large-scale computation and optimization of communication and control.

In essence, PI is simply a repeated application of rollout, in which the rollout policy at each iteration becomes the base policy for the next iteration. Usually, in a standard multiagent rollout policy, all agents are allowed to influence the rollout algorithm at once ("all-agents-at-once" policy). Now, in a new study published in the *IEEE/CAA Journal of Automatica Sinica*, Prof. Bertsekas has come up with an approach that might be a game changer.



Conceptual structure of a multiagent system. Each agent is supplied a "state info" from the "cloud" at each stage of decision making, which then performs local computations based on this state information to apply their controls. Credit: IEEE/CAA Journal of Automatica Sinica, Image courtesy: Dimitri Bertsekas

In his paper, Prof. Bertsekas focused on applying PI to problems with a multiple-component control, each component selected by a different agent. He assumed that all agents had perfect state information and shared it among themselves. He then reformulated the problem by trading off control space complexity with state space complexity. Additionally, instead of an all-agents-at-once policy, he adopted an agent-by-agent policy wherein only one agent was allowed to execute a rollout algorithm at a time, with coordinating information provided by the other agents.

The result was impressive. Instead of an exponentially growing

complexity, Prof. Bertsekas found only a linear growth in computation with the number of agents, leading to a dramatic reduction in the computation cost. Moreover, the computational simplification did not sacrifice the quality of the improved policy, performing at par with the standard rollout algorithm.

Prof. Bertsekas then explored exact and approximate PI algorithms using the new version of agent-by-agent [policy](#) improvement and repeated application of rollout. For highly [complex problems](#), he explored the use of neural networks to encode the successive rollout policies, and to precompute signaling policies that coordinate the parallel computations of different [agents](#).

Overall, Prof. Bertsekas is optimistic about his findings and future prospects of his approach. "The idea of agent-by-agent rollout can be applied to challenging multidimensional control problems, as well as deterministic discrete/combinatorial optimization problems, involving constraints that couple the controls of different stages," he observes. He has published two books on RL, one of which, titled "Rollout, Policy Iteration, and Distributed Reinforcement Learning" soon to be published by Tsinghua Press, China, deals with the subject of his study in detail.

The new approach to multiagent systems might very well revolutionize how complex sequential decision problems are solved.

More information: Dimitri Bertsekas. Multiagent Reinforcement Learning: Rollout and Policy Iteration, *IEEE/CAA Journal of Automatica Sinica* (2021). [DOI: 10.1109/JAS.2021.1003814](https://doi.org/10.1109/JAS.2021.1003814)

Provided by Chinese Association of Automation

Citation: New algorithm makes it easier for computers to solve decision making problems (2021, April 28) retrieved 10 April 2024 from <https://techxplore.com/news/2021-04-algorithm-easier-decision-problems.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.