Towards a Proof Theory for Multi-Agent Logics with Irrevocable Strategies

Thomas Ågotnes[1], **Valentin Goranko**[2], and **Wojciech Jamroga**[3]

[1] Bergen University College
[2] University of the Witwatersrand
[3] Clausthal University of Technology
emails: `tag@hib.no, goranko@maths.wits.ac.za,`
`wjamroga@in.tu-clausthal.de`

Abstract. Logics such as Alternating-time Temporal Logic (ATL) enable us to reason about the strategic abilities of agents. However, strategies in the semantics of ATL are, in a certain precise sense, *revocable*. On the other hand, *irrevocable strategies* are often naturally assumed, e.g., in game theory. In a recent work [1] we have introduced versions of ATL employing irrevocable strategies. This abstract provides a concise account of the logic of irrevocable strategies and describes initial work towards a proof theory for it.

1 Introduction

Logics for game-like scenarios, such as the Alternating-time Temporal Logic ATL [2], have received much interest recently [2, 9], as logical foundations for multi-agent systems. The main semantic assumption behind ATL is that a system at a given time is in one of several possible states, and that the next state of the system is determined by the current state and the actions chosen by each of κ *agents* present in the system. ATL involves strategic quantifiers (called *cooperation modalities*), such as $\langle\langle C\rangle\rangle\mathcal{X}$ and $\langle\langle C\rangle\rangle\mathcal{G}$ where C is a set of agents. Formula $\langle\langle C\rangle\rangle\mathcal{X}\phi$ is intended to mean that coalition C can achieve ϕ in the ne$\mathcal{X}$t state of the system, or, in more detail, that the agents in C can choose their strategies so that, if they use these strategies then ϕ will be true in the next state – no matter what the agents outside C do. Similarly, $\langle\langle C\rangle\rangle\mathcal{G}\phi$ means that C can force ϕ to be true in all future states ($\mathcal{G}$lobally).

Below, we first give a quick review of ATL, and illustrate the fact that strategies in ATL can be seen as *revocable*. We then define a new semantics for ATL with *irrevocable* strategies, and then discuss preliminary results towards a proof theory.

2 Alternating-time Temporal Logic

Here we recall briefly the semantic and syntactic basics of ATL. For full details, see any of [2–4].

A *concurrent game structure* (CGS) is a tuple $M = (\Sigma, \Pi, Q, \pi, Act, d, \delta)$ where $\Sigma = \{1, \dots, \kappa\}$, for $\kappa > 0$, is the set of *agents (players)*; Π is a set of *atomic propositions*; Q is a set of *states*; $\pi : Q \to \wp(\Pi)$ is the *labeling* function; Act is a (usually finite) set of *actions*; d is a mapping such that for each player $i \in \Sigma$ and state $q \in Q$, $d_i(q) \subseteq Act$ is the non-empty set of actions available to player i in q; and δ is the *transition function*, mapping every pair $(q \in Q, \alpha \in D(q))$, where $D(q) = d_1(q) \times \cdots \times d_\kappa(q)$ is the set of *joint actions* at q, to an *outcome state* $\delta(q, \alpha) \in Q$.

For $\alpha \in D(q)$, let α_i denote the i-th component of α. Likewise, $\alpha(C)$ denotes the projection of α onto $C \subseteq \Sigma$, and $D(q, C)$ denotes the projection of $D(q)$ onto C. Whenever necessary, we may write $D(M, q)$ and $D(M, q, C)$ in order to indicate the CGS explicitly. A *pointed CGS* is a pair (M, q) where M is a CGS and q is a state in M.

Given a CGS M, $C \subseteq \Sigma$, a state q in M, and a tuple of actions $\alpha(C) \in D(q, C)$, one for each agent in C, we denote by $out(M, q, \alpha(C))$ the set of outcome states of all joint actions extending $\alpha(C)$. Formally,

$$out(M, q, \alpha(C)) = \{\delta(q, \alpha') \mid \alpha' \in D(q),\ \alpha'(C) = \alpha(C)\}.$$

In particular, $out(M, q, \alpha) = \{\delta(q, \alpha)\}$. Also, we write $out(M, q)$ for the set $out(M, q, \alpha(\emptyset))$ of all possible outcome states from q; when M is fixed, we write $out(q)$.

A *computation* λ is an infinite sequence of states; $\lambda = q_0 q_1 \cdots$, where for each $j \geq 0$ there is a joint action $\alpha \in D(q_j)$ such that $\delta(q_j, \alpha) = q_{j+1}$. By $\lambda[j]$ we will denote the element (q_j) in λ with index j ; respectively $\lambda|j$ will denote the initial segment of λ ending with $\lambda[j]$. Such an initial segment will be called a *finite computation*. The last state of a finite computation σ will be denoted by $l(\sigma)$. Whenever suitable we will regard a state q as a one step computation, and then $l(q) = q$. A simple (or memoryless) *strategy* for a player i is a function $f_i : Q \rightarrow Act$ with $f_i(q) \in d_i(q)$ for each $q \in Q$. That is, the strategy maps each state to an action for player i. A *memory-based strategy* for a player i is a function $f_i : Q^+ \rightarrow Act$ with $f_i(\sigma) \in d_i(l(\sigma))$, i.e., it maps possible histories of the play to i's choices. Clearly, memoryless strategies can be seen as special cases of memory-based strategies, where $f_i(q_1 \dots q_n)$ only depends on the last state q_n. A *joint strategy* for $C \subseteq \Sigma$ is a tuple of strategies, one per $i \in C$; by $Str(M, C)$ we denote the set of joint memoryless strategies for C in M. We then denote by $f_C(q)$ the tuple of respective actions $f_i(q)$ for $i \in C$, and adopt the notation $out(M, q, f_C)$ for $out(M, q, f_C(q))$. Given a state q and a joint memoryless strategy f_C for C, $comp(M, q, f_C)$ denotes the set of possible computations starting in state q where the agents in C use the strategies f_C. Formally, $\lambda \in comp(M, q, f_C)$ iff $\lambda[0] = q$ and for all $j \geq 0$, $\lambda[j+1] \in out(M, \lambda[j], f_C)$.

A language for ATL is determined by the set of atomic propositions Π and the set of agents Σ, and will be denoted by ATL(Π, Σ). The formulae of ATL(Π, Σ) are defined recursively as follows, where $p \in \Pi$ and $C \subseteq \Sigma$:

$$\phi ::= \top \mid p \mid \neg\phi \mid \phi \wedge \phi \mid \langle\langle C \rangle\rangle \mathcal{X}\phi \mid \langle\langle C \rangle\rangle \mathcal{G}\phi \mid \langle\langle C \rangle\rangle \phi\mathcal{U}\phi$$

In addition to the standard propositional connectives, we use $\langle\langle C \rangle\rangle \mathcal{F}\phi$ for $\langle\langle C \rangle\rangle (\top\mathcal{U}\phi)$, and write $\mathbf{A}\mathcal{X}\phi$, $\mathbf{A}\mathcal{G}\phi$, $\mathbf{A}(\phi_1\mathcal{U}\phi_2)$ respectively for $\langle\langle \emptyset \rangle\rangle \mathcal{X}\phi$, $\langle\langle \emptyset \rangle\rangle \mathcal{G}\phi$, $\langle\langle \emptyset \rangle\rangle (\phi_1\mathcal{U}\phi_2)$; $\mathbf{E}\mathcal{X}\phi$, $\mathbf{E}\mathcal{G}\phi$, $\mathbf{E}(\phi_1\mathcal{U}\phi_2)$ respectively for $\langle\langle \Sigma \rangle\rangle \mathcal{X}\phi$, $\langle\langle \Sigma \rangle\rangle \mathcal{G}\phi$ and $\langle\langle \Sigma \rangle\rangle (\phi_1\mathcal{U}\phi_2)$; and $[\![C]\!]\mathcal{X}\phi$ and $[\![C]\!]\mathcal{G}\phi$ respectively for the duals $\neg\langle\langle C \rangle\rangle \mathcal{X}\neg\phi$ and $\neg\langle\langle C \rangle\rangle \mathcal{F}\neg\phi$. Truth of a formula ψ in a state q of a CGS M is defined as usual for the Boolean connectives and via the following clauses for the strategic temporal operators:[4]

$$M, q \models_{\text{ATL}} \langle\langle C \rangle\rangle \mathcal{X}\phi \;\Leftrightarrow\; \exists f_C \in \mathit{Str}(M, C) \forall \lambda \in \mathit{comp}(M, q, f_C)(M, \lambda[1] \models_{\text{ATL}} \phi)$$

$$M, q \models_{\text{ATL}} \langle\langle C \rangle\rangle \mathcal{G}\phi \;\Leftrightarrow\; \exists f_C \in \mathit{Str}(M, C) \forall \lambda \in \mathit{comp}(M, q, f_C) \forall j \geq 0 (M, \lambda[j] \models_{\text{ATL}} \phi)$$

$$M, q \models_{\text{ATL}} \langle\langle C \rangle\rangle \phi_1 \mathcal{U} \phi_2 \;\Leftrightarrow$$

$$\exists f_C \in \mathit{Str}(M, C) \forall \lambda \in \mathit{comp}(M, q, f_C) \exists j \geq 0 ((M, \lambda[j] \models_{\text{ATL}} \phi_2) \;\&$$

$$\forall 0 \leq k < j (M, \lambda[k] \models_{\text{ATL}} \phi_1)).$$

Now, $M \models_{\text{ATL}} \phi$ if $M, q \models_{\text{ATL}} \phi$ for every $q \in Q$; ϕ is ATL-*valid* if $M \models_{\text{ATL}} \phi$ for every CGS M.

In the semantics of ATL, being compositional, strategies can be *revoked* by agents. For instance, the truth of $\langle\langle a \rangle\rangle \mathcal{G}\phi$ requires existence of a strategy f_a for the agent a that guarantees that ϕ remains true along every computation compatible with f_a, but when evaluating the truth of ϕ at a state of such computation in order to verify the truth of $\langle\langle a \rangle\rangle \mathcal{G}\phi$, we *no longer assume that the agent actually adheres to that strategy*. Philosophically, this means that in the semantics of ATL the agents are not "committed" to the strategies claimed to exists by a given cooperation modality, but they are free to change their minds as soon as the next choice point (i.e., next cooperation modality) is encountered. The following example demonstrates the effect of revocable strategies.

Example 1. We are given a system with a shared resource, and are interested in reasoning about whether agent a has access to the resource. Let p denote the fact that agent a controls the resource. The ATL formula $\langle\langle a \rangle\rangle \mathcal{X} p$ expresses the fact that a is able to obtain control of the resource in the next moment, if she chooses to. Now imagine that agent a does not need to access the resource all the time, but she would like to be able to control the resource any time she needs it. This can be expressed in ATL by formula $\langle\langle a \rangle\rangle \mathcal{G} \langle\langle a \rangle\rangle \mathcal{X} p$, saying that a has a strategy which guarantees that, in any future state of the system, a can always force the next state to be one where a controls the resource.

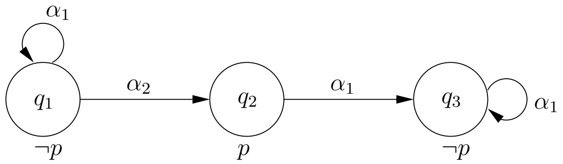

Fig. 1. System M_0 with a single agent a. The transitions between states are labeled by the actions chosen by a.

Consider system M_0 from Figure 1. We have that $M_0, q_1 \models \langle\langle a \rangle\rangle \mathcal{X} p$: a can choose action α_2, which guarantees that p is true next. But we also have that $M_0, q_1 \models \langle\langle a \rangle\rangle \mathcal{G} \langle\langle a \rangle\rangle \mathcal{X} p$:

---

[4] Note that memoryless and memory-based strategies yield equivalent semantics for the standard ATL [7].

a's strategy in this case is to always choose α_1, which guarantees that the system will stay in q_1 forever and, as we have seen, $M_0, q_1 \models \langle\langle a \rangle\rangle \mathcal{X} p$. However, this system does not have the exact property we had in mind because, by following that strategy, the agent a dooms herself to *never access the resource* – in which case it is maybe counter-intuitive that $\langle\langle a \rangle\rangle \mathcal{X} p$ should be true. In other words, a can ensure that she is forever *able* to access the resource – but only by never *actually* accessing it. Indeed, while a can force the *possibility* of achieving p to be true forever, the actual achievement of p destroys that possibility.

We argue that, while there is nothing "wrong" with the standard ATL semantics, there are many situations in which the "committed" interpretation of strategies we introduce next is more appropriate. For instance, a strategy in game theory is usually understood as a *complete* plan that prescribes the player's behaviour in all conceivable situations, and for all future moments. Under this interpretation, $\langle\langle a \rangle\rangle \mathcal{G} \langle\langle a \rangle\rangle \mathcal{X} p$ *cannot* hold in M_0, q_1, because a can only play according to the first or the second strategy, but not to both of them at the same time. Irrevocable strategies are often naturally assumed, not only in game theory, but also in controller synthesis (the controller is an irrevocable strategy), and in planning in AI, where the actual achievement of certain subgoals may affect the possibility of achieving other subgoals.

3 Model Updates and Semantics with Irrevocable Strategies

We start by considering the memoryless case of irrevocable strategies. First, let us spell out again the intuition behind the semantics with irrevocable strategies: a formula such as $\langle\langle C \rangle\rangle \mathcal{X} \phi$ means that there is a joint strategy for C such that, if the actions of the agents from C are thereafter fixed by the choices prescribed by the strategy, ϕ will necessarily be true in the next state; likewise for the other temporal operators. This motivates the notion of a *model update*, similar to the one in [10, 8]: the update of a model by a joint memoryless strategy for C is a model obtained by fixing the choices of every agent in C in each state as prescribed by the strategy for that agent in that state.

Definition 1 (Model Update). *Let M be a CGS, C a coalition, and $f_C \in Str(M, C)$ a memoryless strategy. The update of M by f_C, denoted $M \dagger f_C$, is the same as M, except that the choices of each agent $i \in C$ are fixed by the strategy f_i: $d_i(q) = \{f_i(q)\}$ for each state q.*

The language of the logic IATL – ATL *with irrevocable strategies* – is the same as the language of ATL. Let q be a state in a CGS M. The semantics of the strategic operators in IATL is defined as follows:

$$M, q \models_{\text{IATL}} \langle\langle C \rangle\rangle \mathcal{X} \phi \text{ iff}$$

$$\exists f_C \in Str(M, C) \forall \lambda \in comp(M \dagger f_C, q, f_C)(M \dagger f_C, \lambda[1] \models_{\text{IATL}} \phi)$$

$$M, q \models_{\text{IATL}} \langle\langle C \rangle\rangle \mathcal{G} \phi \text{ iff}$$

$$\exists f_C \in Str(M, C) \forall \lambda \in comp(M \dagger f_C, q, f_C) \forall j \geq 0 (M \dagger f_C, \lambda[j] \models_{\text{IATL}} \phi)$$

$$M, q \models_{\text{IATL}} \langle\langle C\rangle\rangle(\phi_1 \mathcal{U} \phi_2) \text{ iff}$$
$$\exists f_C \in \mathit{Str}(M, C) \forall \lambda \in \mathit{comp}(M \dagger f_C, q, f_C) \exists j \geq 0 (M \dagger f_C, \lambda[j] \models_{\text{IATL}} \phi_2 \;\&$$
$$\forall 0 \leq k < j (M \dagger f_C, \lambda[k] \models_{\text{IATL}} \phi_1))$$

Note, again, that the logic IATL is defined with memoryless strategies. We also define a version of the irrevocable strategies semantics for memory-based strategies, called MIATL (*memory-based* IATL). The language of the logic MIATL is the same as the language of IATL. Unlike in IATL, in MIATL we cannot update the model directly, but must first unfold the model into an (equivalent) tree-like structure, and then eliminate the branches which represent computations which do not conform to the strategy we update with. The tree-unfolding of a CGS M from a state q is denoted $T(M, q)$. Note that a memory-based strategy in M is equivalent to a memoryless strategy in $T(M, q)$. Thus, the MIATL semantics can be defined as follows:

$$M, q \models_{\text{MIATL}} \phi \Leftrightarrow T(M, q), q \models_{\text{IATL}} \phi$$

Intuitively, the MIATL meaning of the cooperation modalities involves pruning the model by a memory-based strategy.

The difference between the above and the original ATL interpretations is that in the former the subformula ϕ of a formula $\langle\langle C\rangle\rangle T\phi$ is evaluated in an updated model, where the actions of group C are fixed, while in the latter the subformula ϕ of $\langle\langle C\rangle\rangle T\phi$ is still evaluated in the original model. Consider model M_0 from Example 1 again. We have that $M_0, q_1 \models_{\text{ATL}} \langle\langle a\rangle\rangle \mathcal{G} \langle\langle a\rangle\rangle \mathcal{X} p$, but $M_0, q_1 \not\models_{\text{IATL}} \langle\langle a\rangle\rangle \mathcal{G} \langle\langle a\rangle\rangle \mathcal{X} p$, and $M_0, q_1 \not\models_{\text{MIATL}} \langle\langle a\rangle\rangle \mathcal{G} \langle\langle a\rangle\rangle \mathcal{X} p$. To that IATL and MIATL are different, it is enough to observe that $M_0, q_1 \models_{\text{MIATL}} \langle\langle a\rangle\rangle \mathcal{X} \langle\langle a\rangle\rangle \mathcal{X} p$, but $M_0, q_1 \not\models_{\text{IATL}} \langle\langle a\rangle\rangle \mathcal{X} \langle\langle a\rangle\rangle \mathcal{X} p$.

4 Towards a Proof Theory

Here we describe first steps towards a proof theory for IATL and MIATL. We sometimes use the terms ATL, IATL, and MIATL more specifically to denote the sets of validities of the respective logics.

4.1 Comparing ATL, IATL, and MIATL

Let us first point out that the semantics of IATL and MIATL coincide for a special class of models.

Definition 2. *A* CGS *is* tree-like *if there is a state (root) from which every state can be reached by a unique finite computation. A typical example of tree-like CGS is the* tree-unfolding $T(M, q)$ *of a given CGS* M *from a given state* q *in it.*

Since every state in a tree-like CGS has a unique history (the path from the root), memoryless and memory-based strategies coincide in tree-like CGSs, and therefore IATL and MIATL are equivalent in them. More precisely:

Proposition 1. *For every pointed tree-like CGS* (M, q) *and an* ATL*-formula* ϕ*:* $M, q \models_{\text{IATL}} \phi$ *iff* $M, q \models_{\text{MIATL}} \phi$*.*

Proposition 2. *If a formula is* MIATL*-satisfiable, then it is* MIATL*-satisfiable in a tree-like CGS.*

Theorem 1.

1. ATL $\not\subseteq$ IATL*, and* ATL $\not\subseteq$ MIATL*;*
2. IATL $\not\subseteq$ ATL*, and* MIATL $\not\subseteq$ ATL*;*
3. IATL $\subseteq$ MIATL*;*
4. MIATL $\not\subseteq$ IATL.

Proof. (Sketch). (1) follows from the fact that IATL and MIATL are not closed under uniform substitution (consider the substitution of $\langle\!\langle \Sigma \rangle\!\rangle \mathcal{X} p \wedge \langle\!\langle \Sigma \rangle\!\rangle \mathcal{X} \neg p$ for p in the ATL axiom $\neg \langle\!\langle \emptyset \rangle\!\rangle \mathcal{X} p \to \langle\!\langle \Sigma \rangle\!\rangle \mathcal{X} p$). As for (2), note that $\langle\!\langle C \rangle\!\rangle \mathcal{X} \langle\!\langle C \rangle\!\rangle \mathcal{X} \phi \leftrightarrow \langle\!\langle C \rangle\!\rangle \mathcal{X} \langle\!\langle \emptyset \rangle\!\rangle \mathcal{X} \phi$, for $C \neq \emptyset$, is a validity of both IATL and MIATL, but not of ATL. (3) follows from Propositions 1 and 2. To show (4) we construct a more complicated formula, valid in MIATL but not in IATL.

Theorem 1 shows that we cannot extend a proof theory for ATL to get a proof theory for either IATL or MIATL (point 1); that we cannot specialise an ATL proof theory either (point 2); and that MIATL extends IATL but not the other way around. The latter point is interesting, because in the case of standard ATL the theorems of the logic defined with memoryless and memory-based strategies, respectively, coincide. In the case of irrevocable strategies, there are strictly more theorems in the case of memory-based strategies.

4.2 Special axioms for IATL

$\langle\!\langle \emptyset \rangle\!\rangle \mathcal{X} \phi$ expresses the fact that ϕ is inevitably true in the next state. We can say that *a coalition C has no choice wrt a formula ϕ* if C can achieve ϕ in the next state precisely when ϕ is inevitable:

$$NoChoice(C, \phi) = \langle\!\langle C \rangle\!\rangle \mathcal{X} \phi \leftrightarrow \langle\!\langle \emptyset \rangle\!\rangle \mathcal{X} \phi.$$

Now, for any formula ϕ and coalitions C, C' we define:

$$Irrevocable(C, C', \phi) = [\![C]\!] \mathcal{G} NoChoice(C', \phi).$$

Thus, after a coalition C has committed to a strategy, all its sub-coalitions have no choice any more. Formally:

Proposition 3. *For any formula ϕ and coalitions C, C' such that $C' \subseteq C$:*

$$\models_{\text{IATL}} Irrevocable(C, C', \phi).$$

The above can be generalized to:

$$\models_{\text{IATL}} [\![C]\!] \mathcal{G} (\langle\!\langle C' \rangle\!\rangle \mathcal{X} \phi \leftrightarrow \langle\!\langle C' \setminus C \rangle\!\rangle \mathcal{X} \phi)$$

for any coalitions C, C'.

We denote the last formula scheme by Normal. Given any formula ϕ, by repeated use of Normal, one can subsequently remove from the internal coalition C' in every subformula $\langle\!\langle C\rangle\!\rangle \ldots \langle\!\langle C'\rangle\!\rangle \ldots$ of ϕ all agents in $C \cap C'$, and thus eventually show the following proposition.

Definition 3. *An* ATL*-formula* ϕ *is* normal *if it contains no nested occurrences of strategic operators with intersecting coalitions, i.e., it contains no subformula of the type* $\langle\!\langle A_1\rangle\!\rangle \ldots \langle\!\langle A_2\rangle\!\rangle \ldots$ *where* $A_1 \cap A_2 \neq \emptyset$*.*

Proposition 4. *Every formula is* IATL*-equivalent to an effectively computable normal formula.*

Consequently, testing IATL-satisfiability of any formula can be effectively reduced to testing IATL-satisfiability of a normal formula, which in turn can be established by suitably modifying the construction of alternating tree automata associated with such formulae, presented in [4].

The importance of normal formulae derives from the fact that their IATL-semantics is essentially compositional and they behave in many respects as in ATL. However, one additional effect of the irrevocable strategies is that once all agents commit to their strategies, only one computation remains possible, i.e. the system becomes *deterministic*, which admits an additional ATL-valid scheme. This yields a new scheme of IATL-valid formulae which are only ATL-valid in deterministic systems:

$$[\![\Sigma]\!]\mathcal{G}([\![\emptyset]\!]\mathcal{X}p \rightarrow \langle\!\langle\emptyset\rangle\!\rangle\mathcal{X}p).$$

Thus, within the scope of $\langle\!\langle\Sigma\rangle\!\rangle$ all cooperation modalities are completely trivialized and can be simply omitted; the result is an LTL formula, evaluated on the unique computation determined by the committed collective strategy of all agents.

4.3 Towards tableaux for IATL and MIATL

Two decision procedures for ATL have been developed in [4] – one using on alternating automata, and the other - based on a finite 'pseudo-model' property for every satisfiable ATL-formula. Both methods can be converted (currently in preparation [5]) into tableau-like proof systems for ATL[5], the relative efficiency of which is still under investigation.

The adaptation of a tableau for ATL to IATL raises some technical difficulties, related to the control of irrevocability of strategies. The modelling within the tableau of the memory-based strategies in MIATL presents an extra complication. These are currently being addressed and will be discussed in a subsequent full paper.

5 Conclusions

This paper is a preliminary report on work towards a proof theory for multi-agent logics with irrevocable strategies. This proof theory is essentially affected by the irrevocable

---

[5] Also, in [6] Hansen develops 'tableau game' for ATL, similar to that constructed by Niwinski and Walukiewicz for the modal mu-calculus, but according to the author her system has a flaw.

strategies as compared to the original ATL; in particular, neither of ATL and IATL subsumes the other. Furthermore, unlike the case of revocable strategies, memory plays an important role in the semantics based on irrevocable strategies, distinguishing the two versions, IATL and MIATL.

6 Acknowledgements

T. Ågotnes acknowledges the financial support of the Research Council of Norway through research grant 166525/V30. V. Goranko thanks Sophie Pinchinat for insightful discussions on the topic of this paper and acknowledges the financial support of the National Research Foundation of South Africa through research grant GUN 2072436.

References

1. T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time Temporal Logics with Irrevocable Strategies, to appear in: Proc. of TARK'2007.
2. R. Alur, T. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *JACM*, 49:672–713, 2002.
3. V. Goranko and W. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2):241–280, 2004.
4. V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of the Alternating-time Temporal Logic. *Theor. Comp. Sci.*, 353:93–117, 2006.
5. V. Goranko, Complete and terminating tableaux for ATL, 2007, in preparation.
6. H. Hansen, Tableau Games for Coalition Logic and Alternating-time Temporal Logic: theory and implementation, master thesis, ILLC, 2004.
7. P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electr. Notes in Theor. Comp. Sci.*, 85(2), 2004.
8. W. van der Hoek, W. Jamroga, and M. Wooldridge. A logic for strategic reasoning. In *Proceedings of AAMAS'05*, pages 157–164, 2005.
9. W. van der Hoek and M. Pauly. Modal logic for games and information. In J. van Benthem, P. Blackburn, and F. Wolter (eds), *Handbook of Modal Logic*, pages 1077–1148. Elsevier, 2007.
10. W. van der Hoek, M. Roberts, and M. Wooldridge. Social laws in alternating time: Effectiveness, feasibility and synthesis. *Synthese*, 2005.