Approximating cost structures in decentralized auctions and optimization for multi-agent task allocation

PhD Proposal

ONERA, Toulouse, France

February 2024

URL

`https://w3.onera.fr/formationparlarecherche/multi-agent-systems-auctions-coordination-approximate-bidding-reinforcement-learning`

Keywords

Multi-agent systems, auctions, coordination, approximate bidding, reinforcement learning

Context and Motivations

In the context of multi-robot or multi-vehicle applications, such as search-and-rescue, urban UAV traffic management, or multi-agent pick-up and delivery, coordination is a key element that provides a multi-perspective and multi-skill approach to problems hardly covered by single robots. Yet, from the performance viewpoint, optimizing resource usage, mission duration and quality, coordination is a hard problem. In fact, when (semi-)autonomous agents have to decide which tasks to fulfill, which path to follow or which multi-agent action to perform, one has to solve hard combinatorial problems such as task allocation [12], resource allocation [2], multi-agent pathfinding [14] or multi-agent trajectory repair [9]. Such problems can be solved using a classical centralized approach, such as combinatorial auctions or mathematical programming [13].

However, in the context of missions where the decisions might be decentralized to (i) improve robustness to disconnections or (ii) improve reactiveness (without waiting for a central authority to make the decision), decentralizing such decision-making is crucial. Moreover, these algorithms strongly rely on the capability for agents to bid over items (tasks/actions/resources/etc.) or bundle of items. In classical auctions or mathematical programming approaches, this bidding requires to be able to valuate each combination (or some of them) of items, or to use more compact representations of sets of interdependent items [4]. This information may take the form of large tables or graphs, may be too large to be computed at fast pace or to be exchanged within unreliable communication infrastructure. Moreover, in the case of multi-agent tasks (requiring several agents to be performed), as in CBGA algorithm [1], or in the case of multi-mode tasks (tasks that can be fulfilled in different manners) as in MM-CBGA algorithm [10], data structures are becoming even larger (to process and to communicate).

Besides auctions, such table-based decision representation is also at the core of some distributed optimization algorithms (DCOP) [6], especially inference-based ones, such as DPOP or Max-Sum, and we aim at providing generic theoretical tools to use both in auctions and DCOPs domains.

In the domain of reinforcement learning, such an approximation and compacting problem has been handled using deep neural networks or other approximation functions to represent and functionally approximate Q-tables for instance [7]. In the domain of multi-agent pick and delivery, attention models have been used to form collectives and allocate tasks [5]. We aim at exploring the use of such techniques in the context of auctions and DCOPs.

Objectives

The idea of this thesis is thus to bring approximation to bidding to implement decentralized auctions and distributed inference with good reactiveness and low communication load. The idea is to first devise approximation schemes for mono-agent mono-mode tasks decentralized allocation, in the consensus-based bundle allocation framework (CBBA) [3], and then to extend the models for multi-agent tasks (CBGA) and multi-mode tasks (MM-CBGA). In a second time, we will investigate how to implement these approximation schemes to the inference-based DCOP framework, in complete (DPOP) and incomplete (Max-Sum) solution methods.

We previously applied auctions and DCOP domains such as Earth observation [9, 10], Unmanned Aircraft System Traffic Management (UTM) [8], and multi-robot missions [11]. We thus envision to implement and evaluate the theoretical and algorithmic contributions of this thesis to one or more of these application domains.

Hosting institution

The *Office national d'études et de recherches aérospatiales* (ONERA) is the French national aerospace research lab. The PhD student will be hosted at ONERA, Toulouse, France (`https://www.onera.fr/en/centers/toulouse`), and will have opportunities to visit IIIA-CSIC, in Barcelona.

Collaboration

The PhD thesis will be jointly supervised by ONERA (Toulouse, France) and IIIA-CSIC (Barcelona, Spain).

Supervisors

- Pr. Gauthier Picard, Directeur de recherche, ONERA
- Dr. Filippo Bistaffa, Tenured Researcher, IIIA-CSIC
- Pr. Juan A. Rodríguez-Aguilar, Research Professor, IIIA-CSIC

Profile and skills required

Master-level with strong skills in AI, optimization, multi-agent, RL

How to apply

Applications (CV, motivation letter, referees, and grades) should be sent to `gauthier.picard@onera.fr, juanantonio.rodriguez@csic.es, filippo.bistaffa@iiia.csic.es` with the subject `"Application to PhD position TIS-DTIS-2024-14"`.

About Toulouse

Beyond its iconic rose-hued brick architecture, Toulouse offers a unique blend of intellectual stimulation, cultural flair, and a relaxed atmosphere that's perfect for thriving during your PhD journey.

Renowned for its top-ranked universities and research institutions, particularly in aerospace and space, Toulouse is a hub for cutting-edge knowledge. You'll have opportunities to collaborate with brilliant minds, access state-of-the-art facilities, and contribute to groundbreaking research that pushes the boundaries of your field.

Toulouse boasts a dynamic student life, where you can connect with a diverse international community and explore the surrounding vineyards and medieval towns on weekend getaways. With its

affordable cost of living, excellent public transportation, and abundance of green spaces, Toulouse offers a comfortable and enjoyable environment to call home during your studies.

More info here: `https://en.wikipedia.org/wiki/Toulouse`

Important dates

- Application deadline: April 1st, 2024
- Interviews: April 15-22, 2024
- Final decision: April 30, 2024
- Start of PhD study: October 2024 (can be adjusted)

References

[1] Luc Brunet, Han-Lim Choi, and Jonathan How. Consensus-based auction approaches for decentralized task assignment. In *AIAA Guidance, Navigation and Control Conference and Exhibit*. American Institute of Aeronautics and Astronautics, June 2008.

[2] Yann Chevaleyre, Paul E. Dunne, Ulrich Endriss, Jérôme Lang, Michel Lemaître, Nicolas Maudet, J. Padget, Steve Phelps, Juan A. Rodríguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica (Slovenia)*, 30:3–31, 2006.

[3] Han-Lim Choi, L. Brunet, and J.P. How. Consensus-based decentralized auctions for robust task allocation. *IEEE Transactions on Robotics*, 25(4):912–926, August 2009.

[4] P. Cramton, Y. Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2005.

[5] Adrià Fenoy, Filippo Bistaffa, and Alessandro Farinelli. An attention model for the formation of collectives in real-world domains. *Artificial Intelligence*, 328:104064, March 2024.

[6] Ferdinando Fioretto, Enrico Pontelli, and William Yeoh. Distributed constraint optimization problems and applications: A survey. *Journal of Artificial Intelligence Research*, 61:623–698, March 2018.

[7] Vincent François-Lavet, Peter Henderson, Riashat Islam, Marc G. Bellemare, and Joelle Pineau. An introduction to deep reinforcement learning. *Foundations and Trends® in Machine Learning*, 11(3–4):219–354, 2018.

[8] Gauthier Picard. Auction-based and distributed optimization approaches for scheduling observations in satellite constellations with exclusive orbit portions. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '22, page 1056–1064, Richland, SC, 5 2022. International Foundation for Autonomous Agents and Multiagent Systems.

[9] Gauthier Picard. Trajectory coordination based on distributed constraint optimization techniques in unmanned air traffic management. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '22, page 1065–1073, Richland, SC, 5 2022. International Foundation for Autonomous Agents and Multiagent Systems.

[10] Gauthier Picard. Multi-agent consensus-based bundle allocation for multi-mode composite tasks. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, page 504–512, Richland, SC, 5 2023. International Foundation for Autonomous Agents and Multiagent Systems.

[11] Felix Quinton, Christophe Grand, and Charles Lesire. Communication-preserving bids in market-based task allocation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, October 2022.

[12] Onn Shehory and Sarit Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1–2):165–200, May 1998.

[13] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, December 2008.

[14] Roni Stern, Nathan Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, T. K. Kumar, Roman Barták, and Eli Boyarski. Multi-agent pathfinding: Definitions, variants, and benchmarks. *Proceedings of the International Symposium on Combinatorial Search*, 10(1):151–158, September 2021.