

Comparing Semantics of Logics for Multi-agent Systems

Valentin Goranko

Department of Mathematics, Rand Afrikaans University

e-mail: vfg@na.rau.ac.za

Wojciech Jamroga

Parlevink Group, University of Twente, Netherlands

Institute of Mathematics, University of Gdansk, Poland

e-mail: jamroga@cs.utwente.nl

Abstract. We draw parallels between several closely related logics that combine – in different proportions – elements of game theory, computation tree logics, and epistemic logics to reason about agents and their abilities. These are: the coalition game logics CL and ECL introduced by Pauly in 2000, the alternating-time temporal logic ATL developed by Alur, Henzinger and Kupferman between 1997 and 2002, and the alternating-time temporal epistemic logic ATEL by van der Hoek and Wooldridge (2002). In particular, we establish some subsumption and equivalence results for their semantics, as well as interpretation of the alternating-time temporal epistemic logic into ATL.

The focus in this paper is on models: alternating transition systems, multi-player game models (alias concurrent game structures) and coalition effectivity models turn out to be intimately related, while alternating epistemic transition systems share much of their philosophical and formal apparatus. Our approach is constructive: we present ways to transform between different types of models and languages.

Keywords: multi-agent systems, game theory, coalition logics, alternating-time temporal logics, epistemic logic.

1. Introduction

In this study we offer a comparative analysis of several recent logical enterprises that aim at modeling multi-agent systems. Most of all, the *coalition game logic* CL and its extended version ECL (Pauly, 2002; Pauly, 2000b; Pauly, 2001), and the *Alternating-time Temporal Logic* ATL (Alur et al., 1997; Alur et al., 1998a; Alur et al., 2002) are studied. These turn out to be intimately related, which is not surprising since all of them deal with essentially the same type of scenarios, viz. a *set of agents* (players, system components) taking actions, simultaneously or in turns, on a common set of states – and thus effecting transitions between these states. The game-theoretic aspect is very prominent in both approaches; furthermore, in both frameworks the agents pursue certain goals with their actions and in that pursuit they can form *coalitions*. In both enterprises the objective is to develop



© 2003 Kluwer Academic Publishers. Printed in the Netherlands.

formal tools for reasoning about such coalitions of agents and their ability to achieve specified outcomes in these action games.

An extension of ATL, called *Alternating-time Temporal Epistemic Logic* (ATEL) was introduced in (van der Hoek and Wooldridge, 2002) in order to enable reasoning about agents acting under incomplete information. Although the semantics for ATEL is still under debate, the original version of that logic is certainly worth investigating. It turns out that, while extending ATL, ATEL can be embedded into the former in the sense that there is a translation of models and formulas of ATEL into ATL that preserves the satisfiability of formulas. This does not imply that logics like ATEL are redundant, of course – in fact, the way of expressing epistemic facts in ATL is purely technical, and the resulting formulas look rather unnatural. Similarly, each of the three alternative semantics for ECL and ATL, investigated here, has its own drawbacks and offers different advantages for practical use.

The rest of the paper is organized as follows: first, we offer a brief summary of the basic concepts from game theory; then we introduce the main “actors” of our study – logics and structures that have been used for modeling multi-agent systems in temporal perspective. In order to make the paper self-contained we have included all relevant definitions from (Pauly, 2002; Pauly, 2001; Alur et al., 1998a; Alur et al., 2002; van der Hoek and Wooldridge, 2002).¹ In Sections 3 and 4 the relationships between these logics and structures are investigated in a formal way. The main results are the following:

- We show that specific classes of multi-player game models are equivalent to some types of alternating transition systems.
- We demonstrate that ATL subsumes CL as well as ECL.
- We show that the three alternative semantics for Alternating-time Temporal Logic and Coalition Logics (based on multi-player game models, alternating transition systems and coalition effectiveness models) are equivalent.
- We show that formulas and models of ATEL can be translated into its fragment ATL.

The paper partly builds on previous work of ours, included in (Goranko, 2001) and (Jamroga, 2003).

¹ We make small notational changes here and there to make the differences and common features between the models and languages clearer and easier to see.

2. Models and logics of strategic ability

The logics studied here have a few things in common. They are intended for reasoning about various aspects of multi-agent systems and multi-player games, they are multi-modal logics, they have been obviously inspired by game theory, and they are based on the temporal logic approach. We present and discuss the logics and their models in this section. A broader survey of logic-based approaches to multi-agent systems can be found in (van der Hoek and Wooldridge, 2003b).

2.1. BASIC INFLUENCES

2.1.1. *Classical game theory*

Logics of agents and action build upon several important concepts from game theory, most of them going back to the 40s and the seminal book (von Neumann and Morgenstern, 1944). We will start with an informal survey of these concepts, following mostly (Hart, 1992). An interested reader is referred to (Aumann and Hart, 1992; Osborne and Rubinstein, 1994) for a more extensive introduction to game theory.

In game theory, a game is usually presented in its extensive and/or strategic form. The *extensive form* defines the game via a tree of possible positions in the game (*states*), game moves (*choices*) available to players, and the outcome (*utility* or *payoff*) that players gain at each of the final states. These games are usually *turn-based*, i.e. every state is assigned a player who controls the choice of the next move, so the players are taking turns. A *strategy* for player a specifies a 's choices at the states controlled by a .

The *strategic form* consists of a matrix that presents the payoffs for all combinations of players' strategies. It presents the whole game in a "snapshot" as if it was played in one single move, while the extensive form emphasizes control and information flow in the game.

EXAMPLE 1. Consider a variant of the *matching pennies* game. There are two players, each with a coin: first a_1 chooses to show the heads (action h) or tails (t), then a_2 does. If both coins are heads up or both coins are tails up, then a_1 wins (and gets score of 1) and a_2 loses (score 0). If the coins show different sides, then a_2 is the winner.

The extensive and strategic forms for this game are shown in Figure 1A. The strategies define agent's choices at all "his" nodes, and are labeled appropriately: $q_1h q_2t$ denotes, for instance, a strategy for a_2 in which the player chooses to show heads whenever the current state of the game is q_1 , and tails at q_2 . Note that – using this strategy – a_2 wins regardless of the first move from a_1 .

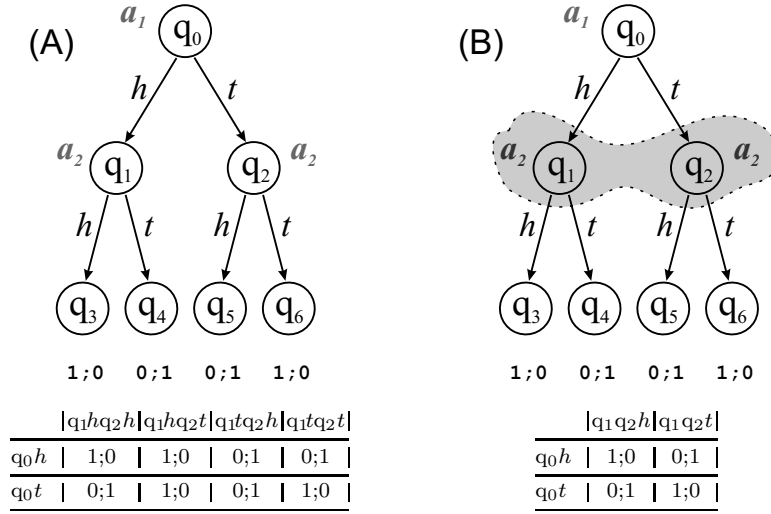


Figure 1. Extensive and strategic form of the matching pennies game: (A) perfect information case; (B) a_1 does not show his coin before the end of the game.

The information available to agents is incomplete in many games. Classical game theory handles this kind of uncertainty through partitioning every player's nodes into so called information sets. An *information set* for player a is a set of states that are indistinguishable for a . Traditionally, information sets are defined only for the states in which a chooses the next step. Now a strategy assigns choices to information sets rather than separate states, because players are supposed to choose the same move for all the situations they cannot distinguish.

EXAMPLE 2. Suppose that a_1 does not show his coin to a_2 before the end of the game. Then nodes q_1 and q_2 belong to the same information set of a_2 , as shown in Figure 1B. No player has a strategy that guarantees his win any more.

A general remark is in order here. The concept of coalitional game traditionally considered in game theory where every possible coalition is assigned a real number (its *worth*), differs somewhat from the one considered here. In this study we are rather concerned with *qualitative* aspects of *game structures* rather than with *quantitative* analysis of specific *games*. It should be clear, however, that these two approaches are in agreement and can be easily put together. Indeed, the intermediate link between them is the notion of (qualitative) *effectivity function* (Pauly, 2002). That notion naturally transfers over to alternating transition systems, thus providing a framework for purely game-theoretic treatment of alternating temporal logics.




Figure 2. Transitions of the variable controller/client system, together with the tree of possible computations.

2.1.2. Computational Tree Logic and epistemic logic

Apart from game theory, the concepts investigated in this paper are strongly influenced by modal logics of computations (such as the *computation tree logic* CTL) and beliefs (*epistemic logic*). CTL (Emerson, 1990; Huth and Ryan, 2000) involves several operators for temporal properties of computations in transition systems: A (*for all paths*), E (*there is a path*), X (*nexttime*), F (*sometime*), G (*always*) and U (*until*). “Paths” refer to alternative courses of events that may happen in the future; nodes on a path denote states of the system in subsequent moments of time along this particular course. Typically, paths are interpreted as sequences of successive states of computations.

EXAMPLE 3. As an illustration, consider a system with a binary variable x . In every step, the variable can retain or change its value. The states and possible transitions are shown in Figure 2. There are two propositions available to observe the value of x : “ $x=0$ ” and “ $x=1$ ”. Then, for example, $EFx=1$ is satisfied in every state of the system: there is a path such that x will have the value of 1 at some moment. However, the above is not true for *every* possible course of action: $\neg AFx=1$.

It is important to distinguish between the *computational structure*, defined explicitly in the model, and the *behavioral structure*, i.e. the model of how the system is supposed to behave in time (Schnoebelen, 2003). In many temporal models the computational structure is finite, while the implied behavioral structure is infinite. The computational structure can be seen as a way of defining the tree of possible (infinite) computations that may occur in the system. The way the computa-

tional structure unravels into a behavioral structure (computation tree) is shown in Figure 2, too.

Epistemic logic offers the notion of *epistemic accessibility relation* that generalizes information sets, and introduces operators for talking about individual and collective knowledge. Section 4 describes them in more detail; a reader interested in a comprehensive exposition on epistemic logic can be also referred to the seminal book by Fagin, Halpern, Moses and Vardi (Fagin et al., 1995), or to (van der Hoek and Verbrugge, 2002) for a survey.

2.2. COALITION GAME LOGICS AND MULTI-PLAYER GAME MODELS

Coalition logic (CL), introduced in (Pauly, 2000b; Pauly, 2002), formalizes reasoning about powers of coalitions in strategic games. It extends the classical propositional logic with a family of (non-normal) modalities $[A]$, $A \subseteq \text{Agt}$, where Agt is a fixed set of players. Intuitively, $[A]\varphi$ means that coalition A can *enforce* an outcome state satisfying φ .

2.2.1. Multi-player strategic game models

Game frames (Pauly, 2002), represent multi-player strategic games where sets of players can form coalitions in attempts to achieve desirable outcomes. Game frames are based on the notion of a *strategic game form* – a tuple $\langle \text{Agt}, \{\Sigma_a \mid a \in \text{Agt}\}, Q, o \rangle$ consisting of:

- a non-empty finite set of *agents* (or *players*) Agt ,
- a family of (non-empty) sets of *actions* (*choices*, *strategies*) Σ_a for each player $a \in \text{Agt}$,
- a non-empty set of *states* Q ,
- an *outcome function* $o : \prod_{a \in \text{Agt}} \Sigma_a \rightarrow Q$ which associates an outcome state in Q to every combination of choices from all the players. By a *collective choice* σ_A we will denote a tuple of choices $\langle \sigma_a \rangle_{a \in A}$ (one for each player from $A \subseteq \text{Agt}$), and we will be writing $o(\sigma_A, \sigma_{\text{Agt} \setminus A})$ with the presumed meaning.

REMARK 1. *Note that the notion of “strategy” in strategic game forms is local, wrapped into one-step actions. It differs from the notion of “strategy” in extensive game forms (used in the semantics of ATL) which represents a global, conditional plan of action. To avoid confusion, we will refer to the local strategies as choices, and use the term collective choice instead of strategy profile from (Pauly, 2002) to denote a combination of simultaneous choices from several players.*




Figure 3. Transitions of the variable controller/client system.

REMARK 2. A strategic game form defines the choices and transitions available at a particular state of the game. If the identity of the state does not follow from the context in an obvious way, we will use indices to indicate which state they refer to.

The set of all strategic game forms for players Agt over states Q will be denoted by Γ_Q^{Agt} . A *multi-player game frame* for a set of players Agt is a pair $\langle Q, \gamma \rangle$ where $\gamma : Q \rightarrow \Gamma_Q^{\text{Agt}}$ is a mapping associating a strategic game form with each state in Q . A *multi-player game model* (MGM) for a set of players Agt over a set of propositions Π is a triple $M = \langle Q, \gamma, \pi \rangle$ where $\langle Q, \gamma \rangle$ is a multi-player game frame, and $\pi : Q \rightarrow \mathcal{P}(\Pi)$ is a *valuation* labeling each state from Q with the set of propositions that are true at that state.

EXAMPLE 4. Consider a variation of the system with binary variable x from Example 3. There are two processes: the controller (or server) s can enforce the variable to retain its value in the next step, or let the client change the value. The client c can request the value of x to be 0 or 1. The players proceed with their choices simultaneously. The states and transitions of the system as a whole are shown in Figure 3.

Again, we should make the distinction between computational and behavioral structures. The multi-player game model unravels into a computation tree in a way analogous to CTL models (cf. Figure 2).

2.2.2. Coalition logic

Formulas of CL are defined recursively as:

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \psi \mid [A]\varphi.$$

where $p \in \Pi$ is a proposition, and $A \subseteq \text{Agt}$ is a group of agents. The semantics of CL can be given via the clauses:

- $M, q \models p$ iff $p \in \pi(q)$ for atomic propositions p ;
- $M, q \models [A]\varphi$ iff there is a collective choice σ_A such that for every collective choice $\sigma_{\text{Agt} \setminus A}$, we have $M, o_q(\sigma_A, \sigma_{\text{Agt} \setminus A}) \models \varphi$.

EXAMPLE 5. Consider the variable client/server system from Example 4. The following CL formulas are valid in this model (i.e. true in every state of it):

1. $(x=0 \rightarrow [s]x=0) \wedge (x=1 \rightarrow [s]x=1)$: the server can enforce the value of x to remain the same in the next step;
2. $x=0 \rightarrow \neg[c]x=1$: c cannot change the value from 0 to 1 on his own;
3. $x=0 \rightarrow \neg[s]x=1$: s cannot change the value on his own either;
4. $x=0 \rightarrow [s, c]x=1$: s and c can cooperate to change the value.

2.2.3. Logics for local and global effectivity of coalitions

In CL, the operators $[A]\varphi$ can express *local effectivity* properties of coalitions, i.e. their powers to force outcomes in single ‘rounds’ of the game. Pauly extends in (Pauly, 2000b) CL to the *Extended Coalition Logic* ECL with iterated operators for *global effectivity* $[A^*]\varphi$ expressing the claim that coalition A has a collective strategy to *maintain the truth of φ throughout the entire game*. In our view, and in the sense of Remark 1, both systems formalize different aspects of reasoning about powers of coalitions: CL can be thought as reasoning about *strategic game forms*, while ECL rather deals with *extensive game forms*, representing sequences of moves, collectively effected by the players’ actions. Since ECL can be embedded as a fragment of ATL (as presented in Section 2.4), we will not discuss it separately here.

2.3. ALTERNATING-TIME TEMPORAL LOGIC AND ITS MODELS

Game-theoretic scenarios can occur in various situations, one of them being *open computer systems* such as computer networks, where the different components can act as relatively autonomous agents, and computations in such systems are effected by their combined actions. The *Alternating-time Temporal Logics* ATL and ATL*, introduced in (Alur et al., 1997), and later refined in (Alur et al., 1998a; Alur et al., 2002), are intended to formalize reasoning about computations in such open systems which can be enforced by coalitions of agents, in a way generalizing logics CTL and CTL*.

2.3.1. The logics ATL and ATL*

In ATL* a class of *cooperation modalities* $\langle\langle A \rangle\rangle$ replaces the path quantifiers E and A. The common-sense reading of $\langle\langle A \rangle\rangle\Phi$ is:

“The group of agents A have a collective strategy to enforce Φ regardless of what all the other agents do”.

ATL is the fragment of ATL^* subjected to the same syntactic restrictions which define CTL as a fragment of CTL^* , i.e. every temporal operator must be immediately preceded by exactly one cooperation modality. The original CTL^* operators E and A can be expressed in ATL^* with $\langle\langle \text{Agt} \rangle\rangle$ and $\langle\langle \emptyset \rangle\rangle$ respectively, but between both extremes one can express much more about the abilities of particular agents and groups of agents. Since model-checking for ATL^* requires 2EXPTIME , but it is linear for ATL, ATL is more useful for practical applications, and we will not discuss ATL^* in this paper. Formally, the recursive definition of ATL formulas is:

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \psi \mid \langle\langle A \rangle\rangle X\varphi \mid \langle\langle A \rangle\rangle G\varphi \mid \langle\langle A \rangle\rangle \varphi \mathcal{U}\psi$$

The ‘‘sometime’’ operator F can be defined in the usual way as: $\langle\langle A \rangle\rangle F\varphi \equiv \langle\langle A \rangle\rangle \top \mathcal{U}\varphi$.

It should be noted that at least three different versions of semantic structures for ATL have been proposed by Alur and colleagues in the last 7 years. The earliest version (Alur et al., 1997), involves definitions of a synchronous turn-based structure and an asynchronous structure in which every transition is controlled by a single agent. The next paper (Alur et al., 1998a) defines general structures called *alternating transition systems* where the agents’ choices are identified with the sets of possible outcomes. In the *concurrent game structures* from (Alur et al., 2002), labels for choices are introduced and the transition function is simplified. The above papers share the same title and they are usually cited incorrectly in the literature as well as citation indices, which may lead to some confusion.

2.3.2. Alternating transition systems

Alternating transition systems – building on the concept of *alternation* developed in (Chandra et al., 1981) – formalize systems of transitions effected by collective actions of all agents involved. In the particular case of one agent (the *system*), alternating transition systems are reduced to ordinary transition systems, and ATL reduces to CTL.

An *alternating transition system* (ATS) is a tuple $T = (\Pi, \text{Agt}, Q, \pi, \delta)$ where:

- Π is a set of (atomic) *propositions*, Agt is a non-empty finite set of *agents*, Q is a non-empty set of *states*, and $\pi : Q \rightarrow \mathcal{P}(\Pi)$ is a *valuation* of propositions;
- $\delta : Q \times \text{Agt} \rightarrow \mathcal{P}(\mathcal{P}(Q))$ is a *transition function* mapping a pair $\langle \text{state}, \text{agent} \rangle$ to a non-empty family of choices of possible next states. The idea is that at state q an agent a chooses a set




Figure 4. An ATS for the controller/client problem

$Q_a \in \delta(q, a)$ thus forcing the outcome state to be from Q_a . The resulting transition leads to a state which is in the intersection of all Q_a for $a \in \text{Agt}$ and so it reflects the mutual will of all agents. Since the system is required to be deterministic (given the state and the agents' decisions), $Q_1 \cap \dots \cap Q_k$ must always be a singleton.²

DEFINITION 1. A state $q_2 \in Q$ is a successor of q_1 if, whenever the system is in q_1 , the agents can cooperate so that the next state is q_2 , i.e. there are choice sets $Q_a \in \delta(q_1, a)$, for each $a \in \text{Agt}$ such that $\bigcap_{a \in \text{Agt}} Q_a = \{q_2\}$. The set of successors of q will be denoted by Q_q^{suc} .

DEFINITION 2. A computation in T is an infinite sequence of states $q_0 q_1 \dots$ such that q_{i+1} is a successor of q_i for every $i \geq 0$. A q -computation is a computation starting from q .

2.3.3. Semantics of ATL based on alternating transition systems

DEFINITION 3. A strategy for agent a is a mapping $f_a : Q^+ \rightarrow \mathcal{P}(Q)$ which assigns to every non-empty sequence of states q_0, \dots, q_n a choice set $f_a(q_0 \dots q_n) \in \delta(q_n, a)$. The function specifies a a 's decisions for every possible (finite) history of system transitions. A collective strategy for a set of agents $A \subseteq \text{Agt}$ is just a tuple of strategies (one per agent from A): $F_A = \langle f_a \rangle_{a \in A}$.

Now, $\text{out}(q, F_A)$ denotes the set of outcomes of F_A from q , i.e. the set of all q -computations in which group A has been using F_A .

² Determinism is not a crucial issue here, as it can be easily imposed by introducing a new, fictitious agent, "Nature", which settles all non-deterministic transitions.

REMARK 3. *This notion of strategy can be specified as “perfect recall strategy”, where the whole history of the game is considered when the choice of the next move is made by the agents. The other extreme alternative is a “memoryless strategy” where only the current state is taken in consideration; further variations on “limited memory span strategies” are possible. While the choice of one or another notion of strategy affects the semantics of the full ATL*, it is not difficult to see that both perfect recall strategies and memoryless strategies eventually yield equivalent semantics for ATL.*

Let $\Lambda[i]$ denote the i th position in computation Λ . The definition of truth of an ATL formula at state q of an ATS $T = \langle \Pi, \text{Agt}, Q, \pi, \delta \rangle$ follows through the below clauses. Informally speaking, $T, q \models \langle\langle A \rangle\rangle \Phi$ iff there exists a collective strategy F_A such that Φ is satisfied for all computations from $\text{out}(F_A, q)$.

(**A, X**) $T, q \models \langle\langle A \rangle\rangle X\varphi$ iff there exists a collective strategy F_A such that for every computation $\Lambda \in \text{out}(q, F_A)$ we have $T, \Lambda[1] \models \varphi$;

(**A, G**) $T, q \models \langle\langle A \rangle\rangle G\varphi$ iff there exists a collective strategy F_A such that for every $\Lambda \in \text{out}(q, F_A)$ we have $T, \Lambda[i] \models \varphi$ for every $i \geq 0$.

(**A, U**) $T, q \models \langle\langle A \rangle\rangle \varphi \mathcal{U} \psi$ iff there exists a collective strategy F_A such that for every $\Lambda \in \text{out}(q, F_A)$ there is $i \geq 0$ such that $T, \Lambda[i] \models \psi$ and for all j such that $0 \leq j < i$ we have $T, \Lambda[j] \models \varphi$.

EXAMPLE 6. An ATS for the variable client/server system is shown in Figure 4. The following ATL formulas are valid in this model:

1. $(x=0 \rightarrow \langle\langle s \rangle\rangle X x=0) \wedge (x=1 \rightarrow \langle\langle s \rangle\rangle X x=1)$: the server can enforce the value of x to remain the same in the next step;
2. $x=0 \rightarrow \neg \langle\langle c \rangle\rangle F x=1 \wedge \neg \langle\langle s \rangle\rangle F x=1$: neither c nor s can change the value from 0 to 1, even in multiple steps;
3. $x=0 \rightarrow \langle\langle s, c \rangle\rangle F x=1$: s and c can cooperate to change the value.

2.3.4. Semantics of ATL based on concurrent game structures and multi-player game models

(Alur et al., 2002) redefines ATL models as *concurrent game structures*:

$$M = \langle k, Q, \Pi, \pi, d, o \rangle,$$

where k is the number of players (so Agt can be taken to be $\{1, \dots, k\}$), the decisions available to player a at state q are labeled with natural

numbers up to $d_a(q)$ (so $\Sigma_a(q)$ can be taken to be $\{1, \dots, d_a(q)\}$); finally, a complete tuple of decisions $\langle \alpha_1, \dots, \alpha_k \rangle$ from all the agents in state q implies a deterministic transition according to the transition function $o(q, \alpha_1, \dots, \alpha_k)$. In a concurrent game structure the type of a strategy function is slightly different since choices are abstract entities indexed by natural numbers now, and a strategy is a mapping $f_a : Q^+ \rightarrow \mathbb{N}$ such that $f_a(\lambda q) \leq d_a(q)$. The rest of the semantics looks exactly the same as for alternating transition systems.

REMARK 4. *Clearly, concurrent game structures are equivalent to Pauly's multi-player game models; they differ from each other only in notation.³ Thus, the ATL semantics can be as well based on MGMs, and the truth definitions look exactly the same as for alternating transition systems (see Section 2.3.3). We leave rewriting the definitions of a strategy, collective strategy and outcome set in terms of multi-player game models to the reader. The next section shows how this shared semantics can be used to show that ATL subsumes coalition logics.*

2.4. EMBEDDING CL AND ECL INTO ATL

It turns out that both CL and ECL are strictly subsumed by ATL in terms of the shared semantics based on multi-player game models.⁴ Indeed, there is a translation of formulas of ECL into ATL, which becomes obvious once the ATL semantic clause (A, X) is rephrased as:

$$[A] T, q \models \langle\langle A \rangle\rangle X\varphi \text{ iff there exists a collective choice } F_A = \{f_a\}_{a \in A} \text{ such that for every collective choice } F_{\text{Agt} \setminus A} = \{f_a\}_{a \in \text{Agt} \setminus A}, \text{ we have } T, s \models \varphi, \text{ where } \{s\} = \bigcap_{a \in A} f_a(q) \cap \bigcap_{a \in \text{Agt} \setminus A} f_a(q)$$

which is precisely the truth-condition for $[A]\varphi$ in the coalition logic CL.

Thus, CL embeds in a straightforward way as a simple fragment of ATL by translating $[A]\varphi$ into $\langle\langle A \rangle\rangle X\varphi$. Accordingly, $[C^*]\varphi$ translates into ATL as $\langle\langle A \rangle\rangle G\varphi$, which follows from the fact that each of $[C^*]\varphi$ and $\langle\langle A \rangle\rangle G\varphi$, is the greatest fixpoint of the same operator over $[C]\varphi$ and $\langle\langle A \rangle\rangle X\varphi$ respectively (see Section 2.5). In consequence, ATL subsumes ECL as the fragment ATL_{XG} involving only $\langle\langle A \rangle\rangle X\varphi$ and $\langle\langle A \rangle\rangle G\varphi$.

³ The only real difference is that the set of states Q and the sets representing agents' choices are explicitly required to be finite in the concurrent game structures, while MGMs and ATs are not constrained this way. However, these requirements are not essential and can be easily omitted if necessary.

⁴ Note that the coalition logic-related notions of choice and collective choice can be readily expressed in terms of alternating transition systems, which immediately leads to a semantics for CL based on ATS, too. Thus, ATL and the coalition logics share the semantics based on alternating transition systems as well.

We will focus on ATL, and will simply regard CL and ECL as its fragments throughout the rest of the paper.

2.5. EFFECTIVITY FUNCTIONS AND COALITION EFFECTIVITY MODELS AS ALTERNATIVE SEMANTICS FOR ATL

As mentioned earlier, game theory usually measures the powers of coalitions *quantitatively*, and characterizes the possible outcomes in terms of *payoff profiles*. That approach can be easily transformed into a *qualitative* one, where the payoff profiles are encoded in the outcome states themselves and each coalition is assigned a *preference order* on these outcome states. Then, the power of a coalition can be measured in terms of *sets of states* in which it can force the actual outcome of the game (i.e. sets for which it is *effective*), thus defining another semantics for ATL, based on so called *coalition effectivity models* (introduced by Pauly for the coalition logics CL and ECL). This semantics is essentially a monotone neighborhood semantics for non-normal multi-modal logics, and therefore it enables the results, methods and techniques already developed for modal logics to be applied here as well.

DEFINITION 4. (Pauly, 2002) *A (local) effectivity function is a mapping of type $e : \mathcal{P}(\text{Agt}) \rightarrow \mathcal{P}(\mathcal{P}(Q))$.*

The idea is that we associate with each set of players the family of outcome sets for which their coalition is effective. However, the notion of effectivity function as defined above is abstract and not every effectivity function corresponds to a real strategic game form. Those which do can be characterized with the following conditions (Pauly, 2002):

1. *Liveness*: for every $A \subseteq \text{Agt}$, $\emptyset \notin e(A)$.
2. *Termination*: for every $A \subseteq \text{Agt}$, $Q \in e(A)$.
3. *Agt-maximality*: if $X \notin e(\text{Agt})$ then $Q \setminus X \in e(\emptyset)$ (if X cannot be effected by the grand coalition of players, then $Q \setminus X$ is inevitable).
4. *Outcome-monotonicity*: if $X \subseteq Y$ and $X \in e(A)$ then $Y \in e(A)$.
5. *Super-additivity*: for all $A_1, A_2 \subseteq \text{Agt}$ and $X_1, X_2 \subseteq Q$, if $A_1 \cap A_2 = \emptyset$, $X_1 \in e(A_1)$, and $X_2 \in e(A_2)$, then $X_1 \cap X_2 \in e(A_1 \cup A_2)$.

We note that super-additivity and liveness imply *consistency of the powers*: for any $A \subseteq \text{Agt}$, if $X \in e(A)$ then $Q \setminus X \notin e(\text{Agt} \setminus A)$.

DEFINITION 5. (Pauly, 2002) *An effectivity function e is called playable if conditions (1)–(5) hold for e .*

\emptyset	$\{s\}$	$\{c\}$	$\{s, c\}$
$\{\{q_0, q_1\}\}$	$\{\{q_0\}, \{q_0, q_1\}\}$	$\{\{q_0\}, \{q_0, q_1\}\}$	$\{\{q_0\}, \{q_1\}, \{q_0, q_1\}\}$

Figure 5. A coalition effectivity function for the variable client/server system.

DEFINITION 6. (Pauly, 2002) *An effectivity function e is the effectivity function of a strategic game form γ if it associates with each set of players A from γ the family of outcome sets $\{Q_1, Q_2, \dots\}$, such that for every Q_i the coalition A has a collective choice to ensure that the next state will be in Q_i .*

THEOREM 5. (Pauly, 2002) *An effectivity function is playable iff it is the effectivity function of some strategic game form.*

EXAMPLE 7. Figure 5 presents a playable effectivity function that describes powers of all the possible coalitions for the variable server/client system from Example 4, and state q_0 .

DEFINITION 7. (Pauly, 2002) *A coalition effectivity frame is a triple $\mathcal{F} = \langle \text{Agt}, Q, E \rangle$ where Agt is a set of players, Q is a non-empty set of states and $E : Q \rightarrow (\mathcal{P}(\text{Agt}) \rightarrow \mathcal{P}(\mathcal{P}(Q)))$ is a mapping which associates an effectivity function with each state. We shall write $E_q(A)$ instead of $E(q)(A)$. A coalition effectivity model (CEM) is a tuple $\mathcal{E} = \langle \text{Agt}, Q, E, \pi \rangle$ where $\langle \text{Agt}, Q, E \rangle$ is a coalition effectivity frame and π is a valuation of the atomic propositions over Q .*

DEFINITION 8. *A coalition effectivity frame (resp. coalition effectivity model) is standard if it contains only playable effectivity functions.*

Thus, coalition effectivity models provide semantics of CL by means of the following truth definition (Pauly, 2002):

$$\mathcal{E}, q \models [A]\varphi \text{ iff } \{s \in \mathcal{E} \mid \mathcal{E}, s \models \varphi\} \in E_q(A).$$

This semantics can be accordingly extended to semantics for ECL (Pauly, 2001) and ATL (Goranko, 2001) by defining effectivity functions for the global effectivity operators in extensive game forms, where they indicate the outcome sets for which the coalitions have long-term *strategies* to effect. This extension can be done using the following fixpoint characterizations of $\langle\langle A \rangle\rangle G\varphi \leftrightarrow \varphi \wedge \langle\langle A \rangle\rangle X\langle\langle A \rangle\rangle G\varphi$, and $\langle\langle A \rangle\rangle \varphi \mathcal{U} \psi \leftrightarrow \psi \vee \varphi \wedge \langle\langle A \rangle\rangle X\langle\langle A \rangle\rangle \varphi \mathcal{U} \psi$ as follows:

$$\begin{aligned} \langle\langle A \rangle\rangle G\varphi &:= \nu \mathbf{Z}. \varphi \wedge \langle\langle A \rangle\rangle X \mathbf{Z}, \\ \langle\langle A \rangle\rangle \varphi \mathcal{U} \psi &:= \mu \mathbf{Z}. \psi \vee \varphi \wedge \langle\langle A \rangle\rangle X \mathbf{Z}. \end{aligned}$$

3. Equivalence of the different semantics for ATL

In this section we compare the semantics for Alternating-time Temporal Logic, based on alternating transition systems and multi-player game models – and show their equivalence (in the sense that we can transform the models both ways while preserving satisfiability of ATL formulas). Further, we show that these semantics are both equivalent to the semantics based on coalition effectivity models.

The transformation from alternating transition systems to multi-player game models is easy: in fact, for every ATS, an isomorphic MGM can be constructed via re-labeling transitions (see Section 3.2). Construction the other way round is more sophisticated: first, we observe that all multi-player game models obtained from alternating transition systems satisfy a special condition we call *convexity* (Section 3.2); then we show that for every convex MGM, an isomorphic ATS can be obtained (Section 3.3). Finally, we demonstrate that for every arbitrary multi-player game model a convex MGM can be constructed that satisfies the same formulas of ATL (Section 3.4).

We show also that the transformations we propose preserve the property of being a turn-based structure, and that they transform injective MGMs into lock-step synchronous ATSS and vice versa.

3.1. SOME SPECIAL TYPES OF ATSS AND MGMS

DEFINITION 9. (Pauly, 2002) *A strategic game form $\langle \text{Agt}, \{\Sigma_a \mid a \in \text{Agt}\}, Q, o \rangle$ is an a -dictatorship if there is a player $a \in \text{Agt}$ who determines the outcome state of the game, i.e.*

$$\forall \sigma_a \in \Sigma_a \exists q \in Q \forall \sigma_{\text{Agt} \setminus \{a\}} o(\sigma_a, \sigma_{\text{Agt} \setminus \{a\}}) = q.$$

*An MGM $\langle Q, \gamma, \pi \rangle$ is turn-based if every $\gamma(q)$ is a dictatorship.*⁵

We note that the notion of a -dictatorship is quite strong: it presumes that *any* choice of the dictator forces a chosen state as the outcome. A meaningful alternative, which one can aptly call *a-leadership*, is when *some* choices of a can force the next state (the “wise choice of the leader”). It should be interesting to investigate whether the dictatorship-based and leadership-based strategic game forms lead to equivalent semantics for ATL.

DEFINITION 10. *A strategic game form is injective if o is injective, i.e. assigns different outcome states to different tuples of choices. An MGM is injective if it contains only injective game forms.*

⁵ In (Pauly, 2002) these game frames are called *dictatorial*, but we disagree with that term. Indeed, at every local step in such game one player determines the move, but these players can be different for the different moves.

EXAMPLE 8. Note that the variable client/server game model from Figure 3 is not injective, because choices $\langle reject, set0 \rangle$ and $\langle reject, set1 \rangle$ always have the same outcome. The model is not turn-based either: s is a leader at both q_0 and q_1 (he can determine the next state with $\sigma_s = reject$), but the outcome of his other choice ($\sigma_s = accept$) depends on the choice of the client. On the other hand, the game tree from Figure 1A can be seen as a turn-based MGM: player a_1 is the dictator at state q_0 , and player a_2 is the dictator at q_1 and q_2 (both players can be considered dictators at q_3, q_4, q_5 and q_6).

DEFINITION 11. (Alur et al., 1997) *An ATS is turn-based synchronous if for every $q \in Q$ there is an agent a who decides upon the next state, i.e. $\delta(q, a)$ consists entirely of singletons.*

Every ATS can be “tightened” by removing from every $Q \in \delta(q, a)$ all states which can never be realized as successors in a transition from q . Every reasonably general criterion should accept such tightening as equivalent to the original ATS.

DEFINITION 12. *An ATS $T = \langle \Pi, \text{Agt}, Q, \pi, \delta \rangle$ is tight if, for every $q \in Q$, $a \in \text{Agt}$ and $Q_a \in \delta(q, a)$, we have $Q_a \subseteq Q_q^{suc}$.*

COROLLARY 6. *For every ATS T there is a tight ATS T' which satisfies the same formulas of ATL.*

DEFINITION 13. *An ATS is lock-step synchronous if the set of successor states Q_q^{suc} of every state q can be labeled with all tuples from some Cartesian product $\prod_{a \in \text{Agt}} Q_a$ so that all choice sets from $\delta(q, a)$ are “hyperplanes” in Q_q^{suc} i.e. sets of the form $\{q_a\} \times \prod_{b \in \text{Agt} \setminus \{a\}} Q_b$, where $q_a \in Q_a$.⁶ In other words, the agents act independently and each of them can only determine its “private” component of the next state. It is worth emphasizing that lock-step synchronous systems closely correspond to the concept of interpreted systems from the literature on reasoning about knowledge (Fagin et al., 1995).*

Note that every lock-step synchronous ATS is tight.

3.2. FROM ALTERNATING TRANSITION SYSTEMS TO MGMS

First, for every ATS $T = \langle \Pi, \text{Agt}, Q, \pi, \delta \rangle$ over a set of agents $\text{Agt} = \{a_1, \dots, a_k\}$ there is an equivalent MGM $M^T = \langle Q, \gamma^T, \pi \rangle$ where, for each $q \in Q$, the strategic game form $\gamma^T(q) = \langle \text{Agt}, \{\Sigma_a^q \mid a \in \text{Agt}\}, o_q, Q \rangle$ is defined in a very simple way:

⁶ The definition in (Alur et al., 1998a) requires the *whole* state space Q to be a Cartesian product of the “local” state spaces; (Lomuscio, 1999) calls such structures “hypercube systems”. We find that requirement unnecessarily strong.

- $\Sigma_a^q = \delta(q, a)$,
- $o_q(Q_{a_1}, \dots, Q_{a_k}) = s$ where $\bigcap_{a_i \in \text{Agt}} Q_{a_i} = \{s\}$.

EXAMPLE 9. Let us apply the transformation to the alternating transition system from Example 6. The resulting MGM is shown in Figure 6. The following proposition states that it satisfies the same ATL formulas as the original system. Note that – as T and M^T include the same set of states Q – the construction preserves validity of formulas (in the model), too.

PROPOSITION 7. *For every alternating transition system T , a state q in it, and an ATL formula φ : $T, q \models \varphi$ iff $M^T, q \models \varphi$.*

The models M^T defined as above share a specific property which will be defined below. First, we need an auxiliary technical notion: a *fusion* of n -tuples $(\alpha_1, \dots, \alpha_n)$ and $(\beta_1, \dots, \beta_n)$ is any n -tuple $(\gamma_1, \dots, \gamma_n)$ where $\gamma_i \in \{\alpha_i, \beta_i\}$, $i = 1, \dots, n$. The following is easy to check.

PROPOSITION 8. *For any game form $\langle \text{Agt}, \{\Sigma_a \mid a \in \text{Agt}\}, Q, o \rangle$, where $\text{Agt} = \{a_1, \dots, a_k\}$, the following two properties of the outcome function $o : \prod_{a \in \text{Agt}} \Sigma_a \rightarrow Q$ are equivalent:*

- (i) *If $o(\sigma_{a_1}, \dots, \sigma_{a_k}) = o(\tau_{a_1}, \dots, \tau_{a_k}) = s$ then $o(\varsigma_{a_1}, \dots, \varsigma_{a_k}) = s$ for every fusion $(\varsigma_{a_1}, \dots, \varsigma_{a_k})$ of $(\sigma_{a_1}, \dots, \sigma_{a_k})$ and $(\tau_{a_1}, \dots, \tau_{a_k})$.*
- (ii) *For every $s \in Q$ there are $\Delta_a \subseteq \Sigma_a$ such that $o^{-1}(s) = \prod_{a \in \text{Agt}} \Delta_a$.*

DEFINITION 14. *A strategic game form $\langle \text{Agt}, \{\Sigma_a \mid a \in \text{Agt}\}, Q, o \rangle$ is convex if the outcome function o satisfies (any of) the two equivalent properties above. A multi-player game model $M = (Q, \gamma, \pi)$ is convex if $\gamma(q)$ is convex for every $q \in Q$.*

PROPOSITION 9. *For every ATS T , the game model M^T is convex.*

Proof: Let M^T be defined as above. If $o_q(Q_{a_1}^1, \dots, Q_{a_k}^1) = o_q(Q_{a_1}^2, \dots, Q_{a_k}^2) = s$ then $s \in Q_a^j$ for each $j = 1, 2$ and $a \in \text{Agt}$, therefore $\bigcap_{a \in \text{Agt}} Q_a^j = \{s\}$ for any fusion $(Q_{a_1}^{j_1}, \dots, Q_{a_k}^{j_k})$ of $(Q_{a_1}^1, \dots, Q_{a_k}^1)$ and $(Q_{a_1}^2, \dots, Q_{a_k}^2)$.

REMARK 10. *Pauly has pointed out that the convexity condition is known in game theory under the name of “rectangularity” and rectangular strategic game forms which are “tight” in sense that their α - and β - effectivity functions coincide are characterized in (Abdou, 1998) as the normal forms of extensive games with unique outcomes.*

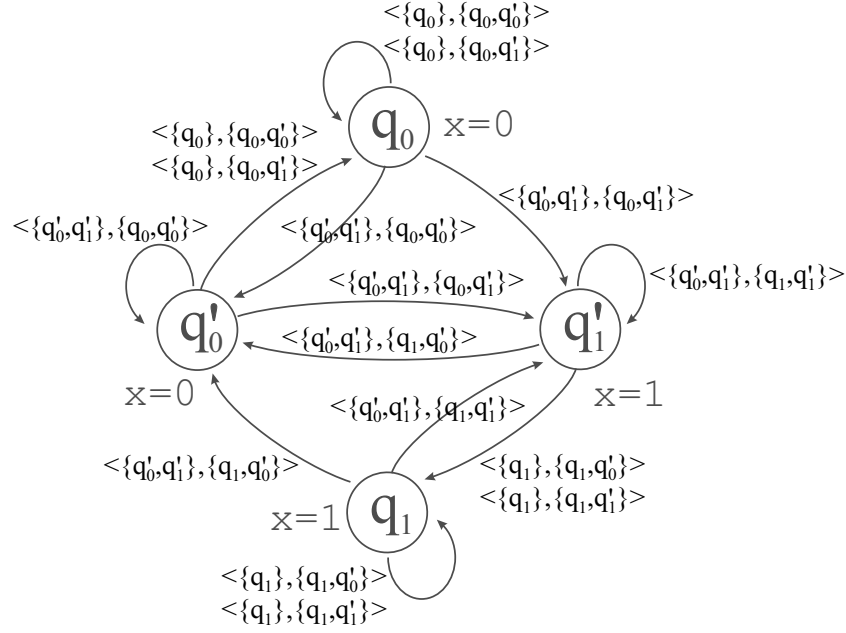


Figure 6. From ATS to a convex game structure: M^T for the system from Figure 4.

PROPOSITION 11.

1. Every turn-based game model is convex.
2. Every injective game model is convex.

Proof:

(1) Let $M = \langle Q, \gamma, \pi \rangle$ be a turn-based MGM for a set of players Agt , and let $d \in \text{Agt}$ be the dictator for $\gamma(q)$, $q \in Q$. Then for every $s \in Q$, we have $o_q^{-1}(s) = \prod_{a \in \text{Agt}} \Delta_a$ where $\Delta_d = \{\sigma_d \in \Sigma_d^q \mid o_q(\dots, \sigma_d, \dots) = s\}$, and $\Delta_a = \Sigma_a^q$ for all $a \neq d$.

(2) is trivial.

Note that the MGM from Figure 6 is convex, although it is neither injective nor turn-based, so the reverse implication does not hold.

3.3. FROM CONVEX MULTI-PLAYER GAME MODELS TO ALTERNATING TRANSITION SYSTEMS

As it turns out, convexity is a sufficient condition if we want to re-label transitions from a multi-player game model back to an alternating transition system. Let $M = \langle Q, \gamma, \pi \rangle$ be a convex MGM over a set of propositions Π , where $\text{Agt} = \{a_1, \dots, a_k\}$, and let $\gamma(q) =$

$\langle \text{Agt}, \{\Sigma_a^q \mid a \in \text{Agt}\}, Q, o_q \rangle$ for each $q \in Q$. We transform it to an ATS $T^M = \langle \Pi, \text{Agt}, Q, \pi, \delta^M \rangle$ with the transition function δ^M defined by

$$\begin{aligned} \delta^M(q, a) &= \{Q_{\sigma_a} \mid \sigma_a \in \Sigma_a^q\}, \\ Q_{\sigma_a} &= \{o_q(\sigma_a, \sigma_{\text{Agt} \setminus \{a\}}) \mid \sigma_b \in \Sigma_b^q, b \neq a\}. \end{aligned}$$

Thus, Q_{σ_a} is the set of states to which a transition may be effected from q while agent a has chosen to execute σ_a . Moreover, $\delta^M(q, a)$ simply collects all such sets. For purely technical reasons we will regard these $\delta^M(q, a)$ as *indexed families* i.e. even if some Q_{σ_1} and Q_{σ_2} are set-theoretically equal, they will be considered different as long as $\sigma_1 \neq \sigma_2$. By convexity of $\gamma(q)$ it is easy to verify that $\bigcap_{a \in \text{Agt}} Q_{\sigma_a} = \{o_q(\sigma_{a_1}, \dots, \sigma_{a_k})\}$ for every tuple $(Q_{\sigma_{a_1}}, \dots, Q_{\sigma_{a_k}}) \in \delta^M(q, a_1) \times \dots \times \delta^M(q, a_k)$. Furthermore, the following propositions hold.

PROPOSITION 12. *For every convex MGM M the ATS T^M is tight.*

PROPOSITION 13. *For every convex MGM M , a state q in it, and an ATL formula φ , $M, q \models \varphi$ iff $T^M, q \models \varphi$.*

Note that the above construction transforms the multi-player game model from Figure 6 exactly back to the ATS from Figure 4. More generally, the constructions converting tight ATSS into convex MGMs and vice versa are mutually inverse, thus establishing a duality between these two types of structures:

PROPOSITION 14.

1. *Every tight ATS T is isomorphic to T^{T^T} .*
2. *Every convex MGM M is isomorphic to M^{T^M} .*

Proof:

1. It suffices to see that $\delta^{M^T}(q, a) = \delta(q, a)$ for every $q \in Q$ and $a \in \text{Agt}$ which is straightforward from the tightness of T .

2. Let $M = \langle Q, \gamma, \pi \rangle$ be a convex MGM and $\gamma(q) = \langle \text{Agt}, \{\Sigma_a^q \mid a \in \text{Agt}\}, Q, o_q \rangle$ for $q \in Q$. For every $\sigma_a \in \Sigma_a^q$ we identify σ_a with Q_{σ_a} defined as above. We have to show that the outcome functions o_q in M and o_q in M^{T^M} agree under that identification. Indeed, $o_q(Q_{\sigma_{a_1}}, \dots, Q_{\sigma_{a_k}}) = s$ iff $\bigcap_{a \in \text{Agt}} Q_{\sigma_a} = \{s\}$ iff $o_q(\sigma_{a_1}, \dots, \sigma_{a_k}) = s$.

The following proposition shows the relationship between structural properties of MGMs and ATSS:

PROPOSITION 15.

1. For every ATS T the game model M^T is injective iff T is lock-step synchronous.
2. For every convex MGM M , the ATS T^M is lock-step synchronous iff M is injective.
3. For every turn-based synchronous ATS T the game model M^T is turn-based. Conversely, if M^T is turn-based for some tight ATS T then T is turn-based synchronous.
4. For every convex MGM M the ATS T^M is turn-based synchronous iff M is turn-based.

Proof:

(1) Let T be lock-step synchronous and $o_q(Q_{a_1}, \dots, Q_{a_k}) = \langle s_{a_1}, \dots, s_{a_k} \rangle$ for some $Q_{a_i} \in \delta(q, a_i)$, $i = 1, \dots, k$. Then $Q_{a_i} = \{s_{a_i}\} \times \prod_{a \in \text{Agt} \setminus \{a_i\}} Q_a$ where $Q_q^{suc} = \prod_{a \in \text{Agt}} Q_a$, whence the injectivity of M^T . Conversely, if M^T is injective then every state $s \in Q_q^{suc}$ can be labeled with the unique tuple $\langle Q_{a_1}, \dots, Q_{a_k} \rangle$ such that $o_q(Q_{a_1}, \dots, Q_{a_k}) = s$, i.e. Q_q^{suc} is represented by $\prod_{a \in \text{Agt}} \delta(q, a)$, and every $Q_{a_i} \in \delta(q, a_i)$ can be identified with $\{Q_{a_i}\} \times \prod_{a \in \text{Agt} \setminus \{a_i\}} \delta(q, a)$.

(2) If M is injective then Q_q^{suc} can be labeled by $\prod_{a \in \text{Agt}} \Sigma_a^q$ where every $Q_{\sigma_{a_i}} \in \delta(q, a_i)$ is identified with $\{\sigma_{a_i}\} \times \prod_{a \in \text{Agt} \setminus \{a_i\}} \delta(q, a)$. Conversely, if T^M is lock-step synchronous then every two different Q_{σ_a} and Q'_{σ_a} from $\delta(q, a)$ must be disjoint, whence the injectivity of M .

(3) and (4): the proofs are straightforward.

3.4. EQUIVALENCE BETWEEN THE SEMANTICS FOR ATL BASED ON ATS AND MGM

So far we have shown how to transform alternating transition systems to convex multi-player game models, and vice versa. Unfortunately, not every MGM is convex. However, for every MGM we can construct a convex multi-player game model that satisfies the same formulas of ATL. This can be done by creating distinct copies of the original states for different incoming transitions, and thus “storing” the knowledge of the previous state and the most recent choices from the agents in the new states. Since the actual choices are present in the label of the resulting state, the new transition function is obviously injective. It is also easy to observe that the below construction preserves not only satisfiability, but also validity of formulas (in the model).

PROPOSITION 16. *For every MGM $M = \langle Q, \gamma, \pi \rangle$ there is an injective (and hence convex) MGM $M' = \langle Q', \gamma', \pi' \rangle$ which satisfies the same formulas of ATL.*

Proof: For every $q = \langle \text{Agt}, \{\Sigma_a^q \mid a \in \text{Agt}\}, Q, o_q \rangle$ we define $Q_q = \{q\} \times \prod_{a \in \text{Agt}} \Sigma_a^q$ and let $Q' = Q \cup \bigcup_{q \in Q} Q_q$. Now we define γ' as follows:

- for $q \in Q$, we define $\gamma'(q) = \langle \text{Agt}, \{\Sigma_a^q \mid a \in \text{Agt}\}, O^q, Q' \rangle$, and $O^q(\sigma_{a_1}, \dots, \sigma_{a_k}) = \langle q, \sigma_{a_1}, \dots, \sigma_{a_k} \rangle$;
- for $\sigma = \langle q, \sigma_{a_1}, \dots, \sigma_{a_k} \rangle \in Q_q$, and $s = o_q(\sigma_{a_1}, \dots, \sigma_{a_k})$, we define $\gamma'(\sigma) = \gamma'(s)$;
- finally, $\pi'(q) = \pi(q)$ for $q \in Q$, and $\pi'(\langle q, \sigma_{a_1}, \dots, \sigma_{a_k} \rangle) = \pi(o_q(\sigma_{a_1}, \dots, \sigma_{a_k}))$ for $\langle q, \sigma_{a_1}, \dots, \sigma_{a_k} \rangle \in Q_q$.

The model M' is injective and it can be proved by a straightforward induction that for every ATL formula φ :

- $M', q \models \varphi$ iff $M, q \models \varphi$ for $q \in Q$, and
- $M', \langle \sigma_{a_1}, \dots, \sigma_{a_k} \rangle \models \varphi$ iff $M, o_q(\sigma_{a_1}, \dots, \sigma_{a_k}) \models \varphi$ for $\langle \sigma_{a_1}, \dots, \sigma_{a_k} \rangle \in Q_q$.

Thus, the restriction of the semantics of ATL to the class of injective (and hence to convex, as well) MGMs does not introduce new validities. Since every ATS can be reduced to an equivalent tight one, we obtain the following.

COROLLARY 17. *For every ATL formula φ the following are equivalent:*

1. φ is valid in all (tight) alternating transition systems.
2. φ is valid in all (injective) multi-player game models.

We note that the above construction preserves validity and satisfiability of ATL* formulas, too.

EXAMPLE 10. We can apply the construction to the controller from Example 4, and obtain a convex MGM equivalent to the original one in the context of ATL. The result is displayed in Figure 7. The labels for the transitions can be easily deduced from their target states. Rewriting the game model into an isomorphic ATS, according to the construction from Section 3.3 (see Figure 8), completes the transformation from an arbitrary multi-player game model to an alternating transition system for which the same ATL formulas hold.




Figure 7. Construction of a convex multi-player game model equivalent to the MGM from Figure 3.

δ	$ $	$q_0, q_{0r0}, q_{0r1}, q_{0a0}, q_{1a0}$	$ $	$q_1, q_{1r0}, q_{1r1}, q_{1a1}, q_{0a1}$	$ $
s	$ $	$\{\{q_{0r0}, q_{0r1}\}, \{q_{0a0}, q_{0a1}\}\}$	$ $	$\{\{q_{1r0}, q_{1r1}\}, \{q_{1a0}, q_{1a1}\}\}$	$ $
c	$ $	$\{\{q_{0r0}, q_{0a0}\}, \{q_{0r1}, q_{0a1}\}\}$	$ $	$\{\{q_{1r0}, q_{1a0}\}, \{q_{1r1}, q_{1a1}\}\}$	$ $

Figure 8. ATS-style transition function for the convex game model from Figure 7.

3.5. ATS OR MGM?

Rajeev Alur stated⁷ that the authors of ATL switched from alternating transition systems to concurrent game structures mostly to improve understandability of the logic and clarity of the presentation. Indeed, identifying actions with their outcomes may make things somewhat artificial and unnecessarily complicated. In particular, we find the convexity condition which ATSs impose too strong and unjustified in many situations. For instance, consider the following variation of the ‘Chicken’ game: two cars running against each other on a country road and each of the drivers, seeing the other car, can take any of the actions: “drive straight”, “swerve to the left” and “swerve to the right”. Each of the combined actions for the two drivers: \langle drive straight, swerve to

⁷ Private communication.

the left) and $\langle \textit{swerve to the right}, \textit{drive straight} \rangle$ leads to a non-collision outcome, while each of their fusions $\langle \textit{drive straight}, \textit{drive straight} \rangle$ and $\langle \textit{swerve to the left}, \textit{swerve to the right} \rangle$ leads to a collision. Likewise, in the “Coordinated Attack” scenario (Fagin et al., 1995) any non-coordinated one-sided attack leads to defeat, while the coordinated attack of both armies, which is a fusion of these, leads to a victory. Thus, the definition of outcome function in coalition games is more general and flexible in our opinion.

Let us consider the system from Example 4 again. The multi-player game model (or concurrent game structure) from Figure 3 looks natural and intuitive. Unfortunately, there is no “isomorphic” ATS that fits the system description. In consequence, an ATS modeling the same situation must be larger (Jamroga, 2003). The above examples show that correct alternating transition systems are more difficult to come up with directly than multi-player game models, and usually they are more complex, too. This should be especially evident when we consider open systems. Suppose we need to add another client process to the ATS from Example 6. It would be hard to extend the existing transition function in a straightforward way so that it still satisfies the formal requirements (all the intersections of choices are singletons). Designing a completely new ATS is probably an easier solution.

Another interesting issue is extendibility of the formalisms. Game models incorporate explicit labels for agents’ choices – therefore the labels can be used, for instance, to restrict the set of valid strategies under uncertainty (Jamroga and van der Hoek, 2003).

3.6. COALITION EFFECTIVITY MODELS AS EQUIVALENT ALTERNATIVE SEMANTICS FOR ATL

Effectivity functions and coalition effectivity models were introduced in Section 2.5, including a characterization of these effectivity functions which describe abilities of agents and their coalitions in actual strategic game forms (playable effectivity functions, Theorem 5). We are going to extend the result to correspondence between multi-player game models and standard coalition effectivity models (i.e. the coalition effectivity models that contain only playable effectivity functions).

Every MGM $M = \langle Q, \gamma, \pi \rangle$ for the set of players Agt corresponds to a CEM $\mathcal{E}^M = \langle \text{Agt}, Q, E^M, \pi \rangle$, where for every $q \in Q$, $X \subseteq Q$ and $A \subseteq \text{Agt}$, we have

$$X \in E_q^M(A) \text{ iff } \exists \sigma_A \forall \sigma_{\text{Agt} \setminus A} \exists s \in X \ o(\sigma_A, \sigma_{\text{Agt} \setminus A}) = s.$$

The choices refer to the strategic game form $\gamma(q)$. Conversely, by Theorem 5, for every standard coalition effectivity model \mathcal{E} there is a

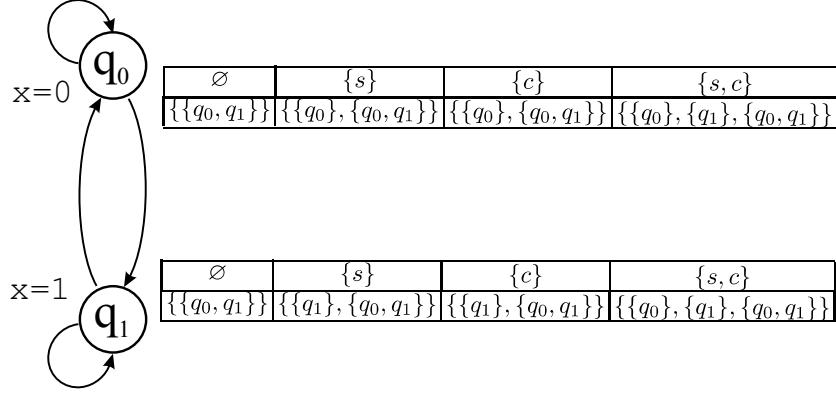


Figure 9. Coalition effectivity model for the variable client/server system

multi-player game model M such that \mathcal{E} is isomorphic to \mathcal{E}^M . Again, by a straightforward induction on formulas, we obtain:

PROPOSITION 18. *For every MGM M , a state q in it, and an ATL formula φ , we have $M, q \models \varphi$ iff $\mathcal{E}^M, q \models \varphi$.*

EXAMPLE 11. Let M be the multi-player game model from Example 4 (variable client/server system). Coalition effectivity model \mathcal{E}^M is presented in Figure 9.

By Proposition 18 and Corollary 17, we eventually obtain:

THEOREM 19. *For every ATL formula φ the following are equivalent:*

1. φ is valid in all (tight) alternating transition systems,
2. φ is valid in all (injective) multi-player game models,
3. φ is valid in all standard coalition effectivity models.

Thus, the semantics of ATL based on alternating transition systems, multi-player game models, and standard coalition effectivity models are equivalent. We note that, while the former two semantics are more concrete and natural, they are mathematically less elegant and suitable for formal reasoning about ATL, while the semantics based on coalition effectivity models is essentially a monotone neighborhood semantics for multi-modal logics. The combination of these semantics was used in (Goranko and van Drimmelen, 2003) to establish a complete axiomatization of ATL.

4. ATEL: adding knowledge to strategies and time

Alternating-time Temporal Epistemic Logic ATEL (van der Hoek and Wooldridge, 2002; van der Hoek and Wooldridge, 2003a) enriches the picture with epistemic component. ATEL adds to ATL operators for representing agents' knowledge: $K_a\varphi$ reads as "agent a knows that φ ". Additional operators $E_A\varphi$, $C_A\varphi$, and $D_A\varphi$ refer to "everybody knows", common knowledge, and distributed knowledge among the agents from A . Thus, $E_A\varphi$ means that every agent in A knows that φ holds, while $C_A\varphi$ means not only that the agents from A know that φ , but they also know that they know that, and know that they know that they know it, etc. The distributed knowledge modality $D_A\varphi$ denotes a situation in which, if the agents could combine their individual knowledge together, they would be able to infer that φ is true.

4.1. AETS AND SEMANTICS OF EPISTEMIC FORMULAS

Models for ATEL are called *alternating epistemic transition systems* (AETS). They extend alternating transition systems with epistemic accessibility relations $\sim_1, \dots, \sim_k \subseteq Q \times Q$ for modeling agents' uncertainty:

$$\mathcal{T} = \langle \text{Agt}, Q, \Pi, \pi, \sim_{a_1}, \dots, \sim_{a_k}, \delta \rangle.$$

These are assumed to be equivalence relations. Agent a 's epistemic relation is meant to encode a 's inability to distinguish between the (global) system states: $q \sim_a q'$ means that, while the system is in state q , agent a cannot really determine whether it is in q or q' . Then:

$$\mathcal{T}, q \models K_a\varphi \text{ iff for all } q' \text{ such that } q \sim_a q' \text{ we have } \mathcal{T}, q' \models \varphi$$

REMARK 20. *Since the epistemic relations are required to be equivalences, the epistemic layer of ATEL refers indeed to agents' knowledge rather than beliefs in general. We suggest that this requirement can be relieved to allow ATEL for other kinds of beliefs as well. In particular, the interpretation of ATEL into ATL we propose in Section 4.4 does not assume any specific properties of the accessibility relations.*

Relations \sim_A^E , \sim_A^C and \sim_A^D , used to model group epistemics, are derived from the individual accessibility relations of agents from A . First, \sim_A^E is the union of the relations, i.e. $q \sim_A^E q'$ iff $q \sim_a q'$ for some $a \in A$. In other words, if everybody knows φ , then no agent may be unsure about the truth of it, and hence φ should be true in all the states that cannot be distinguished from the current state by even one member of the group. Next, \sim_A^C is defined as the transitive closure of




Figure 10. An AETS for the modified controller/client problem. The dotted lines display the epistemic accessibility relations for s and c .

\sim_A^E . Finally, \sim_A^D is the intersection of all the \sim_a , $a \in A$: if any agent from A can distinguish q from q' , then the whole group can distinguish the states in the sense of distributed knowledge. The semantics of group knowledge can be defined as below (for $\mathcal{K} = C, E, D$):

$$\mathcal{T}, q \models \mathcal{K}_A \varphi \text{ iff for all } q' \text{ such that } q \sim_A^{\mathcal{K}} q' \text{ we have } \mathcal{T}, q' \models \varphi$$

The time complexity of model checking for ATEL is still polynomial. (van der Hoek and Wooldridge, 2003a).

EXAMPLE 12. Let us consider another variation of the variable controller example: the client can try to add 1 or 2 (modulo 3) to the value of x ; the server can still accept or reject the request (Figure 10). The dotted lines show that c cannot distinguish being in state q_0 from being in q_1 , while s isn't able to discriminate q_0 from q_2 . Some formulas that are valid for this AETS are shown below:

1. $x=1 \rightarrow K_s x=1$,
2. $x=2 \rightarrow E_{s,c} \neg x=1 \wedge \neg C_{s,c} \neg x=1$,
3. $x=0 \rightarrow \langle\langle s \rangle\rangle X x=0 \wedge \neg K_s \langle\langle s \rangle\rangle X x=0$,
4. $x=2 \rightarrow \langle\langle s, c \rangle\rangle X (x=0 \wedge \neg E_{s,c} x=0)$.

4.2. EXTENDING MULTI-PLAYER GAME MODELS AND COALITION EFFECTIVITY MODELS TO INCLUDE KNOWLEDGE

Multi-player game models and coalition effectivity models can be augmented with epistemic accessibility relations in a similar way, giving

way to multi-player epistemic game models $\mathcal{M} = \langle Q, \gamma, \pi, \sim_{a_1}, \dots, \sim_{a_k} \rangle$ and epistemic coalition effectivity models $\mathcal{E} = \langle \text{Agt}, Q, E, \pi, \sim_{a_1}, \dots, \sim_{a_k} \rangle$ for a set of agents $\text{Agt} = \{a_1, \dots, a_k\}$ over a set of propositions Π . Semantic rules for epistemic formulas remain the same as in Section 4.1 for both kinds of structures. The equivalence results from Section 3 can be extended to ATEL and its models. In particular, Theorem 19 yields an immediate corollary for ATEL semantics:

COROLLARY 21. *For every ATEL formula φ the following are equivalent:*

1. φ is valid in all (tight) alternating epistemic transition systems,
2. φ is valid in all (injective) multi-player epistemic game models,
3. φ is valid in all standard epistemic coalition effectivity models.

We will use multi-player epistemic game models throughout the rest of this chapter for the convenience of presentation they offer.

4.3. PROBLEMS WITH ATEL

One of the main challenges in ATEL is the question how, given an explicit way to represent agents' knowledge, this should interfere with the agents' available strategies. What does it mean that an agent has a strategy to enforce φ , if it involves making different choices in states that are epistemically indistinguishable for the agent, for instance? Moreover, agents are assumed some epistemic capabilities when making decisions, and other for epistemic properties like $K_a\varphi$. The interpretation of knowledge operators refers to the agents' capability to distinguish one *state* from another; the semantics of $\langle\langle A \rangle\rangle$ allows the agents to base their decisions upon *sequences* of states. These relations between complete vs. incomplete information on one hand, and perfect vs. imperfect recall on the other, has been studied in (Jamroga and van der Hoek, 2003). It was also argued that, when reasoning about what an agent can *enforce*, it seems more appropriate to require the agent to know his winning strategy rather than to know only that such a strategy exists.⁸ Two variations of ATEL were proposed as solutions: Alternating-time Temporal Observational Logic (ATOL) for agents with bounded memory and syntax restricted in a way similar to

⁸ This problem is closely related to the distinction between knowledge *de re* and knowledge *de dicto*. The issue is well known in the philosophy of language (Quine, 1956), as well as research on the interaction between knowledge and action (Moore, 1985; Morgenstern, 1991; Wooldridge, 2000).

CTL, and full Alternating-time Temporal Epistemic Logic with Recall (ATEL-R*) where agents were able to memorize the whole game.

The issue of a philosophically consistent semantics for Alternating-time Temporal Logic with epistemic component is still under debate, and it is rather beyond the scope of this paper. We believe that analogous results to those presented here about ATEL can be obtained for logics like ATOL and ATEL-R* and their models.

4.4. INTERPRETATIONS OF ATEL INTO ATL

ATL is trivially embedded into ATEL since all ATL formulas are also ATEL formulas. Moreover, every multi-player game model can be extended to a multi-player epistemic game model by defining all epistemic accessibility relations to be the equality, i.e. all agents have no uncertainty about the current state of the system – thus embedding the semantics of ATL in the one for ATEL, and rendering the former a reduct of the latter.

Interpretation the other way is more involved. We will first construct a satisfiability preserving interpretation of the fragment of ATEL without distributed knowledge (we will call it ATEL_{CE}), and then we will show how it can be extended to the whole ATEL, though at the expense of some blow-up of the models. The interpretation we propose has been inspired by (Schild, 2000). We should also mention (van Otterloo et al., 2003), as it deals with virtually the same issue. Related work is discussed in more detail at the end of the section.

4.4.1. *Idea of the interpretation*

ATEL consists of two orthogonal layers. The first one, inherited from ATL, refers to what agents can achieve in temporal perspective, and is underpinned by the structure defined via transition function o . The other layer is the epistemic component, reflected by epistemic accessibility relations. Our idea of the translation is to leave the original temporal structure intact, while extending it with additional transitions to “simulate” epistemic accessibility links. The “simulation” – like the one in (van Otterloo et al., 2003) – is achieved through adding new “epistemic” agents, who can enforce transitions to epistemically accessible states. Unlike in that paper, though, the “moves” of epistemic agents are orthogonal to the original temporal transitions (“action” transitions): they lead to special “epistemic” copies of the original states rather than to the “action” states themselves, and no new states are introduced into the course of action. The “action” and “epistemic” states form separate strata in the resulting model, and are labeled accordingly to distinguish transitions that implement different modalities.




Figure 11. New model: “action” vs. “epistemic” states, and “action” vs. “epistemic” transitions. Note that the game frames for “epistemic” states are *exact* copies of their “action” originals: the “action” transitions from the epistemic layer lead back to the “action” states.

The interpretation consists of two independent parts: a transformation of models and a translation of formulas. First, we propose a construction that transforms every multi-player epistemic game model \mathcal{M} for a set of agents $\{a_1, \dots, a_k\}$, into a (pure) multi-player game model \mathcal{M}^{ATL} over a set of agents $\{a_1, \dots, a_k, e_1, \dots, e_k\}$. Agents a_1, \dots, a_k are the original agents from \mathcal{M} (we will call them “real agents”). Agents e_1, \dots, e_k are “epistemic doubles” of the real agents: the role of e_i is to “point out” the states that were epistemically indistinguishable from the current state for agent a_i in \mathcal{M} . Intuitively, $K_{a_i}\varphi$ could be then replaced with a formula like $\neg\langle\langle e_i \rangle\rangle X\neg\varphi$ that rephrases the semantic definition of K_a operator from Section 4.1. As \mathcal{M}^{ATL} inherits the temporal structure from \mathcal{M} , temporal formulas might be left intact. However, it is not as simple as that.

Note that agents make their choices simultaneously in multi-player game models, and the resulting transition is a result of all these choices. In consequence, it is not possible that an epistemic agent e_i can enforce an “epistemic” transition to state q , and at the same time a group of real agents A is capable of executing an “action” transition to q' . Thus, in order to distinguish transitions referring to different modalities, we introduce additional states in model \mathcal{M}^{ATL} . States $q_1^{e_i}, \dots, q_n^{e_i}$ are exact copies of the original states q_1, \dots, q_n from Q except for one thing: they satisfy a new proposition e_i , added to enable identifying moves

of epistemic agent e_i . Original states q_1, \dots, q_n are still in \mathcal{M}^{ATL} to represent targets of “action” moves of the real agents a_1, \dots, a_k . We will use a new proposition act to label these states. The type of a transition can be recognized by the label of its target state (cf. Figure 11).

Now, we must only arrange the interplay between agents’ choices, so that the results can be interpreted in a direct way. To achieve this, every epistemic agent can choose to be “passive” and let the others decide upon the next move, or may select one of the states indistinguishable from q for an agent a_i (to be more precise, the epistemic agents do select the epistemic copies of states from Q^{e_i} rather than the original action states from Q). The resulting transition leads to the state selected by the *first* non-passive epistemic agent. If all the epistemic agents decided to be passive, the “action” transition chosen by the real agents follows.

For such a construction of \mathcal{M}^{ATL} , we can finally show how to translate formulas from ATEL to ATL:

- $K_{a_i}\varphi$ can be rephrased as $\neg\langle\{e_1, \dots, e_i\}\rangle X(e_i \wedge \neg\varphi)$: the epistemic moves to agent e_i ’s epistemic states do not lead to a state where φ fails. Note that player e_i can select a state of his if, and only if, players e_1, \dots, e_{i-1} are passive (hence their presence in the cooperation modality). Note also that $K_{a_i}\varphi$ can be as well translated as $\neg\langle\{e_1, \dots, e_k\}\rangle X(e_i \wedge \neg\varphi)$ or $\neg\langle\{a_1, \dots, a_k, e_1, \dots, e_k\}\rangle X(e_i \wedge \neg\varphi)$: when e_i decides to be active, choices from a_1, \dots, a_k and e_{i+1}, \dots, e_k are irrelevant.
- $\langle\langle A \rangle\rangle X\varphi$ becomes $\langle\langle A \cup \{e_1, \dots, e_k\}\rangle\rangle X(\text{act} \wedge \varphi)$ in a similar way.
- To translate other temporal formulas, we must require that the relevant part of a path runs only through “action” states (labeled with act proposition). Thus, $\langle\langle A \rangle\rangle G\varphi$ can be rephrased as $\varphi \wedge \langle\langle A \cup \text{Agt}^e \rangle\rangle X \langle\langle A \cup \text{Agt}^e \rangle\rangle G(\text{act} \wedge \varphi)$. Note that a simpler translation with $\langle\langle A \cup \text{Agt}^e \rangle\rangle G(\text{act} \wedge \varphi)$ is incorrect: the initial state of a path does not have to be an action state, since $\langle\langle A \rangle\rangle G\varphi$ can be embedded in an epistemic formula. A similar method applies to the translation of $\langle\langle A \rangle\rangle \varphi \mathcal{U} \psi$.
- Translation of common knowledge refers to the definition of relation \sim_A^C as the transitive closure of relations \sim_{a_i} : $C_A\varphi$ means that all the (finite) sequences of appropriate epistemic transitions must end up in a state where φ is true.

The only operator that does not seem to lend itself to a translation according to the above scheme is the distributed knowledge operator D , for which we seem to need more “auxiliary” agents. Thus, we will begin with presenting details of our interpretation for ATEL_{CE} – a

reduced version of ATEL that includes only common knowledge and “everybody knows” operators for group epistemics. Section 4.4.3 shows how to modify the translation to include distributed knowledge as well.

We note that an analogous interpretation into ATL can be proposed for the propositional version of BDI logic based on CTL.

4.4.2. Interpreting models and formulas of ATEL_{CE} into ATL

Given a multi-player epistemic game model $\mathcal{M} = \langle Q, \gamma, \pi, \sim_{a_1}, \dots, \sim_{a_k} \rangle$ for a set of agents $\text{Agt} = \{a_1, \dots, a_k\}$ over a set of propositions Π , we construct a new game model $\mathcal{M}^{\text{ATL}} = \langle Q', \gamma', \pi' \rangle$ over a set of agents $\text{Agt}' = \text{Agt} \cup \text{Agt}^e$, where:

- $\text{Agt}^e = \{e_1, \dots, e_k\}$ is the set of epistemic agents;
- $Q' = Q \cup Q^{e_1} \cup \dots \cup Q^{e_k}$, where $Q^{e_i} = \{q^{e_i} \mid q \in Q\}$. We assume that $Q, Q^{e_1}, \dots, Q^{e_k}$ are pairwise disjoint. Further we will be using the more general notation $S^{e_i} = \{q^{e_i} \mid q \in S\}$ for any $S \subseteq Q$.
- $\Pi' = \Pi \cup \{\text{act}, e_1, \dots, e_k\}$, and $\pi'(p) = \pi(p) \cup \bigcup_{i=1, \dots, k} \pi(p)^{e_i}$ for every proposition $p \in \Pi$. Moreover, $\pi'(\text{act}) = Q$ and $\pi'(e_i) = Q^{e_i}$.

For every state q in \mathcal{M} , we translate the game frame $\gamma(q) = \langle \text{Agt}, \{\Sigma_a^q \mid a \in \text{Agt}\}, Q, o \rangle$ to $\gamma'(q) = \langle \text{Agt}', \{\Sigma_a^{q'} \mid a \in \text{Agt}'\}, Q', o' \rangle$:

- $\Sigma_a^{q'} = \Sigma_a^q$ for $a \in \text{Agt}$: choices of the “real” agents do not change;
- $\Sigma_{e_i}^{q'} = \{\text{pass}\} \cup \text{img}(q, \sim_{a_i})^{e_i}$ for $i = 1, \dots, k$, where $\text{img}(q, R) = \{q' \mid qRq'\}$ is the image of q with respect to relation R .
- the new transition function is defined as follows:

$$o'_q(\sigma_{a_1}, \dots, \sigma_{a_k}, \sigma_{e_1}, \dots, \sigma_{e_k}) = \begin{cases} o_q(\sigma_{a_1}, \dots, \sigma_{a_k}) & \text{if } \sigma_{e_1} = \dots = \sigma_{e_k} = \text{pass} \\ \sigma_{e_i} & \text{if } e_i \text{ is the first active} \\ & \text{epistemic agent.} \end{cases}$$

The game frames for the new states are exactly the same: $\gamma'(q^{e_i}) = \gamma'(q)$ for all $i = 1, \dots, k, q \in Q$.

EXAMPLE 13. A part of the resulting structure for the epistemic game model from Figure 10 is shown in Figure 12. All the new states, plus the transitions going out of q_2 are presented. The wildcard “*” stands for any action of the respective agent. For instance, $\langle \text{reject}, *, \text{pass}, \text{pass} \rangle$ represents $\langle \text{reject}, \text{set0}, \text{pass}, \text{pass} \rangle$ and $\langle \text{reject}, \text{set1}, \text{pass}, \text{pass} \rangle$.




Figure 12. Construction for the multi-player epistemic game model from Figure 10.

Now, we define a translation of formulas from ATEL_{CE} to ATL corresponding to the above described interpretation of ATEL models into ATL models:

$$\begin{aligned}
tr(p) &= p, & \text{for } p \in \Pi \\
tr(\neg\varphi) &= \neg tr(\varphi) \\
tr(\varphi \vee \psi) &= tr(\varphi) \vee tr(\psi) \\
tr(\langle\langle A \rangle\rangle X\varphi) &= \langle\langle A \cup \text{Agt}^e \rangle\rangle X(\text{act} \wedge tr(\varphi)) \\
tr(\langle\langle A \rangle\rangle G\varphi) &= tr(\varphi) \wedge \langle\langle A \cup \text{Agt}^e \rangle\rangle X \langle\langle A \cup \text{Agt}^e \rangle\rangle G(\text{act} \wedge tr(\varphi)) \\
tr(\langle\langle A \rangle\rangle \varphi \mathcal{U} \psi) &= tr(\psi) \vee (tr(\varphi) \wedge \langle\langle A \cup \text{Agt}^e \rangle\rangle X \langle\langle A \cup \text{Agt}^e \rangle\rangle \\
&\quad (\text{act} \wedge tr(\varphi)) \mathcal{U} (\text{act} \wedge tr(\psi))) \\
tr(K_{a_i} \varphi) &= \neg \langle\langle \{e_1, \dots, e_i\} \rangle\rangle X (e_i \wedge \neg tr(\varphi)) \\
tr(E_A \varphi) &= \neg \langle\langle \text{Agt}^e \rangle\rangle X (\bigvee_{a_i \in A} e_i \wedge \neg tr(\varphi)) \\
tr(C_A \varphi) &= \neg \langle\langle \text{Agt}^e \rangle\rangle X \langle\langle \text{Agt}^e \rangle\rangle (\bigvee_{a_i \in A} e_i) \mathcal{U} (\bigvee_{a_i \in A} e_i \wedge \neg tr(\varphi))
\end{aligned}$$

LEMMA 22. For every ATEL_{CE} formula φ , model \mathcal{M} , and “action” state $q \in Q$, we have $\mathcal{M}^{\text{ATL}}, q \models tr(\varphi)$ iff $\mathcal{M}^{\text{ATL}}, q^{e_i} \models tr(\varphi)$ for every $i = 1, \dots, k$.

Proof sketch (structural induction on φ): It suffices to note that $tr(\varphi)$ can only contain propositions $\text{act}, e_1, \dots, e_k$ in the scope of $\langle\langle A \rangle\rangle X$ for some $A \subseteq \text{Agt}^e$. Besides, the propositions from φ are true in q iff they are true in q^{e_1}, \dots, q^{e_k} and the game frames for $q, q^{e_1}, \dots, q^{e_k}$ are the same.

LEMMA 23. For every ATEL_{CE} formula φ , model \mathcal{M} , and a state $q \in Q$, we have $\mathcal{M}, q \models \varphi$ iff $\mathcal{M}^{\text{ATL}}, q \models \text{tr}(\varphi)$.

Proof: The proof follows by structural induction on φ . We will show that the construction preserves the truth value of φ for two cases: $\varphi \equiv \langle\langle A \rangle\rangle X\psi$ and $\varphi \equiv C_A\psi$. The cases of $\langle\langle A \rangle\rangle G\psi$ and $\langle\langle A \rangle\rangle \psi \mathcal{U} \vartheta$ can be reduced to the case for $\langle\langle A \rangle\rangle X\psi$ using the fact that these operators are fixpoints (resp. greatest and least) of certain operators defined in terms of $\langle\langle A \rangle\rangle X\psi$ (see Section 2.5). For lack of space we omit the details. An interested reader can tackle the other cases in an analogous way.

case $\varphi \equiv \langle\langle A \rangle\rangle X\psi$, $\text{ATEL}_{\text{CE}} \Rightarrow \text{ATL}$. Let $\mathcal{M}, q \models \langle\langle A \rangle\rangle X\psi$, then there is σ_A such that for every $\sigma_{\text{Agt} \setminus A}$ we have $o_q(\sigma_A, \sigma_{\text{Agt} \setminus A}) \models \psi$. By induction hypothesis, $\mathcal{M}^{\text{ATL}}, o_q(\sigma_A, \sigma_{\text{Agt} \setminus A}) \models \text{tr}(\psi)$; also, $\mathcal{M}^{\text{ATL}}, o_q(\sigma_A, \sigma_{\text{Agt} \setminus A}) \models \text{act}$. Thus, $\mathcal{M}^{\text{ATL}}, o'_q(\sigma_A, \sigma_{\text{Agt} \setminus A}, \text{pass}_{\text{Agt}^e}) = o_q(\sigma_A, \sigma_{\text{Agt} \setminus A}) \models \text{act} \wedge \text{tr}(\psi)$, where pass_C denotes the strategy where every agent from $C \subseteq \text{Agt}^e$ decides to be passive. In consequence, $\mathcal{M}^{\text{ATL}}, q \models \langle\langle A \cup \text{Agt}^e \rangle\rangle X \text{tr}(\psi)$.

case $\varphi \equiv \langle\langle A \rangle\rangle X\psi$, $\text{ATL} \Rightarrow \text{ATEL}_{\text{CE}}$. $\mathcal{M}^{\text{ATL}}, q \models \langle\langle A \cup \text{Agt}^e \rangle\rangle X (\text{act} \wedge \text{tr}(\psi))$, so there is $\sigma_{A \cup \text{Agt}^e}$ such that for every $\sigma_{\text{Agt}' \setminus (A \cup \text{Agt}^e)} = \sigma_{\text{Agt} \setminus A}$ we have $\mathcal{M}^{\text{ATL}}, o'_q(\sigma_{A \cup \text{Agt}^e}, \sigma_{\text{Agt} \setminus A}) \models \text{act} \wedge \text{tr}(\psi)$. Note that $\mathcal{M}^{\text{ATL}}, o'_q(\sigma_{A \cup \text{Agt}^e}, \sigma_{\text{Agt} \setminus A}) \models \text{act}$ only when $\sigma_{A \cup \text{Agt}^e} = \langle\sigma_A, \text{pass}_{\text{Agt}^e}\rangle$, else the transition would lead to an epistemic state. Thus, $o'_q(\sigma_{A \cup \text{Agt}^e}, \sigma_{\text{Agt} \setminus A}) = o_q(\sigma_A, \sigma_{\text{Agt} \setminus A})$, and hence $\mathcal{M}^{\text{ATL}}, o_q(\sigma_A, \sigma_{\text{Agt} \setminus A}) \models \text{tr}(\psi)$. By the induction hypothesis, $\mathcal{M}, o_q(\sigma_A, \sigma_{\text{Agt} \setminus A}) \models \psi$ and finally $\mathcal{M}, q \models \langle\langle A \rangle\rangle X\psi$.

case $\varphi \equiv C_A\psi$, $\text{ATEL}_{\text{CE}} \Rightarrow \text{ATL}$. We have $\mathcal{M}, q \models C_A\psi$, so for every sequence of states $q_0 = q, q_1, \dots, q_n$, $q_i \sim_{a_{j_i}} q_{i+1}$, $a_{j_i} \in A$ for $i = 0, \dots, n-1$, it is true that $\mathcal{M}, q_n \models \psi$. Consider the same q in \mathcal{M}^{ATL} . The shape of the construction implies that for every sequence $q'_0 = q, q'_1, \dots, q'_n$ in which every q_{i+1} is a successor of q_i and every $q_{i+1} \in Q^{e_{j_i}}$, $e_{j_i} \in A^e$, we have $\mathcal{M}^{\text{ATL}}, q'_n \models \text{tr}(\psi)$ (by induction and Lemma 22). Moreover, $\mathcal{M}^{\text{ATL}}, q'_i \models e_{j_i}$ for $i \geq 1$, hence $\mathcal{M}^{\text{ATL}}, q'_i \models \bigvee_{a_j \in A} e_j$. Note that the above refers to all the sequences that can be enforced by the agents from Agt^e , and have $\bigvee_{a_j \in A} e_j$ true along the way (from q'_1 on). Thus, Agt^e have no strategy from q such that $\bigvee_{a_j \in A} e_j$ holds from the next state on, and $\text{tr}(\psi)$ is eventually false:
 $\mathcal{M}^{\text{ATL}}, q \not\models_{\text{ATL}} \langle\langle \text{Agt}^e \rangle\rangle X \langle\langle \text{Agt}^e \rangle\rangle (\bigvee_{a_j \in A} e_j) \mathcal{U} (\bigvee_{a_j \in A} e_j \wedge \neg \text{tr}(\psi))$,
 which proves the case.

case $\varphi \equiv C_A\psi$, $\text{ATL} \Rightarrow \text{ATEL}_{\text{CE}}$. We have $\mathcal{M}^{\text{ATL}}, q \models \neg \langle\langle \text{Agt}^e \rangle\rangle X \langle\langle \text{Agt}^e \rangle\rangle (\bigvee_{a_j \in A} e_j) \mathcal{U} (\bigvee_{a_j \in A} e_j \wedge \neg \text{tr}(\psi))$, so

for every σ_{Agt^e} there is $\sigma_{\text{Agt}' \setminus \text{Agt}^e} = \sigma_{\text{Agt}}$ such that $o'_q(\sigma_{\text{Agt}^e}, \sigma_{\text{Agt}}) = q' \in Q'$ and $\mathcal{M}^{ATL}, q' \models \neg \langle \langle \text{Agt}^e \rangle \rangle (\bigvee_{a_j \in A} e_j) \mathcal{U} (\bigvee_{a_j \in A} e_j \wedge \neg tr(\psi))$. In particular, this implies that the above holds for all epistemic states q' that are successors of q in \mathcal{M}^{ATL} , also the ones that refer to agents from A (*).

Suppose that $\mathcal{M}, q \not\models C_A \psi$ (**). Let us now take the action counterpart $q'_{\text{act}} \in Q$ of q' . By (*), (**) and properties of the construction, q'_{act} occurs also in \mathcal{M} , and there must be a path $q_0 = q, q_1 = q'_{\text{act}}, \dots, q_n, q_i \in Q$, such that $q_i \sim_{a_{j_i}} q_{i+1}$ and $\mathcal{M}, q_n \not\models_{ATEL} \psi$. Then, $\mathcal{M}^{ATL}, q_n \not\models_{ATL} tr(\psi)$ (by induction). This means also that we have a sequence $q_0 = q, q'_1 = q', \dots, q'_n$ in \mathcal{M}^{ATL} , in which every $q'_i \in Q^{e_{j_i}}, a_{j_i} \in A$, is an epistemic counterpart of q_i . Thus, for every $i = 1, \dots, n$: $\mathcal{M}^{ATL}, q'_i \models e_{j_i}$, so $\mathcal{M}^{ATL}, q'_i \models \bigvee_{a_j \in A} e_j$. Moreover, $\mathcal{M}^{ATL}, q_n \not\models_{ATL} tr(\psi)$ implies that $\mathcal{M}^{ATL}, q'_n \not\models_{ATL} tr(\psi)$ (by Lemma 22), so $\mathcal{M}^{ATL}, q'_n \models \neg tr(\psi)$. Thus, $\mathcal{M}^{ATL}, q' \models \langle \langle \text{Agt}^e \rangle \rangle (\bigvee_{a_j \in A} e_j) \mathcal{U} (\bigvee_{a_j \in A} e_j \wedge \neg tr(\psi))$, which contradicts (*).

As an immediate corollary of the last two lemmata we obtain:

THEOREM 24. *For every ATEL_{CE} formula φ and model \mathcal{M} , φ is satisfiable (resp. valid) in \mathcal{M} iff $tr(\varphi)$ is satisfiable (resp. valid) in \mathcal{M}^{ATL} .*

Note that the construction used above to interpret ATEL_{CE} in ATL has several nice complexity properties:

- The vocabulary (set of propositions Π) only increases linearly (and certainly remains finite).
- The set of states in an ATEL -model grows linearly, too: if model \mathcal{M} has n states, then \mathcal{M}^{ATL} has $n' = (k+1)n = O(kn)$ states.
- Let m be the number of transitions in \mathcal{M} . We have $(k+1)m$ action transitions in \mathcal{M}^{ATL} . Since the size of every set $img(q, \sim_a)$ can be at most n , there may be no more than kn epistemic transitions per state in \mathcal{M}^{ATL} , hence at most $(k+1)nk$ in total. Because $m \leq n^2$, we have $m' = O(k^2n^2)$.
- Only the length of formulas may suffer an exponential blow-up, because $tr(\langle \langle A \rangle \rangle G\varphi)$ involves two occurrences of $tr(\varphi)$, and the translation of $\langle \langle A \rangle \rangle \varphi \mathcal{U} \psi$ involves two occurrences of both $tr(\varphi)$ and $tr(\psi)$.⁹ Thus, every nesting of $\langle \langle A \rangle \rangle G\varphi$ and $\langle \langle A \rangle \rangle \varphi \mathcal{U} \psi$ roughly dou-

⁹ We thank an anonymous referee for pointing this out.

bles the size of the translated formula in the technical sense. However, the number of *different subformulas* in the formula only increases linearly. Note that the automata-based methods for model checking (Alur et al., 2002) or satisfiability checking (van Drimelen, 2003) for ATL are based on an automaton associated with the given formula, built from its “subformulas closure” – and their complexity depends on the number of different subformulas in the formula rather than number of symbols.¹⁰

Since the complexity of transforming \mathcal{M} to \mathcal{M}^{ATL} is no worse than $O(n^2)$, and the complexity of ATL model checking algorithm from (Alur et al., 2002) is $O(nml)$, the interpretation defined above can be used, for instance, for an efficient reduction of model checking of ATEL_{CE} formulas to model checking in ATL.

4.4.3. Interpreting models and formulas of full ATEL

Now, in order to interpret the full ATEL we modify the construction by introducing new epistemic agents (and states) indexed not only with individual agents, but with all possible non-empty coalitions:

$$\begin{aligned} \text{Agt}^e &= \{e_A \mid A \subseteq \text{Agt}, A \neq \emptyset\} \\ Q' &= Q \cup \bigcup_{A \subseteq \text{Agt}, A \neq \emptyset} Q^{e_A}, \end{aligned}$$

where Q and all Q^{e_A} are pairwise disjoint. Accordingly, we extend the language with new propositions $\{e_A \mid A \subseteq \text{Agt}\}$. The choices for complex epistemic agents refer to the (epistemic copies of) states accessible via distributed knowledge relations: $\Sigma'_{e_A} = \{pass\} \cup \text{img}(q, \sim_A^D)^{e_A}$. Then we modify the transition function (putting the strategies from epistemic agents in any predefined order):

$$o'_q(\sigma_{a_1}, \dots, \sigma_{a_k}, \dots, \sigma_{e_A}, \dots) = \begin{cases} o_q(\sigma_{a_1}, \dots, \sigma_{a_k}) & \text{if all } \sigma_{e_A} = pass \\ \sigma_{e_A} & \text{if } e_A \text{ is the first active} \\ & \text{epistemic agent} \end{cases}$$

Again, the game frames for all epistemic copies of the action states are the same. The translation for all operators remain the same as well

¹⁰ In fact, we can avoid the exponential growth of formulas in the context of satisfiability checking by introducing a new propositional variable p and requiring that it is universally equivalent to $tr(\varphi)$. We obtain

$$tr(\langle\langle A \rangle\rangle G\varphi) = \langle\langle \emptyset \rangle\rangle G(p \leftrightarrow tr(\varphi)) \wedge (p \wedge \langle\langle A \cup \text{Agt}^e \rangle\rangle X \langle\langle A \cup \text{Agt}^e \rangle\rangle G(\text{act} \wedge p)),$$

plus a similar translation for $\langle\langle A \rangle\rangle \varphi \mathcal{U} \psi$.

An analogous method can be proposed for model checking. To translate $\langle\langle A \rangle\rangle G\varphi$, we first use the algorithm from (Alur et al., 2002) and model-check $tr(\varphi)$ to find the states $q \in Q$ in which $tr(\varphi)$ holds. Then we update the model, adding a new proposition p that holds exactly in these states, and we model-check $(p \wedge \langle\langle A \cup \text{Agt}^e \rangle\rangle X \langle\langle A \cup \text{Agt}^e \rangle\rangle G(\text{act} \wedge p))$ as the translation of $\langle\langle A \rangle\rangle G\varphi$ in the new model. We can tackle $tr(\langle\langle A \rangle\rangle \varphi \mathcal{U} \psi)$ in a similar way.

(just using $e_{\{i\}}$ instead of e_i) and the translation of D_A is:

$$tr(D_A \varphi) = \neg \langle\langle \text{Agt}^e \rangle\rangle X(e_A \wedge \neg tr(\varphi)).$$

The following result can now be proved similarly to Theorem 24.

THEOREM 25. *For every ATEL formula φ and model \mathcal{M} , φ is satisfiable (resp. valid) in \mathcal{M} iff $tr(\varphi)$ is satisfiable (resp. valid) in \mathcal{M}^{ATL} .*

This interpretation requires (in general) an exponential blow-up of the original ATEL model (in the number of agents k). We suspect that this may be inevitable – if so, this tells something about the inherent complexity of the epistemic operators. For a specific ATEL formula φ , however, we do not have to include all the epistemic agents e_A in the model – only those for which D_A occurs in φ . Also, we need epistemic states only for these coalitions. Note that the number of such coalitions is never greater than the length of φ . Thus, the “optimized” transformation gives us a model with $n' = O((k+l)n)$ states and $m' = O((k+l)n^2)$ transitions, while the new formula $tr(\varphi)$ is again only linearly longer than φ (in the sense explained in Section 4.4.2). In consequence, we can still use the ATL model checking algorithm for polynomial model checking of ATEL formulas – the complexity of such process is $O(kl(k+l)^2n^3)$.

4.4.4. Related work

The interpretation presented in this section has been inspired by (Schild, 2000) in which a propositional variant of the BDI logic (Rao and Georgeff, 1991) was proved to be subsumed by propositional μ -calculus. We use a similar method here to show a translation from ATEL models and formulas to models and formulas of ATL that preserves satisfiability. ATL (just like μ -calculus) is a multimodal logic, where modalities are indexed by agents (programs in the case of μ -calculus). It is therefore possible to “simulate” the epistemic layer of ATEL by adding new agents (and hence new cooperation modalities) to the scope. Thus, the general idea of the interpretation is to translate modalities of one kind to additional modalities of another kind.

Similar translations are well known within modal logics community, including translation of epistemic logic into Propositional Dynamic Logic, translation of dynamic epistemic logic without common knowledge into epistemic logic (Gerbrandy, 1999) etc. A work particularly close to ours is included in (van Otterloo et al., 2003). In that paper, a reduction of ATEL model checking to model checking of ATL formulas is presented, and the epistemic accessibility relations are handled in a similar way to our approach, i.e. with use of additional “epistemic”

agents. We believe, however, that our translation is more general, and provides more flexible framework in many respects:

1. The algorithm from (van Otterloo et al., 2003) is intended only for *turn-based acyclic* transition systems, which is an essential limitation of its applicability. Moreover, the set of states is assumed to be finite (hence only finite trees are considered). There is no restriction like this in our method.
2. The language of ATL/ATEL is distinctly reduced in (van Otterloo et al., 2003): it includes only “sometime” (F) and “always” (G) operators in the temporal part (neither “next” nor “until” are treated), and the individual knowledge operator K_a (the group knowledge operators C, E, D are absent).
3. The translation of a model in (van Otterloo et al., 2003) depends heavily on the formula one wants to model-check, while in the algorithm presented here, formulas and models are translated independently (except for the sole case of efficient translation of distributed knowledge).
4. Our intuition is that our interpretation is also more general in the sense that it can work in contexts other than model checking. We plan to apply the same translation scheme to reduce the satisfiability problem from ATEL to ATL, for instance.

5. Concluding Remarks

We have presented a comparative study of several logics that combine elements of game theory, temporal logics and epistemic logics, and demonstrated their relationship. Still, these enterprises differ in their motivations and agendas. We wanted to show them as parts of a bigger picture, so that one can compare them, appreciate their similarities and differences, and choose the system most suitable for the intended applications.

Notably, the systems studied here can benefit from many ideas and results, both technical and conceptual, borrowed from each other. Indeed, ATL has already benefited from being related to coalitional games, as concurrent game structures provide a more general (and natural) semantics than alternating transition systems. Moreover, coalition effectivity models are mathematically simpler and more elegant, and provide technically handier semantics, essentially based on neighborhood semantics for non-normal modal logics (Parikh, 1985; Pauly,

2000a). Furthermore, the pure game-theoretical perspective of coalition logics can offer new ideas to the framework of open multi-agent systems and computations formalized by ATL. For instance, fundamental concepts in game theory, such as *preference relations between outcomes*, and *Nash equilibria* have their counterparts in concurrent game structures (and, more importantly, in the alternating-time logics) which are unexplored yet.

On the other hand, the language and framework of ATL has widened the perspective on coalitional games and logics, providing a richer and more flexible vocabulary to talk about abilities of agents and their coalitions. The *alternating refinement relations* (Alur et al., 1998b) offer an appropriate notion of bisimulation between ATSS and thus can suggest an answer to the question “*When are two coalition games equivalent?*”.¹¹ Also, a number of technical results on expressiveness and complexity, as well as realizability and model-checking methods from (Alur et al., 2002; Alur et al., 1998b) can be transferred to coalition games and logics. And there are some specific aspects of computations in open systems, such as *controllability* and *fairness constraints*, which have not been explored in the light of coalition games.

There were a few attempts to generalize ATL by including imperfect information in its framework: *ATL with incomplete information* in (Alur et al., 2002), ATEL, ATOL, ATEL-R* etc. It can be interesting to see how these attempts carry over to the framework of CL. Also, stronger languages like ATL* and alternating-time μ -calculus can provide more expressive tools for reasoning about coalition games.

In conclusion, we see the main contribution of the present study as casting a bridge between several logical frameworks for multi-agent systems, and we hope to trigger a synergetic effect from their mutual influence.

Acknowledgements

Valentin Goranko acknowledges the financial support during this research provided by the National Research Foundation of South Africa and the SASOL Research Fund of the Faculty of Science at Rand Afrikaans University. He would like to thank Johan van Benthem for sparking his interest in logical aspects of games, Marc Pauly for stimulating discussions on coalition games and valuable remarks on this paper, Moshe Vardi for useful suggestions and references, and Philippe Balbiani for careful reading and corrections on an earlier version of

¹¹ Cf. the paper “When are two games the same” in (van Benthem, 2000).

the text. Wojciech Jamroga would like to thank Mike Wooldridge and Wiebe van der Hoek who inspired him to have a closer look at logics for multi-agent systems; and Barteld Kooi as well as Wiebe van der Hoek (again) for their remarks and suggestions.

References

- Abdou, J.: 1998, ‘Rectangularity and Tightness: A Normal Form Characterization of Perfect Information Extensive Game Forms’. *Mathematics and Operations Research* **23**(3).
- Alur, R., T. A. Henzinger, and O. Kupferman: 1997, ‘Alternating-Time Temporal Logic’. In: *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*. pp. 100–109, IEEE Computer Society Press. Available at citeseer.nj.nec.com/170985.html.
- Alur, R., T. A. Henzinger, and O. Kupferman: 1998a, ‘Alternating-Time Temporal Logic’. *Lecture Notes in Computer Science* **1536**, 23–60. Available at citeseer.nj.nec.com/174802.html.
- Alur, R., T. A. Henzinger, and O. Kupferman: 2002, ‘Alternating-Time Temporal Logic’. *Journal of the ACM* **49**, 672–713. Updated, improved, and extended text. Available at www.cis.upenn.edu/~alur/Jacm02.pdf.
- Alur, R., T. A. Henzinger, O. Kupferman, and M. Vardi: 1998b, ‘Alternating Refinement Relations’. In: *CONCUR’98*.
- Aumann, R. and S. Hart (eds.): 1992, *Handbook of Game Theory with Economic Applications, Volume 1*. Elsevier/North-Holland.
- Chandra, A., D. Kozen, and L. Stockmeyer: 1981, ‘Alternation’. *Journal of the ACM* **28**(1), 114–133.
- Emerson, E. A.: 1990, ‘Temporal and Modal Logic’. In: J. van Leeuwen (ed.): *Handbook of Theoretical Computer Science*, Vol. B. Elsevier Science Publishers, pp. 995–1072.
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi: 1995, *Reasoning about Knowledge*. MIT Press: Cambridge, MA.
- Gerbrandy, J.: 1999, ‘Bisimulations on Planet Kripke’. Ph.D. thesis, University of Amsterdam.
- Goranko, V.: 2001, ‘Coalition Games and Alternating Temporal Logics’. In: J. van Benthem (ed.): *Proc. of the 8th Conference on Theoretical Aspects of Rationality and Knowledge (TARK VIII)*. pp. 259–272, Morgan Kaufmann. Corrected version. Available at <http://general.rau.ac.za/math/goranko/papers/glt131.pdf>.
- Goranko, V. and G. van Drimmelen: 2003, ‘Complete Axiomatization of the Alternating-time Temporal Logic’. Submitted manuscript.
- Hart, S.: 1992, ‘Games in Extensive and Strategic Forms’. In: R. Aumann and S. Hart (eds.): *Handbook of Game Theory with Economic Applications, Volume 1*. Elsevier/North-Holland, pp. 19–40.
- Huth, M. and M. Ryan: 2000, *Logic in Computer Science: Modelling and reasoning about systems*. Cambridge University Press.
- Jamroga, W.: 2003, ‘Some Remarks on Alternating Temporal Epistemic Logic’. In: B. Dunin-Keplicz and R. Verbrugge (eds.): *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS 2003)*. pp. 133–140.
- Jamroga, W. and W. van der Hoek: 2003, ‘Agents that Know how to Play’. Submitted.

- Lomuscio, A.: 1999, ‘Information Sharing among Ideal Agents’. Ph.D. thesis, University of Birmingham.
- Moore, R.: 1985, ‘A Formal Theory of Knowledge and Action’. In: J. Hobbs and R. Moore (eds.): *Formal Theories of the Commonsense World*. Ablex Publishing Corp.
- Morgenstern, L.: 1991, ‘Knowledge and the Frame Problem’. *International Journal of Expert Systems* **3**(4).
- Osborne, M. and A. Rubinstein: 1994, *A Course in Game Theory*. MIT Press.
- Parikh, R.: 1985, ‘The Logic of Games and its Applications’. *Ann. of Discrete Mathematics* **24**, 111–140.
- Pauly, M.: 2000a, ‘Game Logic for Game Theorists’. Technical Report INS-R0017, CWI.
- Pauly, M.: 2000b, ‘A Logical Framework for Coalitional Effectivity in Dynamic Procedures’. In: *Proceedings of the Conference on Logic and the Foundations of Game and Decision Theory (LOFT4)*. to appear in Bulletin of Economics Research.
- Pauly, M.: 2001, ‘Logic for Social Software’. Ph.D. thesis, University of Amsterdam.
- Pauly, M.: 2002, ‘A Modal Logic for Coalitional Power in Games’. *Journal of Logic and Computation* **12**(1), 149–166.
- Quine, W.: 1956, ‘Quantifiers and Propositional Attitudes’. *Journal of Philosophy* **53**, 177–187.
- Rao, A. and M. Georgeff: 1991, ‘Modeling Rational Agents within a BDI-Architecture’. In: *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*. pp. 473–484.
- Schild, K.: 2000, ‘On the Relationship Between BDI Logics and Standard Logics of Concurrency’. *Autonomous Agents and Multi Agent Systems* pp. 259–283.
- Schnoebelen, P.: 2003, ‘The Complexity of Temporal Model Checking’. In: *Advances in Modal Logics, Proceedings of AiML 2002*. World Scientific.
- van Benthem, J.: 2000, ‘Logic and Games’. Technical Report X-2000-03, ILLC. Lecture notes, available at <http://www.illc.uva.nl/Publications/Lists/www-allpubl-2000.html>.
- van der Hoek, W. and R. Verbrugge: 2002, ‘Epistemic Logic: A Survey’. *Game Theory and Applications* **8**, 53–94.
- van der Hoek, W. and M. Wooldridge: 2002, ‘Tractable Multiagent Planning for Epistemic Goals’. In: C. Castelfranchi and W. Johnson (eds.): *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-02)*. pp. 1167–1174, ACM Press, New York.
- van der Hoek, W. and M. Wooldridge: 2003a, ‘Cooperation, Knowledge and Time – Alternating-time Temporal Epistemic Logic and its Applications’. *Studia Logica* **75**(1), 125–157.
- van der Hoek, W. and M. Wooldridge: 2003b, ‘Towards a Logic of Rational Agency’. *Logic Journal of the IGPL* **11**(2), 135–160.
- van Drimmelen, G.: 2003, ‘Satisfiability in Alternating-time Temporal Logic’. In: *Proceedings of LICS’2003*.
- van Otterloo, S., W. van der Hoek, and M. Wooldridge: 2003, ‘Knowledge as Strategic Ability’. *Electronic Lecture Notes in Theoretical Computer Science* **85**(2).
- von Neumann, J. and O. Morgenstern: 1944, *Theory of Games and Economic Behaviour*. Princeton University Press: Princeton, NJ.
- Wooldridge, M.: 2000, *Reasoning about Rational Agents*. MIT Press : Cambridge, Mass.