Co-Design of multi-agent heterogeneous robotic systems for search tasks

Maximilian Stralz*, Meshal Alharbi*, Yujun Huang* and Gioele Zardini*
*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology

***Abstract*—In this work, we present a framework which allows to reason about multi-agent heterogeneous robotic systems in search tasks. Leveraging a monotone theory of Co-Design, our framework jointly designs hardware and software components of multi-agent systems. Thereby, making the fleet itself an optimization objective, rather than optimizing procedures for already existing fleets. A case study on coverage tasks showcases the power of the framework in extracting optimal solutions and reasoning about the trade-offs in design decisions such as cost, energy consumption and coverage time.**

I. INTRODUCTION

Multi-agent robotic systems have gained increasing prominence over the past decades, enabled by the maturity of single-robot platforms and their effectiveness in complex tasks [1]. Applications span warehouse logistics [2, 3], domestic service robots [4], and safety-critical domains such as gas-leak detection [5] and de-mining [4]. A central challenge is area coverage, i.e., coordinating sensor-equipped fleets to monitor or survey environments, which is vital in search and rescue (SAR), forest fire monitoring [6], industrial inspection [7], and disaster response [8]. Research in this field largely focuses on coverage path planning (CPP) while typically assuming fixed, often homogeneous, fleets [1, 4].

Contributions – We introduce a framework for formally modeling multi-agent robotic systems in search tasks, explicitly accounting for both homogeneous and heterogeneous fleets. Through a case study, we show that the proposed framework can integrate, but is not limited to, CPP algorithms with a monotone theory of Co-Design [9, 10], thereby enabling the joint optimization of fleet composition, robot design, and planner selection. This holistic perspective provides a principled way to explore trade-offs among energy, cost, and task completion time, supporting well-informed design decisions.

Related work – In [11], Hu et al. propose a framework for large-scale heterogeneous multi-robot coverage which decomposes the search domain into sub-regions, employing a evolution-guided generative adversarial network to efficiently solve the multi-robot task allocation (MRTA) problem by producing Pareto optimal team allocations, hereafter referred to as the MRTA planner. Path planning within each sub-region is then performed in a decentralized manner, where each team optimizes its trajectories by minimizing an ergodic metric, ensuring spatially uniform coverage. Kapoutsis et al. addressed homogeneous multi-robot coverage by introducing the Divide Areas Algorithm for Optimal Multi-Robot Coverage Path Planning (DARP), which partitions the environment into balanced, connected regions based on the fleets' initial positions. The planner then applies the spanning tree coverage (STC) algorithm within each region to generate optimal, non-backtracking paths [12]. In [13] Kazemdehbashi proposed

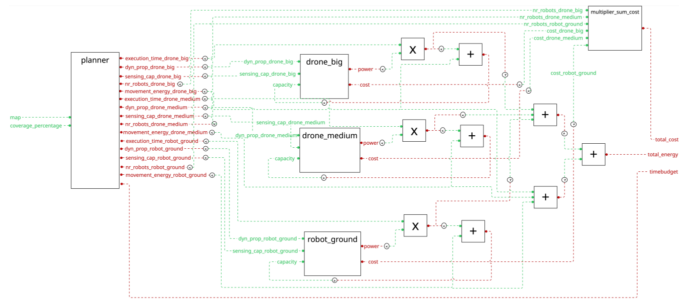

Fig. 1. Co-design diagram for heterogeneous multi-robot coverage, modeling fleet composition, robot design, and planner choice. Green wires denote functionalities (map, i.e. the search space, and coverage percentage), while red wires denote resources (cost, energy, and time). The objective is to derive Pareto optimal fleet designs, satisfying coverage requirements while minimizing resource usage.

the adaptive grid-based decomposition (AGD) algorithm for heterogeneous Unmanned Aerial Vehicle (UAV) fleets in maritime SAR, where grid cell sizes are adapted to the UAVs' sensing footprints, and coupled it with a mixed-integer programming (MIP) model to generate coverage paths. In [9, 10], Zardini and Censi introduce *Co-Design*, a framework that formalizes each subsystem as a design problem which requires resources and provides functionalities. These systems can be interconnected to form larger problems, allowing multiple subsystems to be jointly optimized, yielding Pareto optimal trade-offs across complex systems. Co-Design has proven effective in domains such as in Formula 1, balancing physical configuration and race tactics [14], or in the design of autonomous vehicles (AVs), where task-driven Co-Design jointly optimizes hardware, perception, planning, and computation under resource and performance constraints [15].

II. CO-DESIGN FRAMEWORK

While this work focuses on heterogeneous multi-robot fleets for coverage tasks, the underlying modeling framework is general and applicable to a broader class of multi-agent systems and domain tasks. Fig. 1 illustrates the Co-Design diagram of our framework. In this representation, green wires denote functionalities and red wires denote resources that each block respectively provides or requires. The diagram, together with Co-Design's optimization algorithm [10], determines optimal fleet configurations, including robot types, design parameters, and planner choice, that achieve a required coverage level while minimizing cost, energy, and time. Accordingly, the functionalities are map and coverage_percentage, and the resources are total_cost, total_energy, and timebudget. Querying the diagram with a *FixFunMinRes* query [10] (e.g., requiring a coverage percentage of at least 75% on a given search space) yields the complete set of Pareto optimal design solutions together with their associated trade-offs. The following subsections introduce the constituent blocks, i.e., the individual design problems, that compose our framework.

Robot modeling – A fleet is defined as a combination of robot types, each specified by the number of units per type. While all types share a common interface, their capabilities differ: they provide dynamic properties, battery capacity, and sensing capabilities, and require cost and power. Unlike approaches that assume fixed robot designs, we explicitly optimize the internal design of each type. A robot is modeled as three interconnected modules - battery, actuation, and sensing - each contributing design variables that the optimizer can tune to explore trade-offs. The actuation module imposes a payload limit, which in turn creates a feedback loop as the total weight depends on the combined weight of all modules. This ensures that any chosen configuration is dynamically feasible. Similarly, the total cost and idle power consumption are obtained by summing contributions across modules. The instantaneous power consumption is modeled as $P(t) = P_{idle} + c_{vel}v(t)^2 + c_{acc}|a(t)|$, where $v(t)$ and $a(t)$ denote velocity and acceleration and c_{vel} and c_{acc} are constants of the actuation module. This formulation defines the inner optimization problem: *For each robot type, determine the set of feasible and optimal design choices as well as their associated trade-offs, given the task requirements.*

Planner modeling – The planner design problem requires the fleet configuration (number of robots per type and configuration), the energy required by each type to follow the assigned plan, the execution time, and a global time budget for task completion. In return, the planner provides the achieved coverage percentage of the search space, defined as the union of all areas swept by the robots' sensing footprints. Connecting the planner with the robot design problems forms the outer optimization loop. The planner depends on robot properties and imposes energy and execution-time requirements, ensuring that each robot can feasibly complete its assigned trajectory. This introduces a second feedback loop: The total energy required by a robot is given by the movement energy plus idle power integrated over execution time, ensuring that the chosen battery capacity is sufficient to finish the task. Finally, the total cost and total energy for completing the mission are obtained by summing the contributions of all robots, yielding the global resource measures.

III. Experimental Results

Experimental setup – We evaluate our framework in a rectangular search space of 400m x 800m, considering three robot types (big drone, medium drone, and ground robot), with up to five units of each type available for fleet composition. Design variations are restricted to the actuation and battery modules, while sensing is fixed to one option per type. Battery configurations are adopted from [10], and the set of available actuation/sensing modules is summarized in Table I.

The planner block is populated with options by running simulations across all robot and fleet configurations using the AGD planner (with an adaptation of the MIP approach with a Christofides traveling salesman problem (TSP) approximation [16]), the MRTA planner, and the DARP planner introduced in Section I. To ensure physical feasibility, a motion planner is applied on top of each coverage plan, enforcing the dynamic

Type	Actuation module (v, c_{vel}, c_{acc}; cost, P_{idle})	Sensing module (radius, cost, P_{idle})
Big drone	20, 10, 20; 4000, 250 23, 14, 22; 4200, 280	50, 200, 2
Medium drone	10, 4, 7; 500, 22 14, 5, 9; 520, 24	12, 100, 1
Ground robot	5, 1, 3; 4000, 250 7, 2, 5; 4000, 250	30, 100, 15

TABLE I
AVAILABLE MODULES IN THE EXPERIMENTS. $v : m/s$, $c_{vel} : Ws^2/m^2$, $c_{acc} : Ws^2/m$, $cost : USD$, $P_{idle} : W$, $radius : m$.

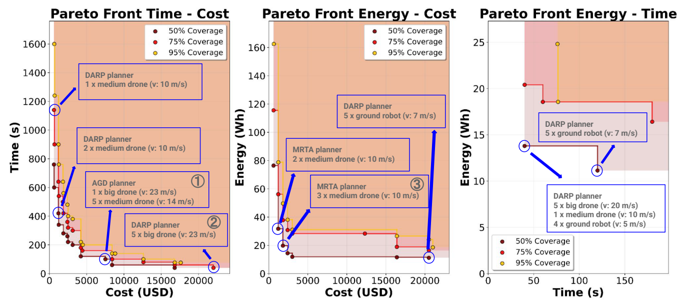

Fig. 2. Pareto fronts resulting from querying our framework to cover a rectangular search space of size 400m x 800m with coverage of 50% (dark red), 75% (red) and 95% (yellow), respectively.

constraints of the robots and verifying that the resulting trajectories are executable in practice.

Results — With the models and design choices fully instantiated, we queried the framework to obtain Pareto optimal solutions for coverage requirements of 50%, 75%, and 95%. The resulting Pareto fronts, shown in Fig. 2, highlight the trade-offs among the requirements. As expected, allocating five large drones achieves rapid coverage but incurs very high cost (2), whereas relying on a single drone significantly reduces cost but leads to longer completion times. At 1, the AGD planner yields the optimal solution, consistent with its design for heterogeneous fleets, while at 2 DARP is selected due to the homogeneous fleet composition. A further notable point is 3, where the MRTA planner along with two medium drones in their less powerful variant are chosen, illustrating its energy-awareness through slower prescribed speeds. Overall, the results demonstrate how the framework enables systematic exploration of trade-offs: Determining when powerful robots are needed, when cheaper designs suffice, and where similar performance can be achieved at far lower cost.

IV. Discussion and Conclusion

In this work, we present a Co-Design framework for analyzing heterogeneous multi-agent robotic systems in search tasks. Our framework adopts a holistic perspective by jointly considering planner selection, fleet composition, and robot design. Thus, our approach provides a structured exploration of the trade-offs associated with each design decision. The results demonstrate the framework's effectiveness in supporting informed reasoning about system-level design choices. Future work aims to incorporate more realistic sensing modules, explicitly modeling uncertainty, and relaxing restrictions on the search space geometry. Moreover, the framework can be further exploited by expanding the query space to include more diverse scenarios, robot types, and module configurations.

References

[1] Randa Almadhoun, Tarek Taha, Lakmal Seneviratne, and Yahya Zweiri. A survey on multi-robot coverage path planning for model reconstruction and mapping. *SN Applied Sciences*, 1(8):847, 2019.

[2] Peter R Wurman, Raffaello D'Andrea, and Mick Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI magazine*, 29(1):9–9, 2008.

[3] Trevor Standley. Finding optimal solutions to cooperative pathfinding problems. In *Nat. Conf. on Artificial Intelligence (AAAI)*, volume 24, pages 173–178, 2010.

[4] Enric Galceran and Marc Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61:1258–1276, 12 2013. doi: 10.1016/j.robot.2013.09.004.

[5] Pal-Stefan Murvay and Ioan Silea. A survey on gas leak detection and localization techniques. *Journal of Loss Prevention in the Process Industries*, 25:966–973, 11 2012. doi: 10.1016/j.jlp.2012.05.010.

[6] Luis Merino, Fernando Caballero, J. Ramiro Martinez-de Dios, Ivan Maza, and Anibal Ollero. An unmanned aircraft system for automatic forest fire monitoring and measurement. *Journal of Intelligent and Robotic Systems*, 65:533–548, 01 2012. doi: 10.1007/s10846-011-9560-x.

[7] Andreas Breitenmoser, Fabien Tâche, Gilles Caprari, Roland Siegwart, and Roland Moser. Magnebike—toward multi climbing robots for power plant inspection. pages 1713–1720, 01 2010. doi: 10.1145/1838194.1838200.

[8] Wesley Debusk. Unmanned aerial vehicle systems for disaster relief: Tornado alley. 04 2009. doi: 10.2514/6.2009-2045.

[9] Gioele Zardini. *Co-Design of Complex Systems: From Autonomy to Future Mobility Systems*. PhD thesis, ETH Zürich, Switzerland, 2023. URL https://www.research-collection.ethz.ch/handle/20.500.11850/648075. Awarded the ETH Medal for Outstanding Doctoral Thesis.

[10] Andrea Censi. A mathematical theory of co-design. *CoRR*, abs/1512.08055, 2015. URL http://arxiv.org/abs/1512.08055.

[11] Jiaheng Hu, Howard Coffin, Julian Whitman, Matthew Travers, and Howie Choset. *Large-Scale Heterogeneous Multi-robot Coverage via Domain Decomposition and Generative Allocation*, pages 52–67. 12 2022. ISBN 978-3-031-21089-1. doi: 10.1007/978-3-031-21090-7_4.

[12] Athanasios Ch. Kapoutsis, Savvas A. Chatzichristofis, and Elias B. Kosmatopoulos. Darp: Divide areas algorithm for optimal multi-robot coverage path planning. *Journal of Intelligent & Robotic Systems*, 86(3):663–680, 2017.

[13] Sina Kazemdehbashi. Adaptive grid-based decomposition for uav-based coverage path planning in maritime search and rescue, 2024. URL https://arxiv.org/abs/2412.00899.

[14] Marc-Philippe Neumann, Giona Fieni, Francesca Furia, Alberto Cerofolini, Vittorio Ravaglioli, Christopher H. Onder, and Gioele Zardini. Strategic co-design in formula 1: Balancing physical configuration and race tactics. 2025. URL https://www.research-collection.ethz.ch/handle/20.500.11850/708582.

[15] Dejan Milojevic, Gioele Zardini, Mirian Elser, Andrea Censi, and Emilio Frazzoli. Codei: Resource-efficient task-driven co-design of perception and decision making for mobile robots applied to autonomous vehicles. *IEEE Transactions on Robotics*, 41:2727–2748, 2025. doi: 10.1109/TRO.2025.3552347. URL https://ieeexplore.ieee.org/document/10930687. Additionally Selected for Presentation at IEEE/RSJ IROS 2025.

[16] Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report 388, Graduate School of Industrial Administration, Carnegie Mellon University, 1976.