Knowledge-Based Strategies for Multi-Agent Teams Playing Against Nature (Extended abstract)

Dilian Gurov
KTH Royal Institute of Technology
Stockholm, Sweden
dilian@kth.se

Valentin Goranko
Stockholm University
Stockholm, Sweden
valentin.goranko@philosophy.su.se

Edvin Lundberg
Rocker AB
Stockholm, Sweden
edvin_lundberg@msn.com

ACM Reference Format:
Dilian Gurov, Valentin Goranko, and Edvin Lundberg. 2021. Knowledge-Based Strategies for Multi-Agent Teams Playing Against Nature (Extended abstract). In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), London, UK, May 3–7, 2021*, IFAAMAS, 3 pages.

In this work[1], we are interested in the strategy synthesis problem for teams (or coalitions) of agents that have to accomplish a given common objective, while acting under imperfect information and under various other natural assumptions. In particular, we are interested in the notion of *knowledge* of agents in that context, and how it affects the strategic abilities of a team.

Motivation. Strategy synthesis for teams of agents is a complex problem. In general, if no bound is put on the size of the memory of the agents, the strategy synthesis problem is undecidable for coalitions of two or more agents in the presence of imperfect information, even for some basic classes of objectives (see, e.g., [4]).

When the information is imperfect, agents typically need to maintain and use a finite abstraction of the history in order to be able to achieve the objective. We refer to this information, suitably structured for use, as "(dynamic) knowledge". We refer to strategies that are directly based on knowledge, i.e., strategies that map knowledge states to actions and update the knowledge state during play, as "knowledge-based strategies".

To achieve certain objectives, agents may have to maintain "higher-order" knowledge (i.e., knowledge about each other's knowledge). Intuitively, the higher the order (or nesting depth) of knowledge, the higher the strategic abilities of the coalition.

For a bounded order of knowledge, the space of potential knowledge states is finite, and the synthesis problem of knowledge-based strategies, unlike that of perfect-recall strategies, is decidable. It is this class of strategies and their synthesis that we investigate here.

Approach. We study the above problem in the context of *multi-player games with imperfect information against Nature* (MAGIIAN, for short). We make the following assumptions on the games:

(1) the game arena is discrete, finite, and known to the agents,
(2) certain game states are indistinguishable for certain agents, thus modelling the "imperfect information",
(3) the agents may or may not see each others' actions,
(4) the agents cannot communicate with each other,
(5) the agents may or may not know each others' strategies.

We believe that the less the agents know or observe, i.e., the higher their uncertainty about the current state-of-affairs, the higher the impact is of having higher-order knowledge. Thus, the case where agents cannot observe each others' actions and do not know each others' strategies is a natural starting point for studying also the less restricted cases, which we discuss in the full paper. We only consider knowledge representations with bounded memory (since the unbounded case gives rise to undecidability results).

Inspired by a subset construction on single-agent games against Nature, namely the Knowledge-Based Subset Construction (or KBSC for short), which reduces games with imperfect information in a strategy-preserving fashion to "expanded" games of perfect information (see, e.g., [1, 5]), we choose a first-order knowledge representation based on sets of game locations. The semantic interpretation of such a set is "the best estimate the agent can make about the current state-of-affairs". In the single-agent case this representation turns out to be sufficient for the class of parity objectives, as shown in [1]. Then, we represent higher-order knowledge by nesting recursively such sets of locations in a suitable fashion.

Also inspired by the KBSC, we investigate the correspondence between knowledge-based strategies and memoryless, observation-based strategies in expanded games resulting from applying a generalised, multi-agent version of the KBSC, which we introduce here and call the MKBSC[2]. The locations of the expanded games are conceptually joint knowledge states of the agents. We propose a generic scheme for extending the KBSC to the multi-agent setting, which is independent of the concrete knowledge representation and the concrete assumptions on what the agents can know and observe. Our generic scheme consists of four stages:

(1) **Projection**: for each agent $i \in \mathsf{Agt}$, compute the individual views of the input game G, based on what the agent knows, does and sees. This stage results in *n* single-agent games with imperfect information.
(2) **Expansion**: expand each of the individual views with the KBSC. The results are *n* single-agent games with perfect information.
(3) **Composition**: combine the individual expansions by using a product construction, resulting in a single multi-agent game with perfect information.

[1]This extended abstract is based on a full paper available from https://arxiv.org/abs/2012.14851, currently under journal submission.

[2]An earlier version of this construction, of which the MKBSC construction presented here is an improvement, was developed in the Master's thesis [2].

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, London, UK.

(4) **Partition**: define each agent's observations as induced by the composition product, reflecting their local knowledge. The final result is a multi-agent game with imperfect information.

We define a concrete instantiation of that scheme in the full paper. The MKBSC can be iterated, essentially computing higher-order knowledge (i.e., incrementing the knowledge depth with each iteration). An implementation of the MKBSC as a tool[3] is described in [3].

We refer to the two emerging views on strategies, based on knowledge and based on expanded games, as the "intensional" and the "extensional" view, respectively. The correspondence between the two views is useful in several ways. First, one can reduce the synthesis problem of knowledge-based strategies to the synthesis problem of memoryless, observation-based strategies. Furthermore, the individual observation-based memoryless strategies in the expanded games are simultaneously memoryless strategies in the single-player games of perfect information that are produced at stage (2) of the MKBSC. This can serve as the basis for the design of efficient knowledge-based strategy synthesis algorithms, since strategy synthesis for the latter class is well-studied and has efficient solutions. Second, while strategy synthesis is more conveniently performed on the expanded games, once synthesised, the strategies can be presented to the agents as knowledge-based strategies, without the need for storing the expanded games, but by recomputing the knowledge in the course of the play (i.e., on-the-fly). And third, there is a "knowledge saturation" phenomenon that manifests itself much more explicitly in the extensional view: for some games, the iterated MKBSC "stabilises", producing isomorphic games from some iteration on. In the intensional view, stabilisation corresponds to the existence of a finite knowledge representation that contains the knowledge corresponding to the limit of the MKBSC construction. For games on which the iterated construction eventually stabilises, this opens up the possibility to achieve and utilise common knowledge.

Contributions. This work includes the following results and contributions. First, we propose a formal notion of higher-order knowledge with a representation and semantic interpretation, and a notion of knowledge update. Based on this, we provide a formal notion of knowledge-based strategies. Next, we develop a generalisation of the KBSC to the case of multiple agents, as a scheme that can be reused for other similar expansions. The construction in effect computes knowledge, and its iteration computes higher-order knowledge. For this construction, we show a strategy preservation result for perfect recall strategies with respect to reachability and safety objectives. The proof of this result is constructive, and also reveals how to preserve finite-memory strategies. We then establish a formal relationship between memoryless observation-based strategies in the expanded games, knowledge-based strategies in the original games, and the corresponding finite-memory strategies in the original games, and the equivalence of the latter two. With this we also exhibit formally the duality between the intensional and the extensional views. From this correspondence, we obtain a reduction of the synthesis problem of knowledge-based strategies to that of memoryless observation-based strategies in expanded games. Then, we sketch a heuristic for strategy synthesis, based on the observation made above. Further, we give a formal meaning to the statement that the higher the order of knowledge, the higher the strategic abilities of the team, and argue that this indeed is the case here. (However, this increase is not strict, as will become clear.) Finally, we make the discovery that for some games the iterated MKBSC stabilises, in the sense that from some iteration on it results in isomorphic games. One implication of this is that for stabilising games, the problem of existence of a winning knowledge-based strategy, without a predefined bound on the knowledge nesting depth, is decidable. Another implication is that for stabilising games, one can use the iterated MKBSC to compute common knowledge.

Conclusions and future work. From our results one can draw the following conclusions. First, higher-order knowledge of agents can be based on nesting a first-order notion of knowledge interpreted as "most precise estimate of the current state-of-affairs". The higher the order (i.e., the nesting depth) of knowledge, the higher the strategic abilities of the team. Also, the higher the uncertainty of the players, i.e., the less they observe and know, the higher the benefit from nested knowledge. Next, for the class of knowledge-based strategies considered here (i.e., for the proposed notion of knowledge) and the classes of reachability and safety objectives, for a given bound on the nesting depth of knowledge and under the additional condition that the team of agents has a perfect distributed knowledge (PDK condition), we have an *algorithm* for strategy synthesis; without such a bound it is only a semi-algorithm. Then, there is a *duality* between the extensional and the intensional views. The former is more suitable for strategy synthesis, while the latter can be more convenient in the play, and can also be used to explain the synthesised strategy. And finally, on some games the iterated MKBSC *stabilises*. If there is no winning memoryless observation-based strategy in the stable expansion, then, under the PDK condition, there is no winning knowledge-based strategy of any order in the original game. However, there might still be a winning finite-memory observation-based strategy in the original game.

In future work we plan to characterise the class of objectives that can be achieved with knowledge-based strategies of the type defined here. We also plan to study the strategy synthesis problem after relaxing some of the assumptions made here, such as when agents know each others strategies, or when agents do have some (limited) communication. Further, we plan to study in depth the stabilisation phenomenon of the MKBSC and, in particular, characterise the structural conditions for stabilisation, and investigate the relationship of stable games to common knowledge. Furthermore, we will explore other knowledge representations, comparing the respective classes of objectives that they are sufficient for, and define the corresponding expansions following the general scheme. We will also design strategy synthesis algorithms and heuristics, and investigate their complexity. Further, we will explore temporal (epistemic) logic as a means for defining objectives, and epistemic logic as a means for representing the individual knowledge-based strategies. Finally, we plan to evaluate the practical utility of the strategy synthesis method proposed here, and investigate potential application areas.

[3] Available from `github.com/helmernylen/mkbsc`.

REFERENCES

[1] Chatterjee, K., Doyen, L., Henzinger, T.A., Raskin, J.: Algorithms for omega-regular games with imperfect information. Logical Methods in Computer Science 3(3), 1–23 (2007)

[2] Lundberg, E.: Collaboration in Multi-Agent Games. Tech. rep., KTH Royal Institute of Technology, School of Computer Science and Communication (2017), https://kth.diva-portal.org/smash/get/diva2:1115157/FULLTEXT01.pdf

[3] Nylén, H., Jacobsson, A.: Investigation of a Knowledge-based Subset Construction for Multi-player Games of Imperfect Information. Tech. rep., KTH Royal Institute of Technology, School of Computer Science and Communication (2018), https://kth.diva-portal.org/smash/get/diva2:1221520/FULLTEXT01.pdf

[4] Pnueli, A., Rosner, R.: Distributed reactive systems are hard to synthesize. In: Symposium on Foundations of Computer Science (FOCS'90). pp. 746–757. IEEE Computer Society (1990)

[5] Reif, J.H.: The complexity of two-player games of incomplete information. Computer and System Sciences 29(2), 274–301 (1984)