Scalable Coalition Formation for Extremely Large Collectives

Doctoral Consortium

Neha G. Pusalkar
Oregon State University
Corvallis, USA
pusalkan@oregonstate.edu

ABSTRACT

Coalition formation facilitates coordination of robotic collectives by partitioning robots into task-oriented teams. The combinatorial nature of this problem presents a key challenge when developing algorithms for extremely large scale collectives containing thousands of robots. Deployment in real-world environments exacerbates this complexity by requiring adaptation to dynamic events. This research develops scalable collective coalition formation algorithms for heterogeneous collectives by leveraging the collectives' compositional properties and a game-theoretic framework. Future research proposes to extend these methods to real-world settings by implementing a coalition reformation algorithm that restructures coalitions to address dynamic events. The proposed algorithmic framework demonstrates scalability and enables adaptability, making it viable for real-world missions.

KEYWORDS

Robotic Collectives; Coalition Formation; Hedonic Games

ACM Reference Format:
Neha G. Pusalkar. 2026. Scalable Coalition Formation for Extremely Large Collectives: Doctoral Consortium. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/MGGH7353

1 INTRODUCTION

The fundamental bottleneck in deploying Extremely Large Scale Collectives (ELSCs) stems from the coalition formation and reformation problem's combinatorial complexity. Real-world domains require coalition restructuring in response to dynamic events, but often lack reliable communication networks essential for effective coordination. Distributed approaches (e.g., auctions) coordinate via extensive message passing, but incur significant communication overhead for collectives [6]. Mass produced robots restrict heterogeneity in ELSCs to a small number of unique robot types relative to the collective size, creating an inherent redundancy that can reduce inter-robot communication during coalition formation and reformation. This research leverages this type-based redundancy to develop a distributed and scalable coalition formation algorithm for ELSCs operating under communication constraints. The proposed research will develop a novel distributed algorithm that enables priority-driven coalition reformation to ensure long-duration autonomous operations in dynamic real-world domains.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/MGGH7353

2 RELATED WORK

Existing coalition formation approaches that utilize greedy heuristics [14], optimization techniques [2], market-based auctions [7], and game-theoretic methods [1] scale poorly with collective size [6]. The Group Agent Partitioning and Placing Event with the services extension (GRAPE-S) produces near-optimal solutions for heterogeneous multi-capability collectives with up to 1,000 robots using an anonymous hedonic game [4]. Robots maintain beliefs about all robots' task assignments and join individually profitable coalitions based on the coalition size. GRAPE-S's global belief synchronization necessitates extensive inter-robot communication, which cannot be guaranteed in communication-constrained environments. Centralized hedonic approaches rapidly produce solutions for 2,000 robots, but centralized decision making creates a single point of failure [3]. Additionally, these methods do not support dynamic real-world environments, limiting their applicability to domains where initial robot allocation remains valid throughout the mission.

3 COALITION FORMATION IN ELSCS

ELSC's limited heterogeneity allows adoption of a type-based representation [12, 13]. This research develops LeaderGRAPE-S that associates a leader robot with unique robot types based on robot capabilities, and decomposes the coalition formation problem into multiple leader-centric subproblems [9]. Each leader independently enables belief negotiations among its respective follower subgroup (i.e., robots sharing the leader's capabilities) using GRAPE-S. This decomposition reduces the number of robots participating in each hedonic game, lowering inter-robot negotiation and improving computational and communication efficiency relative to GRAPE-S.

LeaderGRAPE-S's total communication requirement is too high for communication-constrained domains. Leader Negotiation with GRAPE-S (LN-GRAPE-S) addresses this limitation by formulating a hedonic game over a set of leaders ($\mathcal{L}$) [10]. Each leader maintains a belief about each tasks' assigned capabilities and allocates the required followers to tasks yielding the highest marginal utility at each iteration. Leaders broadcast their updated belief states and synchronize their beliefs to a state that has been updated the most, or the most recently updated state. Belief updates and synchronization continue until all leaders are satisfied with their follower's task allocations. LN-GRAPE-S constitutes a finite potential game, guaranteeing Nash stability [10]. LN-GRAPE-S's belief synchronization is abstracted to the leader-level because robots of the same type provide identical marginal contributions, effectively eliminating individual robot-level negotiations [11].

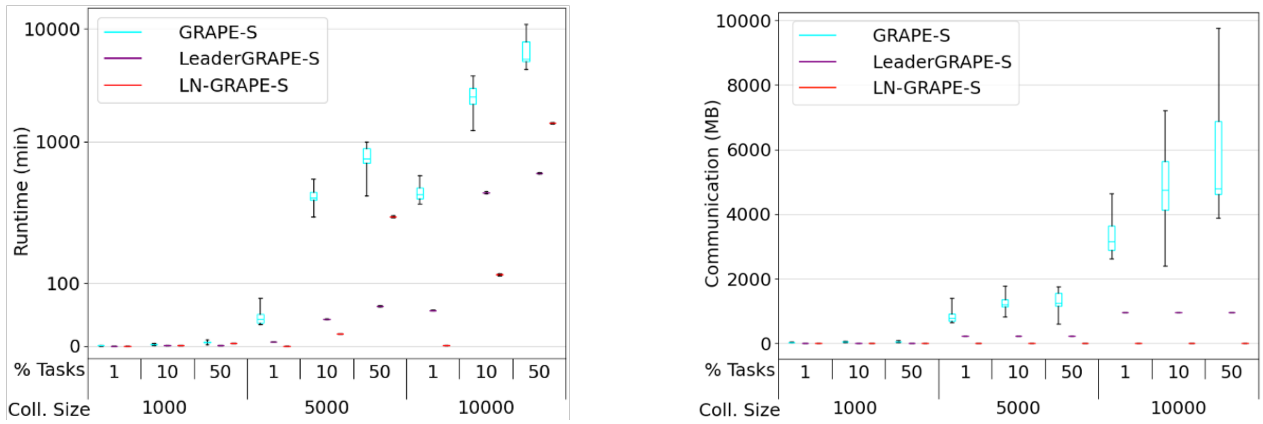

(a) Algorithmic Runtime

(b) Total Communication

Figure 1: Algorithmic runtime (a) and total communication requirement (b) for GRAPE-S, LeaderGRAPE-S, and LN-GRAPE-S, evaluated on ELSCs containing 1,000, 5,000, and 10,000 robots, with each robot providing two capabilities. % Tasks indicates number of tasks expressed as a percentage of the collective size.

4 PRELIMINARY RESULTS

LeaderGRAPE-S and LN-GRAPE-S outperformed GRAPE-S [5] in algorithmic runtime and total communication requirement for collectives with 10 capabilities, where each robot provided two capabilities (Figure 1). LeaderGRAPE-S was on average 11 times faster than GRAPE-S for 10,000 robot collectives with 50% tasks (i.e., 5000 tasks) [9]. The reduced inter-robot negotiations and problem decomposition contribute to LeaderGRAPE-S's faster runtimes. LN-GRAPE-S was on average 7 times faster for 1% and 10% tasks, and on average 3 times slower for 50% tasks relative to LeaderGRAPE-S [10]. LN-GRAPE-S's slower runtimes at higher task percentages result from fewer followers being allocated per iteration due to smaller coalitions, requiring more iterations to converge.

LeaderGRAPE-S's communication overhead was on average 7 times lower than GRAPE-S's, with the maximum total communication requirement being 10 times lower. LN-GRAPE-S's total communication requirement was on average 825 times lower than GRAPE-S, and 121 times lower than LeaderGRAPE-S. LN-GRAPE-S's maximum communication requirement (6.55 MB) was 1490 times and 147 times lower than the maximum values for GRAPE-S (9762.70 MB) and LeaderGRAPE-S (960.87 MB), respectively, for 10,000 robot collectives with 50% tasks [10]. This substantial reduction in LN-GRAPE-S's communication overhead is due to limiting the belief synchronization to a significantly smaller set of leader robots, rather than the entire collective of size N ($|\mathcal{L}| \ll N$), and to using batch allocation, which reduces both message size and the total number of broadcast messages.

LeaderGRAPE-S and LN-GRAPE-S produced optimal solutions for 100% and 96% of all trials, respectively. Problems with fewer optimal allocation options (i.e., tasks with perfectly matching capability deficits) generated suboptimal solutions, with the minimum utility solution exceeding the 50% suboptimality bound established by the comparison algorithm's Nash stable partition [8].

5 PROPOSED RESEARCH

LN-GRAPE-S demonstrates that the leader-level abstraction enables scalability while incurring a low communication overhead. Translating this efficiency to real-world domains requires assessing the algorithm's efficiency in realistic distributed settings. Following this validation, LN-GRAPE-S will be extended to incorporate a coalition restructuring mechanism to enable adaptation to dynamic events including new tasks, changing task requirements, and robot failures. This proposed algorithm will develop a preemptive and priority-driven coalition reformation framework that integrates the leader-follower hierarchy with a multi-objective optimization problem. Leaders (and their respective followers) providing the new task's (i.e., the task triggering the dynamic event) capabilities will be considered for reformation. Leaders will locally evaluate followers by lexicographically ranking them from least to most disruptive to reallocate, utilizing real-world robot and environment characteristics as optimization objectives. These objectives are ranked in decreasing order of importance as current task priority, task execution stage (e.g., interruptible, non-interruptible, resumable), inter-task dependencies, and proximity to the new task. The leaders will negotiate among themselves to determine the number of followers each leader must reallocate. Finally, each leader will reallocate the determined number of followers starting with the first in its ordered list. The proposed coalition reformation approach will be evaluated based on the number of new tasks successfully assigned, as well as the number and priority of the existing tasks disrupted. The leader-follower abstraction will maintain scalability, while the multi-objective formulation will ensure minimal disruption of the existing coalition structure during reallocation. The proposed scalable coalition reformation framework paves the way for deploying ELSCs in dynamic and unstructured domains, where long-term autonomy is critical for effective operation.

ACKNOWLEDGMENTS

This effort was supported by ONR grant N00014-23-1-2171. The views, opinions, and findings are those of the authors and do not represent the official views or policies of ONR.

REFERENCES

[1] Vittorio Bilò, Gianpiero Monaco, and Luca Moscardelli. 2022. Hedonic games with fixed-size coalitions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 9287–9295.

[2] Lei Cao, He shun Tan, Hui Peng, and Ming cong Pan. 2014. Multiple UAVs hierarchical dynamic task allocation based on PSO-FSA and decentralized auction. In *IEEE International Conference on Robotics and Biomimetics*. IEEE, 2368–2373.

[3] Emily Czarnecki and Ayan Dutta. 2021. Scalable hedonic coalition formation for task allocation with heterogeneous robots. *Intelligent Service Robotics* 14, 3 (2021), 501–517.

[4] Grace Diehl. 2023. *Collective coalition formation*. Ph.D. Dissertation. Oregon State University.

[5] Grace Diehl and Julie A Adams. 2023. GRAPE-S: Near real-time coalition formation for multiple service collectives. *arXiv preprint arXiv:2310.12480* (2023).

[6] Grace Diehl and Julie A Adams. 2024. The viability of domain constrained coalition formation for robotic collectives. *Swarm Intelligence* (2024), 1–42.

[7] Muhammad Irfan and Adil Farooq. 2016. Auction-based task allocation scheme for dynamic coalition formations in limited robotic swarms with heterogeneous capabilities. In *International Conference on Intelligent Systems Engineering*. IEEE, 210–215.

[8] Inmo Jang, Hyo-Sang Shin, and Antonios Tsourdos. 2018. Anonymous hedonic game for task allocation in a large-scale multiple agent system. *IEEE Transactions on Robotics* 34, 6 (2018), 1534–1548.

[9] Neha G. Pusalkar and Julie A. Adams. 2025. Leader-based coalition formation for extremely large scale collectives. In *IEEE Conference on Artificial Intelligence*. 701–706.

[10] Neha G. Pusalkar and Julie A. Adams. 2026. Extremely Large Collective Coalition Formation: Scalability. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems*.

[11] Tammar Shrot, Yonatan Aumann, and Sarit Kraus. 2010. On agent types in coalition formation problems. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. 757–764.

[12] Suguru Ueda, Atsushi Iwasaki, Vincent Conitzer, Naoki Ohta, Yuko Sakurai, and Makoto Yokoo. 2018. Coalition structure generation in cooperative games with compact representations. *Autonomous Agents and Multi-Agent Systems* 32 (2018), 503–533.

[13] Suguru Ueda, Makoto Kitaki, Atsushi Iwasaki, and Makoto Yokoo. 2011. Concise characteristic function representations in coalitional games based on agent types. In *International Joint Conference on Artificial Intelligence*, Vol. 11. 393–399.

[14] Wanqing Zhao, Qinggang Meng, and Paul WH Chung. 2015. A heuristic distributed task allocation method for multivehicle multitask problems and its application to search and rescue scenario. *IEEE transactions on cybernetics* 46, 4 (2015), 902–915.