$\mathcal{H}$-EFP: Bridging Efficiency in Multi-Agent Epistemic Planning with Heuristics

Francesco Fabiano[1][0000−0002−1161−0336], Theoderic Platt[1], Son Tran[1][0000−0003−3689−8433], and Enrico Pontelli[1][0000−0002−7753−1737]

New Mexico State University, Las Cruces NM 88003, USA
`{ffabiano,theop,stran,epontell}@nmsu.edu`

Abstract. In this paper, we propose $\mathcal{H}$-EFP, a Portfolio-like *multi-agent epistemic planning* solver that demonstrates scalability potential and tangible performance improvements compared to state-of-the-art epistemic multi-agent planning systems. Ultimately, our goal is to broaden the practical application of multi-agent epistemic reasoning in real-world scenarios by reducing resource demands, potentially enabling its use in modeling situations involving multiple entities sharing information, such as autonomous driving.

Keywords: Multi-Agent Planning · Belief Manipulation · Heuristics.

1 Introduction

Artificial intelligence (AI) research has increasingly focused on planning systems where multiple agents must interact based on shared and private knowledge. Such multi-agent planning scenarios, particularly those involving epistemic reasoning—*i.e.*, reasoning about agents' beliefs and knowledge—are becoming critical for real-world applications such as autonomous driving, robotics, and collaborative AI systems.

However, existing *Multi-Agent Epistemic Planning* (MEP) systems face significant challenges in scalability due to the computational overhead associated with reasoning about agents' nested beliefs. The $m\mathcal{A}^*$ [2] action language is commonly used for representing epistemic planning problems, but solving such problems efficiently remains an open issue. This has led to the development of diverse techniques to reduce the computational load, such as limitations on the depth of nested beliefs [8], or the definition of heuristics [7]. Comprehensive insights into issues and research directions related to epistemic planning have been summarized in the report from a recent Dagstuhl's meeting [1].

Our contribution, the $\mathcal{H}$-EFP planner, builds on these ideas by exploiting diverse heuristics to show that significant gains in efficiency can be made through the use of heuristics, allowing for the practical use of MEP in real-world applications where computational resources are constrained.

2 Background

2.1 Multi-Agent Epistemic Planning

A multi-agent epistemic planning problem is defined as a tuple consisting of a set of fluents (the properties of the world), agents, possible actions, an initial state, and a goal state. Given this tuple we need to fina a *solution*, that is a sequence of actions that transforms the initial state into a goal state, while reasoning about the knowledge and beliefs of the agents involved. MEP extends traditional planning by incorporating epistemic actions, such as communication or observation, which can change what agents know or believe about the world. These actions are formalized in the $m\mathcal{A}^*$ [2] action language, which allows for reasoning about both ontic actions (those that change the state of the world) and epistemic actions (those that change the state of knowledge).

The main challenge in MEP is the sheer size of both the epistemic states, or e-states, and of the search space, the former which grows exponentially with the number of agents and fluents and the latter that increases exponentially with the number of actions.

2.2 Heuristics-Based Reasoning

As mentioned before, planning on multi-agent epistemic domains is a very resource-demanding task. That is why, even if optimizing the knowledge structures is essential, only focusing on such a task may never allow epistemic planners to become tools suited for real-life scenarios. For this reason, we decided to investigate alternative search strategies that may help in reducing the resources needed to solve MEP problems. In particular we focused on developing a MEP solvers that can exploit both the standard *Breadth-First Search* (**BFS**) and *Best-First Search*. These two search approaches are well-known in the planning community and, therefore, we will not provide any details on their implementation.

While **BFS** is an *uninformed* search—*i.e.*, it traverses the space using only information derived by the search-tree and not from the e-states themselves—Best-First Search selects, at each step, the *best* state, that is the one that is, supposedly, closest to the goal. The problem with this last approach lies in finding a good function to calculate the score of each e-state and, therefore, in understanding which e-state is the best one. These functions, known as *heuristics*, have been deeply studied in the planning community and are, nowadays, a standard concept [5, 9]. That is why, in this work, we decided to focus on formalizing some domain-independent heuristics for MEP.

3 $\mathcal{H}$-EFP

As main contribution we present $\mathcal{H}$-EFP, a comprehensive epistemic planner. This planner is heavily inspired by the systems presented in [4, 7]. While it shares with them the ability to comprehensively reason on the full extent of $m\mathcal{A}^*$, $\mathcal{H}$-EFP heavily exploits a Planning Graph-like data structure (referred to as e-PG), tailored for MEP, to derive heuristics. Our definition of e-PG builds on the version introduced by [7], but with significant modifications to the underlying knowledge representation, enabling our data structure to overcome fundamental issues in the original definition that prevented the resolution of certain problem types (*i.e.*, those with goals containing negated belief formulas). Due to space constraints, a formal definition of the e-PG, along with its components, demonstrations, and a graphical example illustrating its properties, can be found at the following link: https://francescofabiano.github.io/resources/.

Let us now present the specific heuristics that allow $\mathcal{H}$-EFP to evaluate the various e-states during the planning process.

- `SUB`: The first, not dependent from e-PG, simply associates a higher evaluation to e-states that satisfy more sub-goals. To improve this heuristic we defined functions that "break" complex goals into a conjunction of simpler ones.
- `C_PG`: This heuristics emulates the one presented in [7], inspired in turn by the classical Planning Graph. e-PG is used to derive the "importance" of each belief formula (its distance from the goal level) and then each e-state is characterized by the sum of the derived belief formulae scores. In particular `C_PG` reflects the *hAdd* heuristics in MEP; that is, for any belief formula ψ, its importance is calculated as the distance in terms of state levels from the initial state that verifies ψ to the state level that entails the goal. The lower the distance, the more important the formula is. This captures the idea that a formula entailed at a state level closer to the goal is more beneficial for reaching the goal, as it is harder to activate than others. We then sum all the formulae of interest that a state activates to evaluate it and we explore the one with the smallest sum.
- `L_PG`: This heuristic calculates the score of an e-state by constructing a Planning Graph from it (as initial state) and calculating the length—the shorter the better—of the constructed e-PG. If an e-PG cannot reach the goal from an e-state, then the e-state is discarded. This behavior is similar to the one adopted by the heuristics *hFF* in classical planning.
- `S_PG`: This heuristics is simply an execution of `C_PG` on every e-state.
- e-$\mathbf{A}^*$: This heuristics exploits the fact that e-PG represents an admissible "relaxation" of the planning process and its length $\leq$ than the length of the plan required to go from the starting e-state to the goal state (when reachable). Therefore, we calculate `L_PG` (that uses the length of e-PG as a heuristic value) from each e-state combining it with its depth to generate an $\mathbf{A}^*$-like admissible heuristic.

While other heuristics can be derived from the information generated by e-PG, we leave this investigation as a future work.

To fully take advantage of the introduced heuristics, $\mathcal{H}$-EFP makes use of a Portfolio-like resolution process. For this section, we will generalize the above heuristics along with *Breadth-First Search* (**BFS**) into the set $\mathcal{H}$. Let us note that **BFS** is also included in this to guarantee complete resolution.

Table 1. Performances on the Collaboration and Communication (**CC**) domain of the various solving approaches of $\mathcal{H}$-EFP and EFP.

Collaboration and Communication with $\|\mathcal{AG}\| = 2$, $\|\mathcal{F}\| = 16$, and $\|\mathcal{A}\| = 40$																		
L	Time (seconds)						Expanded Nodes						Plan Length					
	BFS	L_PG	S_PG	C_PG	SUB	e-$\mathbf{A}^*$	**BFS**	L_PG	S_PG	C_PG	SUB	e-$\mathbf{A}^*$	**BFS**	L_PG	S_PG	C_PG	SUB	e-$\mathbf{A}^*$
3	0.141	0.254	0.253	0.329	**0.084**	0.473	6	**3**	**3**	4	4	**3**	**3**	**3**	**3**	4	**3**	**3**
4	0.649	0.322	0.306	TO	**0.174**	0.842	29	**4**	**4**	-	10	7	**4**	**4**	**4**	-	5	**4**
5	6.493	1.578	0.535	TO	**0.244**	8.138	275	18	**5**	-	8	53	**5**	8	**5**	-	**5**	**5**
6	37.236	2.79	0.396	TO	**0.313**	27.672	1611	23	**6**	-	15	151	**6**	8	**6**	-	7	**6**
7	TO	5.425	2.14	TO	**0.443**	TO	-	49	28	-	**22**	-	-	12	**9**	-	**9**	-

Table 2. Direct comparison of EFP and $\mathcal{H}$-EFP on the Grapevine (**Gr**) domain.

Grapevine with $\|\mathcal{AG}\| = 4$						
$\|\mathcal{F}\|$	$\|\mathcal{A}\|$	L	Time (seconds)		Plan Length	
			EFP	$\mathcal{H}$-EFP	EFP	$\mathcal{H}$-EFP
12	40	2	0.047	**0.012**	**2**	**2**
		3	0.352	**0.017**	**3**	**3**
		4	2.253	**0.018**	**4**	**4**
		5	16.384	**0.072**	**5**	6
		6	38.632	**0.08**	**6**	**6**
16	60	2	0.373	**0.06**	**2**	**2**
		3	4.274	**0.093**	**3**	**3**
		4	43.672	**0.094**	**4**	**4**
		5	TO	**0.236**	-	**5**
		6	TO	**1.016**	-	**7**

The planner's remaining components closely resemble those outlined in the works by [4, 7]. For the sake of conciseness, we direct interested readers to these sources for a comprehensive introduction.

4 Experimental Evaluation

Within this section, we conduct a comparative analysis between $\mathcal{H}$-EFP, the primary contribution of this paper, and EFP 2.0 [4] (referred to as EFP for brevity), which is, to the best of our knowledge, the current state-of-the-art comprehensive epistemic planner. While other MEP solvers exist [3,8], we focus our comparison on EFP to illustrate the potential of enhancing MEP planning with heuristics. EFP shares the same underlying code with $\mathcal{H}$-EFP, this highlights that all performance improvements are directly due to heuristics-based reasoning. We leave a complete investigation of $\mathcal{H}$-EFP's scalability for future work. All experiments were executed on a machine equipped with a 3.00GHz Intel Core i9-13900K processor and 128GB of memory.

The evaluation encompasses several domains recognized as standard benchmarks in the MEP setting: *Assembly Line* (**AL**), *Coin in the Box* (**CB**), *Collab-*

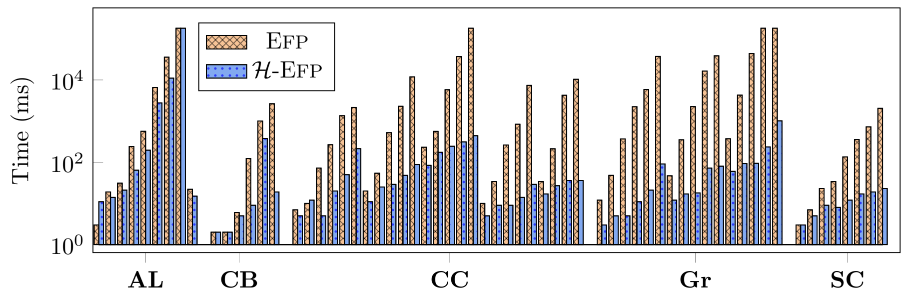

Fig. 1. Plot of the solving times of EFP and $\mathcal{H}$-EFP. The various solving times are grouped for domains.

oration and Communication (**CC**), *Grapevine* (**Gr**), and *Selective Communication* (**SC**). For brevity, detailed descriptions of these domains are omitted but can be found in works such as [4,6] and a comprehensive overview of the experimental results is available here: https://francescofabiano.github.io/resources/.

For the sake of readability, we use the following abbreviations: 'TO' represents Time-Out (solving exceeding 120 seconds), $|\mathcal{AG}|$ signifies the number of agents in the domain, $|\mathcal{F}|$ denotes the number of fluents, $|\mathcal{A}|$ indicates the number of actions, and L highlights the optimal plan length.

Table 1 effectively demonstrates how the incorporation of heuristics increases scalability in the solving process in terms of time and visited nodes, with minimal effects on the plans' quality. The table highlights that as plan length increases, the **BFS** method experiences rapid growth in solving time and explored nodes, eventually rendering it infeasible. Conversely, heuristic-based approaches exhibit significantly lower increases in time requirement and explored nodes relatively to plan length. Among the tested heuristics, `SUB` has the best performance, both in this specific instance and across the other benchmarks. This stems from the minimal overhead in constructing such heuristics. It is noteworthy that $\mathcal{H}$-EFP, leveraging a portfolio-solving approach, consistently selects the most efficient approach in terms of solving times, underscoring its adaptability and performance. This is highlighted in Table 2 that illustrates how $\mathcal{H}$-EFP, thanks to all its components, is able to greatly outperform EFP, especially when the problem complexity increases.

Figure 1 offers a comprehensive overview of the comparison between the two solvers. The visual representation distinctly demonstrates that $\mathcal{H}$-EFP consistently surpasses or, at the very least, equals the state-of-the-art performance exhibited by EFP. This substantiates $\mathcal{H}$-EFP as a superior approach in addressing multi-agent epistemic planning problems.

5 Conclusion

Multi-agent planners, where entities need to plan based on their awareness of other agents' knowledge and beliefs, often face computational inefficiencies. To mitigate this problem, we introduced a Portfolio-style epistemic solver, $\mathcal{H}$-EFP, that showcases promising scalability and substantial performance improvements compared to the state-of-the-art.

While $\mathcal{H}$-EFP is still in its early stages, future work will focus on optimization. Currently, $\mathcal{H}$-EFP's Portfolio-style process keeps heuristic parameters fixed, but expanding it to run all parameter combinations in parallel could reduce tuning time. Additionally, we aim to develop new heuristics from e-PG and enhance e-PG to capture more complex epistemic concepts. Lastly, we plan to explore Reinforcement Learning and Supervised Learning to derive domain-specific heuristics for integration into the portfolio-solving approach.

References

1. Baral, C., Bolander, T., van Ditmarsch, H., McIlraith, S.A.: Epistemic planning. Dagstuhl Reports **7**(6), 1–47 (2017)
2. Baral, C., Gelfond, G., Pontelli, E., Son, T.C.: An action language for multi-agent domains. Artificial Intelligence **302**, 103601 (2022). https://doi.org/https://doi.org/10.1016/j.artint.2021.103601
3. Bolander, T., Dissing, L., Herrmann, N.: DEL-based Epistemic Planning for Human-Robot Collaboration: Theory and Implementation. In: Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning. pp. 120–129 (11 2021). https://doi.org/10.24963/kr.2021/12
4. Fabiano, F., Burigana, A., Dovier, A., Pontelli, E.: EFP 2.0: A multi-agent epistemic solver with multiple e-state representations. In: Proceedings ICAPS, Nancy, France, October 26-30, 2020. pp. 101–109. AAAI Press (2020), https://aaai.org/ojs/index.php/ICAPS/article/view/6650
5. Keyder, E., Geffner, H.: Heuristics for planning with action costs revisited. In: Proceedings of the 2008 Conference on ECAI 2008. pp. 588–592. IOS Press, Amsterdam, The Netherlands, The Netherlands (2008)
6. Kominis, F., Geffner, H.: Beliefs in multiagent planning: From one agent to many. In: Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015, Jerusalem, Israel, June 7-11, 2015. pp. 147–155. AAAI Press (2015), http://www.aaai.org/ocs/index.php/ICAPS/ICAPS15/paper/view/10617
7. Le, T., Fabiano, F., Son, T.C., Pontelli, E.: EFP and PG-EFP: Epistemic forward search planners in multi-agent domains. In: Proceedings ICAPS. pp. 161–170. AAAI Press, Delft, The Netherlands (June 24–29 2018), https://aaai.org/ocs/index.php/ICAPS/ICAPS18/paper/view/17733
8. Muise, C.J., Belle, V., Felli, P., McIlraith, S.A., Miller, T., Pearce, A.R., Sonenberg, L.: Planning over multi-agent epistemic states: A classical planning approach. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA. pp. 3327–3334. AAAI Press (2015), http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9974
9. Russell, S.J., Norvig, P.: Artificial Intelligence - A Modern Approach, Third International Edition. Pearson Education (2010), http://vig.pearsoned.com/store/product/1,1207,store-12521_isbn-0136042597,00.html