Relentful Strategic Reasoning in Alternating-Time Temporal Logic ⋆

Fabio Mogavero[1], Aniello Murano[1], and Moshe Y. Vardi[2]

[1] Universitá degli Studi di Napoli "Federico II", I-80126 Napoli, Italy.
[2] Rice University, Department of Computer Science, Houston, TX 77251-1892, U.S.A.
{mogavero, murano}@na.infn.it vardi@cs.rice.edu

1 Extended Abstract

Multi-agent systems recently emerged as a new paradigm for better understanding distributed systems [Woo01, FHMV95]. In multi-agent systems, different processes can have different goals and the interactions between them may be adversarial or cooperative. Interactions between processes in multi-agent systems can thus be seen as games in the classical framework of game theory, with adversarial coalitions [OR94]. Classical branching-time temporal logics, such as CTL* [EH86], turn out to be of limited power when applied to multi-agent systems. For example, consider the property `Prop`: "processes 1 and 2 cooperate to ensure that a system (having more than two processes) never enters a fail state". It is well known that CTL* cannot express `Prop` [AHK02]. Rather, CTL* can only say whether the set of all agents can or cannot prevent the system from entering a fail state.

To allow the temporal-logic framework to work within the setting of multi-agent systems, Alur, Henzinger, and Kupferman introduced *Alternating-Time Temporal Logic* (ATL*, for short) [AHK02]. This is a generalization of CTL* obtained by replacing the path quantifiers, "E" (*there exists*) and "A" (*for all*), with "*cooperation modalities*" of the form $\langle\!\langle A \rangle\!\rangle$ and $[\![A]\!]$, where A is a set of *agents*, which can be used to represent the power that a coalition of agents has to achieve certain results. In particular, these modalities express selective quantifications over those paths that can be effected as outcomes of infinite games between the coalition and its complement. ATL* formulas are interpreted over *game structures*, which model a set of interacting processes. Given a game structure $\mathcal{G}$ and a set A of agents, the ATL* formula $\langle\!\langle A \rangle\!\rangle \psi$ is satisfied at a state s iff there is a *strategy* for the agents in A such that, no matter the strategy that is executed by agents not in A, the resulting outcome of the interaction satisfies ψ at s. Coming back to the previous example, one can see that the property `Prop` can be expressed by the ATL* formula $\langle\!\langle \{1,2\} \rangle\!\rangle \mathsf{G}\, \neg \mathit{fail}$, where G is the classical temporal modality "*globally*".

Traditionally, temporal logics are *memoryless*: once a path in the underlying structure (usually a computation tree) is quantified at a given state, the computation that led to that state is forgotten [KV06]. In the case of ATL*, we have even more: the logic is also "relentless", in the sense that the agents are not able to formulate their strategies depending on the history of the computation; when $\langle\!\langle A \rangle\!\rangle \psi$ is asserted in a state s, its truth is independent of the path that led to s. Inspired by a work on *strong cyclic planning* [DTV00], Pistore and Vardi proposed a logic that can express the spectrum between the strong goal $\mathsf{A}\psi$ and the weak goal $\mathsf{E}\psi$ in planning [PV07]. A novel aspect of the Pistore-Vardi logic is that it is "*memoryful*", in the sense that the satisfiability of a formula at a state s depends on the future as well as on the past, i.e., the trace starting from the initial state and leading to s. Nevertheless, this logic does not have a standard temporal logical syntax (for example, it is not closed under conjunction and disjunction). Also, it is less expressive than CTL*. This has lead Kupferman and Vardi [KV06] to introduce a memoryful variant of CTL* (mCTL*, for short), which unifies in a common framework both CTL* and the Pistore-Vardi logic. Syntactically, mCTL* is obtained from CTL* by simply adding a special proposition *present*, which is needed to emulate the ability of CTL* to talk about the "present" time. Semantically, mCTL* is obtained from CTL* by reinterpreting the path quantifiers of the logic to be memoryful.

Recently, ATL* has become very popular in the context of multi-agent system planning [vdHW02, Jam04]. In such a framework, a memoryful enhancement of ATL* enables "relentful" planning, that is, agents can relent and change their goals, depending on their history[3]. That is, when a specific goal at a certain state is checked, agents may learn from the past to change their goals. Note that this does not mean that agents change their strategy, but that they can choose a strategy that allows them to change their goals. For example, consider the ATL* formula $\langle\!\langle \emptyset \rangle\!\rangle \mathsf{G} \langle\!\langle A \rangle\!\rangle \psi$. In the memoryful framework, this formula is satisfied by a game structure $\mathcal{G}$ (at its starting node) iff for each possible trace (history) ρ the agents in A can ensure that the evolution of $\mathcal{G}$ that extends ρ satisfies ψ from the start state.

In this paper, we introduce and study the logic mATL*, a memoryful extension of ATL*. Thus, mATL* can be thought of as a fusion of mCTL* and ATL* in a common framework. Similarly to mCTL*, the syntax of mATL* is obtained from ATL* by simply adding a special proposition *present*. Semantically, mATL* is obtained from ATL* by reinterpreting the path quantifiers of the logic to be memoryful. More specifically, for a game

⋆ An extended version of this paper is going to appear in the proceedings of LPAR-16 2010 [MMV10].

[3] In Middle English to relent means to melt. In modern English it is used only in the combination of "relentless".

structure $\mathcal{G}$, the mATL* formula $\langle\langle A \rangle\rangle \psi$ holds at a state s of $\mathcal{G}$ if there is a strategy for agents in A such that, no matter which is the strategy of the agents not in A, the outcome of the game, obtained by *extending* the execution trace of the system ending in s, satisfies ψ. As an example, consider the situation in which the agents in a set A have the goal to satisfy p until q holds, unless they see r in which case they change their goal to satisfy u until v holds from the *start* of the computation. This is difficult to handle in ATL*, since the specification depends on the past. On the other hand, it can be easily handled in mATL*, with the formula $\langle\langle A \rangle\rangle ((p \, \mathsf{U} \, (q \vee r)) \wedge \mathsf{G} \, f)$, where f is $r \rightarrow \langle\langle A \rangle\rangle (u \, \mathsf{U} \, v)$.

In the paper, we also consider an extension of mATL* with *past operators* (mpATL*, for short). As for classical temporal logics, past operators allow reasoning about the past in a computation [LPZ85]. In mpATL*, we can further require that coalitions of agents had a memoryful goal in the past. In more details, we can write a formula whose satisfaction, at a state s, depends on the trace starting from the initial state and leading to a state s' occurring before s. Coming back to the previous example, by using P as the dual of F, we can change the alternative goal f of agents in A to be $r \rightarrow \mathsf{P} \, (h \wedge \langle\langle A \rangle\rangle (u \, \mathsf{U} \, v))$, which requires that once r occurs at a state s, at a previous state s' of s in which h holds, the subformula u until v from the start of the computation must be true.

An important contribution of this work is to show for the first time a clear and complete picture of the relationships among ATL* and its extensions with memoryful quantification and past modalities, which goes beyond the expressiveness results obtained in [KV06] for mCTL*. Since memoryfulness refers to behavior from the start of the computation, which occurred in the past, memoryfulness is intimately connected to the past. Indeed, we prove this formally. We study the expressive power and the succinctness of mATL* w.r.t. ATL*, as well as the memoryless fragment of mpATL* (i.e., the extension of ATL* with past modalities), which we call pATL*. We show that the three logics have the same expressive power, but both mATL* and pATL* are at least exponentially more succinct than ATL*. As for m^{-}ATL* (where the minus stands for the variant of the logic without the "present" proposition but the path interpretation is still memoryful), we prove that it is strictly less expressive than ATL*. Conversely, we prove that pATL* is equivalent to p^{-}ATL*, but exponentially more succinct.

From an algorithmic point of view, we examine two decision problems for mpATL*, *model checking* and *satisfiability*. We show that model checking is not easier than satisfiability and in particular that both are 2EXPTIME-COMPLETE, as for ATL*. We recall that this is not the case for mCTL*, where the model checking is EXPSPACE-COMPLETE, while satisfiability is 2EXPTIME-COMPLETE. For upper bounds, we follow an *automata-theoretic approach* [KVW00]. To this aim, we introduce a new automaton model, the *agent-action tree automata with satellites* (AGCTAS, for short), which extends both *automata over concurrent game structures* in [SF06] and *alternating automata with satellites* in [KV06], in a common setting. Then we reduce the decision problems for mpATL* to the emptiness problem of such automata. For technical convenience, AGCTAS states are partitioned into states regarding the satellite and those regarding the rest of the automaton, which we call the *main automaton*. The complexity results then come from the fact that mpATL* formulas can be translated into an AGCTAS with an exponential number of states for the main automaton and doubly exponential number of states for the satellite, and from the fact that the emptiness problem for AGCTAS is solvable in EXPTIME w.r.t. both the size of the main automaton and the logarithm of the size of the satellite.

References

[AHK02] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *JACM*, 49(5):672–713, 2002.

[DTV00] M. Daniele, P. Traverso, and M.Y. Vardi. Strong Cyclic Planning Revisited. In *ECP'99*, pages 35–48, 2000.

[EH86] E.A. Emerson and J.Y. Halpern. "Sometimes" and "Not Never" Revisited: On Branching Versus Linear Time. *JACM*, 33(1):151–178, 1986.

[FHMV95] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge.* MIT Press, Cambridge, Mass., 1995.

[Jam04] W. Jamroga. Strategic Planning Through Model Checking of ATL Formulae. In *ICAISC'04*, LNCS 3070, pages 879–884. Springer-Verlag, 2004.

[KV06] O. Kupferman and M.Y. Vardi. Memoryful Branching-Time Logic. In *LICS'06*, pages 265–274. IEEE Computer Society, 2006.

[KVW00] O. Kupferman, M.Y. Vardi, and P. Wolper. An Automata Theoretic Approach to Branching-Time Model Checking. *JACM*, 47(2):312–360, 2000.

[LPZ85] O. Lichtenstein, A. Pnueli, and L.D. Zuck. The Glory of the Past. In *LP'85*, pages 196–218, 1985.

[MMV10] F. Mogavero, A. Murano, and M.Y. Vardi. Relentful Strategic Reasoning in Alternating-Time Temporal Logic. In *LPAR'10*, 2010.

[OR94] M.J. Osborne and A. Rubinstein. *A course in game theory.* MIT Press, 1994.

[PV07] M. Pistore and M.Y. Vardi. The Planning Spectrum - One, Two, Three, Infinity. *JAIR'07*, 30:101–132, 2007.

[SF06] S. Schewe and B. Finkbeiner. Satisfiability and Finite Model Property for the Alternating-Time μ-Calculus. In *CSL'06*, pages 591–605, 2006.

[vdHW02] W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In *AAMAS'02*, pages 1167–1174, 2002.

[Woo01] M.J. Woolridge. *Introduction to Multiagent Systems.* John Wiley & Sons, 2001.