Multiagent Matroid Upgrading: Greedy is Fair and Efficient

Extended Abstract

Qingwen Ma
East China Normal University
School of Software Engineering
Shanghai, China
51265902177@stu.ecnu.edu.cn

Chao Peng
East China Normal University
School of Software Engineering
Shanghai, China
cpeng@sei.ecnu.edu.cn

Changfeng Xu
East China Normal University
School of Software Engineering
Shanghai, China
51265902120@stu.ecnu.edu.cn

Chenyang Xu
East China Normal University
School of Software Engineering
Shanghai, China
cyxu@sei.ecnu.edu.cn

Ruilong Zhang
City University of Hong Kong
(Dongguan)
Department of Computer Science
Dongguan, China
ruilong.zhang@cityu-dg.edu.cn

ABSTRACT

This paper introduces a general multiagent matroid upgrading problem that models a broad class of real-world resource allocation tasks. In this setting, there are multiple agents and a ground set of elements, where each element is assigned to a specific agent and has two associated costs: a default cost and a reduced (upgraded) cost. Upgrading an element lowers its cost to the upgraded value, while non-upgraded elements retain their default costs. Each agent is associated with its own matroid, with the goal of finding a minimum-cost basis. The central task is to select at most k elements to upgrade so as to minimize a non-decreasing convex function over the agents' minimum basis costs, capturing both efficiency and fairness objectives in multiagent systems.

We show that the problem is polynomial-time solvable and that an optimal solution can be obtained via a simple greedy algorithm. Our analysis exploits the structural properties of matroids to establish the existence of optimal substructures, thereby ensuring that greedy upgrading yields optimal outcomes. Building on this insight, we can further extend our result to more general settings, such as scenarios with interval fairness constraints, where the number of elements upgraded for each agent is required to lie within a specified interval.

KEYWORDS

Matroid upgrading; Multiagent systems; Greedy algorithms; Fair resource allocation

ACM Reference Format:
Qingwen Ma, Chao Peng, Changfeng Xu, Chenyang Xu, and Ruilong Zhang. 2026. Multiagent Matroid Upgrading: Greedy is Fair and Efficient: Extended Abstract. 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/LREX9402

All authors (ordered alphabetically) have equal contributions and are corresponding authors.

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/LREX9402

1 INTRODUCTION

The allocation of limited resources among multiple agents is a fundamental problem in multi-agent systems, with applications ranging from communication networks [1] and distributed computing [15] to robotics [16] and transportation [14]. In such scenarios, resources are typically scarce, and different agents may face their own structural constraints on how resources can be utilized, e.g., maintaining feasibility with respect to connectivity, capacity, or diversity requirements [4]. At the same time, the system usually needs to consider both efficiency, which seeks to maximize overall benefit, and fairness, which aims to balance the outcomes among agents [11]. These challenges motivate the development of models that integrate efficiency- and fairness-oriented objectives with combinatorial constraints in a multi-agent resource allocation framework.

In this paper, we study a general class of *matroid upgrading* problems in multiagent systems, which captures a variety of applications such as network upgrading and neural network compression. In the model, there is a ground set of elements E, where each element $e \in E$ is associated with two non-negative costs: a default cost $\hat{c}(e)$ and an upgraded cost $\check{c}(e)$, with $\hat{c}(e) \geq \check{c}(e)$. Initially, each element incurs its default cost. However, if selected for upgrading, its cost is reduced to the upgraded value.

The elements are partitioned into n disjoint groups $\{E^{(i)}\}_{i\in[n]}$, with each group assigned to an agent. Each agent $i \in [n]$ is associated with a matroid[1] $\mathcal{M}^{(i)} = (E^{(i)}, \mathcal{I}^{(i)})$, and the objective is to find a minimum-cost basis within the matroid. We can select a subset $S \subseteq E$ of at most k elements to upgrade, with the objective of minimizing $\sum_{i\in[n]} F_i\big(\delta_S(\mathcal{M}^{(i)})\big)$, where $\delta_S(\mathcal{M}^{(i)})$ denotes the minimum cost of a basis in $\mathcal{M}^{(i)}$ given that only elements in S are upgraded, and each F_i is a non-decreasing convex function. Formally, the multiagent matroid upgrading problem (MMUP) can be

[1] A matroid is a set system $(E, \mathcal{I})$ with $\mathcal{I} \subseteq 2^E$ such that (i) $\emptyset \in \mathcal{I}$; (ii) for each $S \in \mathcal{I}$, all S's subsets are also in $\mathcal{I}$; (iii) If $A, B \in \mathcal{I}$ with $|A| < |B|$, then $\exists e \in B \setminus A$ such that $\{e\} \cup A \in \mathcal{I}$.

written as:

$$\min_{S \subseteq E, |S| \leq k} \sum_{i \in [n]} F_i \left(\delta_S(\mathcal{M}^{(i)}) \right)$$
$$\text{s.t.} \quad \delta_S(\mathcal{M}^{(i)}) = \min_{B_i \in \mathcal{B}^{(i)}} \sum_{e \in B_i} c_S(e) \quad \forall i \in [n] \,, \tag{MMUP}$$

where $\mathcal{B}^{(i)}$ denotes the set of bases of matroid $\mathcal{M}^{(i)}$, and the cost $c_S(e)$ is defined as $c_S(e) = \check{c}(e)$ if $e \in S$, and $c_S(e) = \hat{c}(e)$ otherwise. We remark that the objective function captures a variety of aggregation goals, such as the weighted sum of agents' utilities when overall efficiency is prioritized, or the ℓ_q norm of the utilities when fairness across agents is the primary concern.

2 RELATED WORK

There are several other upgrade models studied in the literature. For example, the scheduling with testing introduced by [8] is closely related to the upgrading framework considered in this paper. Recent works [2, 3, 9] further explore this model. Among them, the most closely related is [7], which studies a single-machine scheduling problem where each job has an upper and lower processing time, and testing a job reveals its lower bound. This mirrors the upgrade operation in our setting. The goal is to select at most k jobs to test in order to minimize the total completion time. They present an FPTAS for this problem. However, due to the different nature of constraints, the techniques used in our work diverge substantially from theirs.

Other related work focuses on optimization in network upgrades. For instance, [5] investigates how to efficiently find alternative edges when the cost of a network link changes, and [13] studies dynamic algorithms for maintaining a minimum spanning tree as the graph evolves. Notably, these works are concerned with adapting to changes in edge costs or network structure after upgrades have occurred. In contrast, our model focuses on deciding which edges to upgrade in order to optimize a global objective.

3 MAIN RESULTS

We formalize the *multi-agent matroid upgrading problem* (MMUP), which integrates efficiency- and fairness-driven objectives with combinatorial constraints arising from matroid structures. While resource allocation in multi-agent systems is generally challenging, we identify a structured subclass grounded in matroid theory that admits efficient and provably optimal solutions via a simple greedy approach.

Main Theorem. *Given any MMUP instance, there exists a general greedy algorithm that computes an optimal solution in polynomial time.*

Our main contribution is a polynomial-time greedy algorithm that solves MMUP optimally. The algorithm is natural and simple: we begin by assuming that all elements are upgraded and compute the minimum-cost basis for each agent. The union of all these bases forms a candidate element set. Then, we iteratively select the element from this candidate set that results in the largest decrease in the leader's objective and include it in the upgrade set, continuing this process until k elements have been selected. Although the algorithm is conceptually simple, proving its optimality is nontrivial. The correctness relies on a crucial structured property of optimal solutions, which we refer to as the *nestedness property*. This property ensures that for any optimal solution involving $k-1$ upgrades, there always exists an optimal solution with k upgrades that contains the former as a subset.

By leveraging the nestedness property, we further show that our results extend to several more general settings. For instance, the objective can be replaced with a minimax form to capture worst-case efficiency among agents, i.e., $\max_{i\in[n]} F_i(\delta_S(\mathcal{M}^{(i)}))$, Alternatively, one can incorporate recently popular interval fairness constraints [6, 12] which require that the number of upgraded elements for each follower must lie within a specified interval. We prove that in both of these generalized models, the greedy algorithm remains optimal.

Connection to Budget-Constrained MST. As a byproduct, we find that our model captures the special case of the budget-constrained minimum spanning tree (MST) problem with $\{0,1\}$ edge weights. This classical problem aims to select a spanning tree of minimum total cost subject to a weight budget constraint. It is known to be NP-hard in general, and existing approaches are based on LP techniques, including Lagrangian relaxation [17] and LP rounding [10]. In contrast, our algorithm is purely combinatorial and implies that the $\{0,1\}$-weight case can be solved in polynomial time. To the best of our knowledge, this positive result was not known before.

4 CONCLUSION

In this paper, we introduce a general multiagent matroid upgrading problem that models a wide range of practical network upgrading scenarios. We prove that a simple greedy algorithm solves this problem optimally. Furthermore, our results extend naturally to more general settings, including those incorporating fairness constraints.

Our work opens several directions for future research. For example, it would be interesting to study a more general setting where the upgrade quota varies for each element. Notably, this setting generalizes the knapsack problem and is therefore NP-hard. Investigating whether greedy algorithms can still provide strong approximation guarantees in this context remains an interesting open problem.

ACKNOWLEDGMENTS

This work is supported by the National Key Research and Development Program of China (2023YFA1009402, 2025YFC2423000), NSFC Programs (62302166, 62161146001, 62372176), Shanghai Key Lab of Trustworthy Computing, Shanghai Frontiers Science Center of Molecule Intelligent Syntheses and Fundamental Research Funds for the Central Universities.

REFERENCES

[1] Manzoor Ahmed, Noor Fatima, Salman Raza, Hamid Ali, Abdul Qayum, Wali Ullah Khan, Muhammad Sheraz, and Teong Chee Chuah. 2025. Optimizing Resource Allocation and Task Offloading in Multi-UAV MEC Networks. *IEEE Access* 13 (2025), 68710–68725. https://doi.org/LREX9402

[2] Susanne Albers and Alexander Eckl. 2020. Explorable Uncertainty in Scheduling with Non-uniform Testing Times. In *WAOA (Lecture Notes in Computer Science, Vol. 12806)*. Springer, 127–142.

[3] Susanne Albers and Alexander Eckl. 2021. Scheduling with Testing on Multiple Identical Parallel Machines. In *WADS (Lecture Notes in Computer Science, Vol. 12808)*. Springer, 29–42.

[4] Sanae Amani and Christos Thrampoulidis. 2021. Decentralized Multi-Agent Linear Bandits with Safety Constraints. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*. AAAI Press, 6627–6635. https://doi.org/LREX9402

[5] David A. Bader and Paul Burkhardt. 2022. A Simple and Efficient Algorithm for Finding Minimum Spanning Tree Replacement Edges. *J. Graph Algorithms Appl.* 26, 1 (2022), 577–588.

[6] Shuang Cui, Kai Han, Shaojie Tang, Feng Li, and Jun Luo. 2024. Fairness in Streaming Submodular Maximization Subject to a Knapsack Constraint. In *KDD*. ACM, 514–525.

[7] Christoph Damerius, Peter Kling, Minming Li, Chenyang Xu, and Ruilong Zhang. 2023. Scheduling with a Limited Testing Budget: Tight Results for the Offline and Oblivious Settings. In *ESA (LIPIcs, Vol. 274)*. 38:1–38:15.

[8] Christoph Dürr, Thomas Erlebach, Nicole Megow, and Julie Meißner. 2020. An Adversarial Model for Scheduling with Testing. *Algorithmica* 82, 12 (2020), 3630–3675.

[9] Mingyang Gong, Randy Goebel, Guohui Lin, and Eiji Miyano. 2022. Improved approximation algorithms for non-preemptive multiprocessor scheduling with testing. *Journal of Combinatorial Optimization* 44, 1 (2022), 877–893.

[10] Fabrizio Grandoni, R. Ravi, Mohit Singh, and Rico Zenklusen. 2014. New approaches to multi-objective optimization. *Math. Program.* 146, 1-2 (2014), 525–554.

[11] Niko A. Grupen, Bart Selman, and Daniel D. Lee. 2022. Cooperative Multi-Agent Fairness and Equivariant Policies. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*. AAAI Press, 9350–9359. https://doi.org/LREX9402

[12] Marwa El Halabi, Federico Fusco, Ashkan Norouzi-Fard, Jakab Tardos, and Jakub Tarnawski. 2023. Fairness in Streaming Submodular Maximization over a Matroid Constraint. In *ICML (Proceedings of Machine Learning Research, Vol. 202)*. PMLR, 9150–9171.

[13] Mao Luo, Huigang Qin, Xinyun Wu, Caiquan Xiong, Dahai Xia, and Yuanzhi Ke. 2024. Efficient Maintenance of Minimum Spanning Trees in Dynamic Weighted Undirected Graphs. *Mathematics* 12, 7 (2024), 1021.

[14] Gunasekaran Manogaran, Jiechao Gao, and Tu N. Nguyen. 2023. Optimizing Resource and Service Allocations for IoT-Assisted Intelligent Transportation Systems. *IEEE Trans. Intell. Transp. Syst.* 24, 11 (2023), 12877–12887. https://doi.org/LREX9402

[15] Elyas Oustad, Abolfazl Younesi, Mohsen Ansari, Sepideh Safari, Mohammad Arman Soleimani, Jörg Henkel, and Alireza Ejlali. 2025. DIST: Distributed Learning-Based Energy-Efficient and Reliable Task Scheduling and Resource Allocation in Fog Computing. *IEEE Trans. Serv. Comput.* 18, 3 (2025), 1336–1351. https://doi.org/LREX9402

[16] Roopsi Rathi, Saurav Dixit, Shweta Bansal, Kaushal Kumar, Natalia Taskaeva, Alexander Yu. Tumanov, and Vinod John. 2022. Stackelberg game approach for resource allocation in device-to-device communication with heterogeneous networks. *Robotics Auton. Syst.* 156 (2022), 104222. https://doi.org/LREX9402

[17] R. Ravi and Michel X. Goemans. 1996. The Constrained Minimum Spanning Tree Problem (Extended Abstract). In *SWAT (Lecture Notes in Computer Science, Vol. 1097)*. Springer, 66–75.