From Coalition Logic to STIT [1]

Jan Broersen [2]

Department of Information and Computing Sciences
Universiteit Utrecht
Utrecht, The Netherlands

Andreas Herzig [3]

Institut de Recherche en Informatique de Toulouse
Université Paul Sabatier
Toulouse, France

Nicolas Troquard [4]

Institut de Recherche en Informatique de Toulouse
Université Paul Sabatier, Toulouse, France
Laboratorio di Ontologia Applicata
Università degli Studi di Trento, Trento, Italy

Abstract

STIT is a logic of agency that has been proposed in the nineties in the domain of philosophy of action. It is the logic of constructions of the form "agent a sees to it that φ". We believe that STIT theory may contribute to the logical analysis of multiagent systems. To support this claim, in this paper we show that there is a close relationship with more recent logics for multiagent systems. We focus on Pauly's Coalition Logic and the logic of the *cstit* operator, as described by Horty. After a brief presentation of Coalition Logic and an adapted discrete-time STIT framework, we introduce a translation from Coalition Logic to STIT, and prove that it is correct.

Key words: multiagent systems, agency, Coalition Logic, STIT theory, modal logic

[1] We would like to thank Marc Pauly for stimulating discussions about Coalition Logic, Laure Vieu and Claudio Masolo for valuable remarks on this paper.
[2] Email: broersen@cs.uu.nl
[3] Email: herzig@irit.fr
[4] Email: troquard@irit.fr

This is a preliminary version. The final version will be published in
Electronic Notes in Theoretical Computer Science
URL: www.elsevier.nl/locate/entcs

1 Introduction

STIT is a logic of agency that has been proposed in the nineties in the domain of philosophy of action [2]. It is the logic of constructions of the form "agent a sees to it that φ".

Several versions of this modality have been studied in the philosophical literature. Here we use the simplest one, viz. the so-called Chellas' STIT operator (*cstit*) [12]. Chellas' STIT operator has been generalized to groups of agents in [13]. Other versions such as the more complex deliberative STIT operator can be defined from Chellas'.

The semantics of the STIT operator is based on branching time temporal structures. In this it differs from the "bringing it about" operator whose semantics is defined in terms of neighborhood models that do not refer to time [16,5,14]. As a consequence it is more natural to study the interaction of agency and time in a STIT setting than in a "bringing it about" setting.

Up to now, the STIT operator has been used mainly in the logical analysis of agency and its relation with deontic concepts [13,12]. Nevertheless we believe that STIT theory may contribute to the logical analysis of multiagent systems. To support this claim, in this paper we show that there is a close relationship with more recent logics for multiagent systems.

We focus on Pauly's Coalition Logic (CL) [15] and the logic of the *cstit* operator, as described by Horty [12]. CL has been introduced to reason about what single agents and groups of agents are able to achieve. $[A]\varphi$ reads "group A can enforce an outcome state satisfying φ". As shown by Goranko in [9], CL is a fragment of Alternating-time Temporal Logic (ATL) that has been proposed by Alur et al. [1]. In this paper we propose a translation from CL to a discrete version of STIT.

In [17], a close examination of the differences and similarities of the models of STIT theory and ATL is undertaken. It is shown that, under the addition of some specific conditions, the models of the two systems can be seen to obey similar properties. However, these properties are not necessarily expressible in the logics of STIT or ATL. So, although, from a philosophical point of view, it may be interesting to look at properties of models as such, here we are essentially interested only in those properties that are expressible in the logics. Where [17] compares the models underlying the logics of ATL (and thus CL) and STIT, we directly compare the logics of both systems.

In section 2 we offer a brief presentation of Coalition Logic. Section 3 deals with an adapted discrete-time STIT framework. Section 4 presents the main result of this note: we describe a translation from CL to STIT, and prove that it is correct. We discuss it in Section 5. Section 6 concludes with some perspectives of investigations.

2 CL-models

In what follows, $\mathcal{A}tm$ represents a set of atomic propositions, and $\mathcal{A}gt$ is the set of agents.

A *game model* is a tuple $\mathcal{M} = \langle W, \{\Sigma_{a,w} | a \in \mathcal{A}gt, w \in W\}, o, v\rangle$, where:

- W is a nonempty set of possible worlds (alias moments or states).
- $\Sigma_{a,w}$ is a set of choices (alias actions) for each agent $a \in \mathcal{A}gt$ and moment $w \in W$. From some (abstract) set of actions, a *particular* choice $\sigma_{A,w}$ of a *group* of agents $A \subseteq \mathcal{A}gt$ in a world w is defined as $\sigma_{A,w} \in \prod_{a\in A} \Sigma_{a,w}$.
- o is a function $o : \prod_{a\in\mathcal{A}gt} \Sigma_{a,w} \mapsto W$ yielding a unique outcome state for every combination of choices by agents in $\mathcal{A}gt$. Thus, if every agent in $\mathcal{A}gt$ opts for an action, the next state of the world is completely determined. Following Pauly, as the occasion arises we slightly generalize the type of the function o, such that it may take two arguments; $o(\sigma_{A,w}, \sigma_{(\mathcal{A}gt\setminus A),w})$ then yields the unique outcome state where the agents in A choose $\sigma_{A,w}$ and the agents in the complementary set $\mathcal{A}gt \setminus A$ choose $\sigma_{(\mathcal{A}gt\setminus A),w}$. Now we can generalize the function o to a function $o : \prod_{a\in A} \Sigma_{a,w} \mapsto 2^W$ mapping for each moment w, the choices of a group of agents A into a set of possible outcome states, by defining: $o(\sigma_{A,w}) = \{o(\sigma_{A,w}, \sigma_{(\mathcal{A}gt\setminus A),w}) \mid \sigma_{(\mathcal{A}gt\setminus A),w} \in \prod_{a\in(\mathcal{A}gt\setminus A)} \Sigma_{a,w}\}$.
- v is a valuation function $v : \mathcal{A}tm \mapsto 2^W$.

Relation with Pauly's original game structures

Pauly defines the semantics of CL using models $\mathcal{M} = (W, E, V)$, where W is a non-empty set of states, E is a playable effectivity function $W \mapsto (2^{\mathcal{A}gt} \mapsto 2^{2^W})$ yielding for every state a function mapping sets of agents A to actions, understood as the set of states A's simultaneous actions result in. Playable effectivity functions are defined to obey some specific conditions, making CL-frames equivalent to game frames (as Pauly proves).

The above definition of game structures differs from Pauly's in two points. First of all, we do not have the agent names as a separate set in the models. Also in the STIT models we define in section 3, contrary to usual practice in STIT-semantics, we do not include the set of agents in the models. This is not necessary, since the agent domains of the functions $\Sigma_{a,w}$ and o are simply assumed to consist of all agents relevant for the interpretation of formulas (like the domain of the valuation function v is assumed to consist of all proposition symbols relevant for the interpretation of formulas). The other difference is that Pauly uses action sets Σ_a, while we make these sets not only relative to agents, but also to worlds (i.e. $\Sigma_{a,w}$). This difference is only cosmetical. Pauly uses one set of choices (choice names) per agent (Σ_a) that is 'reused' in every world. We do not reuse choices (choice names), but use a separate set $\Sigma_{a,w}$ for every agent/world pair instead. The underlying philosophical question is whether or not two choices are always different when performed in different worlds. It is quite easy to see that the two ways of referring to choices do not have any influence on the logic. In CL (and in ATL), the actions (choices) are not made explicit in the object language. Therefore, the logic does not depend on the way we name or refer to actions (choices) in the models. The only difference then seems that in Pauly's setting, the *number* of choices in every state of a model is the same, while in our setting this is not necessarily the case. But also this is not essential, since, without affecting satisfiability, in any of our models we can always introduce dummy choices (e.g. duplicates of existing choices) to make the number of choices equal for each world.

Truth conditions

A formula is evaluated with respect to a model and a moment.

$$\mathcal{M}, w \models p \iff w \in v(p), p \in \mathcal{A}tm$$

$$\mathcal{M}, w \models \neg\varphi \iff \mathcal{M}, w \not\models \varphi$$

$$\mathcal{M}, w \models \varphi \vee \psi \iff \mathcal{M}, w \models \varphi \text{ or } \mathcal{M}, w \models \psi$$

We define the semantic of the modality $[A]\varphi$, whose intuitive interpretation is that the group of agents A can enforce, in one move, an outcome moment satisfying φ, as follows:

$\mathcal{M}, w \models [A]\varphi \iff \exists \sigma_{A,w} \in \prod_{a\in A} \Sigma_{a,w}, \forall w' \in o(\sigma_{A,w}), \mathcal{M}, w' \models \varphi$.

$\models_{CL} \varphi$ denotes that $\mathcal{M}, w \models \varphi$ for every CL-model $\mathcal{M}$ and world w in $\mathcal{M}$.

The following complete axiomatization of CL validities is given in [15]:

(⊥) $\neg[A]\bot$

(⊤) $[A]\top$

(N) $\neg[\emptyset]\neg\varphi \to [\mathcal{A}gt]\varphi$

(M) $[A](\varphi \wedge \psi) \to [A]\varphi$

(S) $[A_1]\varphi \wedge [A_2]\psi \to [A_1 \cup A_2](\varphi \wedge \psi)$ if $A_1 \cap A_2 = \emptyset$

(RE) from $\varphi \equiv \psi$ infer $[A]\varphi \equiv [A]\psi$

Figure 1 shows an example where $\Sigma_{a,w_0} = \{\sigma_a^1, \sigma_a^2\}$. At the moment w_0, the agent a can enforce q (by choosing the action σ_a^1 leading to the moment w_1): $\mathcal{M}, w_0 \models [a]q$.

3 Discrete STIT-models

The semantics of STIT is embedded in the branching time framework. It is based on structures of the form $\langle W, < \rangle$, in which W is a nonempty set of moments, and $<$ is a tree-like ordering of these moments, such that for any w_1, w_2 and w_3 in W, if $w_1 < w_3$ and $w_2 < w_3$, then either $w_1 = w_2$ or $w_1 < w_2$ or $w_2 < w_1$. Moreover, we here constrain the $<$ to be a *discrete* ordering, that is to say that, given a moment w_1, there exists a successor moment w_2 such that $w_1 < w_2$ and there is no moment w_3 such that $w_1 < w_3 < w_2$.

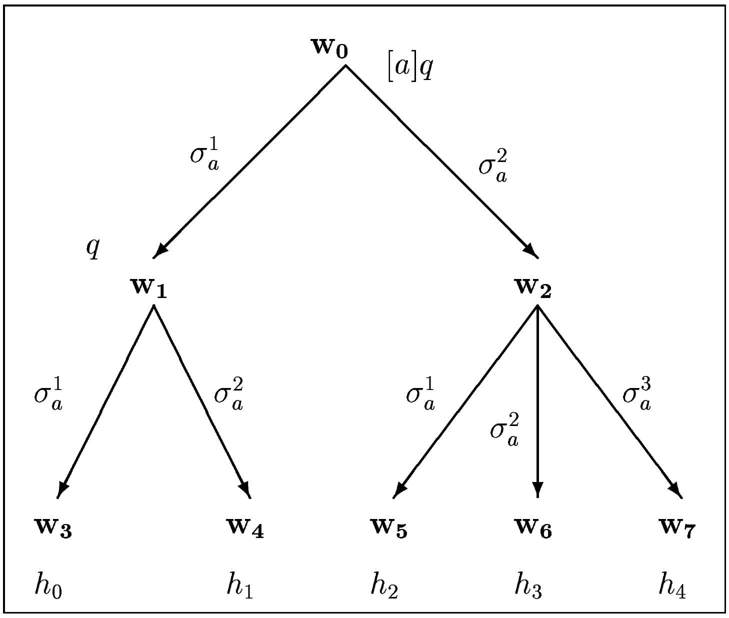

Fig. 1. Example of CL-model.

A maximal set of linearly ordered moments from W is a *history*. Thus, $m \in h$ denotes that the moment m is 'on' the history h. We define $Hist$ as the set of all histories of a STIT structure. $H_w = \{h | h \in Hist, w \in h\}$ denotes the set of histories passing through w. An *index* is a pair w/h, consisting of a moment w and a history h from H_w (i.e. a history and a moment in that history).

A *STIT-model* is a tuple $\mathcal{M} = \langle W, Choice, <, v \rangle$, where:

- $\langle W, < \rangle$ is a branching-time structure.
- $Choice : \mathcal{A}gt \times W \mapsto 2^{2^{Hist}}$ is a function mapping each agent and each moment w into a partition of H_w. The equivalence classes belonging to $Choice_a^w$ can be thought of as possible choices or actions available to a at w. Given a history h, $Choice_a^w(h)$ represents the particular choice from $Choice_a^w$ containing h, or in other words, the particular action performed by a at the index w/h. We must have $Choice_a^w \neq \emptyset$ and $Q \neq \emptyset$ for every $Q \in Choice_a^w$. [5]
- v is valuation function $v : \mathcal{A}tm \mapsto 2^{W \times Hist}$.

In STIT-models, moments may have different valuations, depending on the history they are living in (cf. [13, footnote 2 p. 586]). Thus, at any specific moment, we have different valuations corresponding to the results of the different (non-deterministic) actions possibly taken at that moment.

In order to deal with group agency, Horty defines in [12, section 2.4], the notion of collective choice. Horty first introduces action selection functions s_w

[5] Note that these constraints are not explicited in [13], but they are necessary to ensure validity of the axioms they present.

from $\mathcal{A}gt$ into 2^{H_w} satisfying the condition that for each $w \in W$ and $a \in \mathcal{A}gt$, $s_w(a) \in Choice_a^w$. So, a selection function s_w selects a particular action for each agent at w.

Then, for a given w, $Select_w$ is the set of all selection functions s_w. For every $s_w \in Select_w$, it is assumed that $\bigcap_{a \in \mathcal{A}gt} s_w(a) \neq \emptyset$. This constraint corresponds to the hypothesis that the agents' choices are independent, in the sense that agents can never be deprived of choices due to the choices made by other agents.[6] Moreover, in order to match CL, we assume that $\bigcap_{a \in \mathcal{A}gt} s_w(a)$ is exactly $H_{w'}$ of a next moment w'. As explained in [8], this determinism is not a limitation of the modelling capabilities of the langauge, since we could introduce a neutral agent "nature", in order to accommodate non-deterministic transitions.

Using choice functions s_w, the $Choice$ function can be generalized to apply to groups of agents ($Choice : 2^{\mathcal{A}gt} \times W \mapsto 2^{2^{Hist}}$). A collective choice for a group of agents $A \subseteq \mathcal{A}gt$ is defined as:

$$Choice_A^w = \{\bigcap_{a \in A} s_w(a) | s_w \in Select_w\}$$

Semantics

A formula is evaluated with respect to a model and an index.

$$\mathcal{M}, w/h \models p \iff w/h \in v(p), p \in \mathcal{A}tm.$$

$$\mathcal{M}, w/h \models \neg\varphi \iff \mathcal{M}, w/h \not\models \varphi$$

$$\mathcal{M}, w/h \models \varphi \vee \psi \iff \mathcal{M}, w/h \models \varphi \text{ or } \mathcal{M}, w/h \models \psi$$

Historical necessity at a moment w in a history is defined as truth in all histories passing through w:

$\mathcal{M}, w/h \models \Box\varphi \iff \mathcal{M}, w/h' \models \varphi, \forall h' \in H_w$.

When $\Box\varphi$ holds at w then φ is said to be settled true at w. $\Diamond\varphi$ is defined in the usual way as $\neg\Box\neg\varphi$, and stands for historical possibility.

There are several STIT operators; we here just introduce the so-called Chellas' STIT which is defined as follows:

$\mathcal{M}, w/h \models [A\, cstit : \varphi] \iff \mathcal{M}, w/h' \models \varphi, \forall h' \in Choice_A^w(h)$.

Intuitively it means that group A's current choices ensure φ, whatever other agents outside A do.

As it is shown in [13], both Chellas' STIT and historical necessity are S5 modal operators, and $\models_{STIT} \Box\varphi \rightarrow [A\, cstit : \varphi]$.

As we have discrete time, we can also define the $\mathbf{X}$ operator:

[6] Note that from this constraint it follows that two agents cannot possibly have an identical set of choices at the same moment. It also follows that there are not less than $\prod_{a \in \mathcal{A}gt} |\Sigma_{a,w}|$ histories passing through a moment w. Moreover, at moments where the minimal number of histories satisfies this constraint, choices at future moments will be vacuous.

$\mathcal{M}, w/h \models \mathbf{X}\varphi \iff \exists w' \in h\, (w < w', \mathcal{M}, w'/h \models \varphi, \nexists w'' \in h\, (w < w'' < w'))$.

We note $\models_{STIT} \varphi$ if $\mathcal{M}, w/h \models \varphi$ for every STIT-model $\mathcal{M}$, h in $\mathcal{M}$ and moment w in h.

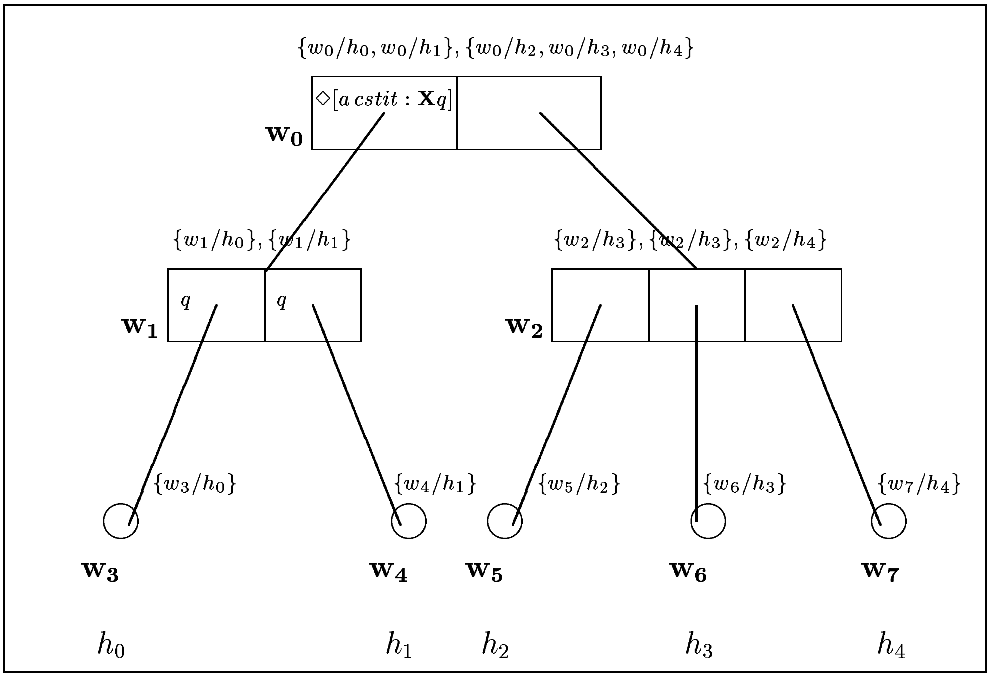

Fig. 2. Example of STIT-model. Indexes of a same choice partition are collected together.

Figure 2 is an example of STIT-model. The moments w_0, w_1 and w_2 are divided in several choices (alias subsets of the sets of histories passing through them). We can see that we have as many indexes in a moment as the number of histories passing through it. $\Diamond[a\, cstit : \mathbf{X}q]$ is true at w_0, because every index in the left choice partition of w_0 (viz. w_0/h_0 and w_0/h_1), leads to a next index where q is true (viz. w_1/h_0 and w_1/h_1).

4 From CL-models to STIT-models

Now we define a translation from CL formulae to $STIT$ formulae:

$$\begin{aligned} tr(p) &:= \Box p, \text{ for } p \in \mathcal{A}tm \\ tr(\neg\varphi) &:= \neg tr(\varphi) \\ tr(\varphi \vee \psi) &:= tr(\varphi) \vee tr(\psi) \\ tr([A]\varphi) &:= \Diamond[A\, cstit : \mathbf{X} tr(\varphi)] \end{aligned}$$

Note that $\models_{STIT} tr(\varphi) \equiv \Box tr(\varphi)$ (the proof uses the fact that the logic of historical necessity $\Box$ is S5).

Theorem 4.1 *If φ is CL-satisfiable then $tr(\varphi)$ is STIT-satisfiable.*

Proof. For any game model $\mathcal{M}_{CL} = \langle W_{CL}, \{\Sigma_{a,w} | a \in \mathcal{A}gt, w \in W_{CL}\}, o, v_{CL}\rangle$, we define $\mathcal{M}'_{CL} = \langle T_{CL}, \{\Sigma_{a,w} | a \in \mathcal{A}gt, w \in T_{CL}\}, o, v_{CL}\rangle$ to be the game model that results from unravelling the function o into a tree (thus, $W_{CL} \subseteq T_{CL}$, where the possible difference between these sets are 'semantically indistinguishable duplicates' of worlds in W_{CL}). From similar results in monotone modal logic [10] and normal modal logic [4], it is immediately clear that the unravelled model is satisfiable if the original model is. The second step is to associate with every game-model $\mathcal{M}'_{CL} = \langle T_{CL}, \{\Sigma_{a,w} | a \in \mathcal{A}gt, w \in T_{CL}\}, o, v_{CL}\rangle$, a STIT model $\mathcal{M}'_{STIT} = \langle W_{STIT}, Choice, <, v_{STIT}\rangle$, satisfying the following conditions:

- $W_{STIT} = T_{CL}$
- $w < u \iff \exists v_1, \ldots, v_n\, (v_1 = w, v_n = u, \forall i < n\, (\exists \sigma_{\mathcal{A}gt,v_i}\, (o(\sigma_{\mathcal{A}gt,v_i}) = v_{i+1})))$
- $\forall a, \forall w, Choice^w_a = \{\{h \mid h \cap o(\sigma_{a,w}) \neq \emptyset\} | \sigma_{a,w} \in \Sigma_{a,w}\}$
  Thus, every element of $Choice^w_a$ collects the histories (recall that these are sets of states) passing through the outcomes $o(\sigma_{a,w})$ for some action $\sigma_{a,w}$.
- $\forall w, \forall h \in H_w, v_{STIT}(w/h) = v_{CL}(w)$

It is straightforward to check that the conditions ensure that an associated model $\mathcal{M}'_{STIT}$ is indeed a discrete STIT-model, and that for any game tree model there is always exactly one such an associated model.

We now prove that any associated STIT model satisfying the conditions is satisfiable if the game tree model it is associated to is satisfiable. That is, we prove (by structural induction on φ) that $\mathcal{M}'_{CL}, w \models \varphi$ only if $\mathcal{M}'_{STIT}, w/h \models tr(\varphi), \forall h \in H_w$. Cases of atomic formulae, negations and disjunctions are trivial, and we here only present the case where $\varphi = [A]\psi$. $\mathcal{M}'_{CL}, w \models [A]\psi$ means that there exists a $\sigma_{A,w}$ such that for all $u \in o(\sigma_{A,w})$ we have $\mathcal{M}'_{CL}, u \models \psi$. So by induction hypothesis, for all $u \in o(\sigma_{A,w})$ and for all $h \in H_u$, $\mathcal{M}'_{STIT}, u/h \models tr(\psi)$. By construction of $Choice^w_A$, this is true only if there is a partition choice $Q \in Choice^w_A$ such that for all histories $h \in Q$ we have $\mathcal{M}'_{STIT}, u/h \models tr(\psi)$. By construction of $<$, this means that $\mathcal{M}'_{STIT}, w/h \models \mathbf{X}tr(\psi)$. We also can deduce that for all $h \in Q$ we have $\mathcal{M}'_{STIT}, w/h \models [A\, cstit : \mathbf{X}tr(\psi)]$. And then for all $h \in H_w$ we have $\mathcal{M}'_{STIT}, w/h \models \Diamond[A\, cstit : \mathbf{X}tr(\psi)]$. ■

Theorem 4.2 *If* $\models_{CL} \varphi$ *then* $\models_{STIT} tr(\varphi)$.

Proof. Instead of semantical proof, we use the axiomatization of [15]: we prove that the translations of the axioms are valid, and that the translated inference rules preserve validity.

It is obvious that the translation of (RE) preserves validity.

The translation of axiom ($\bot$) is $\neg\Diamond[A\, cstit : \mathbf{X}\bot]$, which is equivalent to $\Box\neg[A\, cstit : \mathbf{X}\bot]$. To see that this is valid note first that $\models_{STIT} \neg[A\, cstit : \bot]$ because each element of $Choice^w_a$ is nonempty, for any a and w. ¿From the latter it follows that $\models_{STIT} \neg[A\, cstit : \mathbf{X}\bot]$ (because $\bot \equiv \mathbf{X}\bot$, and because $[A\, cstit : \_]$ satisfies the rule of substitution of equivalences). Then by the necessitation rule for $\Box$ (which is valid in STIT-models) we get $\models_{STIT} \Box\neg[A\, cstit : \mathbf{X}\bot]$.

The translation of axiom ($\top$) is $\Diamond[A\, cstit : \mathbf{X}\top]$. First, $\models_{STIT} [A\, cstit : \top]$ because the rule of necessitation holds for $[A\, cstit : \_]$. Second, $\models_{STIT} [A\, cstit : \mathbf{X}\top]$ (because $\top \equiv \mathbf{X}\top$, and because $[A\, cstit : \_]$ satisfies the rule of substitution of equivalences). Third, $\models_{STIT} \Diamond[A\, cstit : \mathbf{X}\top]$ because $\models_{STIT} \psi \to \Diamond\psi$ holds for any ψ (due to the T-axiom for $\Box$).

The translation of (N) is $\neg\Diamond[\emptyset\, cstit : \mathbf{X}\neg tr(\varphi)] \to \Diamond[\mathcal{A}gt\, cstit : \mathbf{X}tr(\varphi)]$. This is valid in STIT because $\mathbf{X}\neg\psi \equiv \neg\mathbf{X}\psi$ and $\models_{STIT} \neg[\emptyset\, cstit : \neg\psi] \to [\mathcal{A}gt\, cstit : \psi]$ for all ψ, due to our supplementary condition that $Choice^{w}_{\mathcal{A}gt}$ must be a singleton.

The translation of (M) is $\Diamond[A\, cstit : \mathbf{X}(tr(\varphi) \wedge tr(\psi))] \to \Diamond[A\, cstit : \mathbf{X}tr(\varphi)]$. This is STIT-valid first because $\mathbf{X}$ is a normal modal operator, i.e. $\models_{STIT} \mathbf{X}(tr(\varphi) \wedge tr(\psi)) \to \mathbf{X}tr(\varphi)$. And second because $[A\, cstit : \_]$ and $\Box$ are also normal modal operator: from $\gamma_1 \to \gamma_2$ we can infer $[A\, cstit : \gamma_1] \to [A\, cstit : \gamma_2]$ and $\Diamond[A\, cstit : \gamma_1] \to \Diamond[A\, cstit : \gamma_2]$.

The translation of axiom (S) is $\Diamond[A_1\, cstit : \mathbf{X}tr(\varphi)] \wedge \Diamond[A_2\, cstit : \mathbf{X}tr(\psi)] \to \Diamond[A_1 \cup A_2\, cstit : \mathbf{X}(tr(\varphi) \wedge tr(\psi))]$, for $A_1 \cap A_2 = \emptyset$. This is valid in STIT because the choices of a group A in STIT-models are constructed just as the outcome function of A in CL, viz. by pointwise intersection. ■

Theorem 4.3 *φ is satisfiable in CL iff $tr(\varphi)$ is satisfiable in STIT.*

Proof. As an immediate corollary of theorems 4.1 and 4.2. ■

Figure 2 is the STIT-model obtained from the CL-model of Figure 1. At moment w_0, the $Choice^{m}_{a}$ partition is divided in two sets of histories. We have $w_0 < w_1$, and q is settled true at w_1. And $\Diamond[a\, cstit : \mathbf{X}q]$ is true at each index of w_0.

5 Discussion

Figures 1 and 2 can be somewhat misleading. Their purpose was not so much to explain the semantics of Coalition Logic and STIT logic, as to reveal to the essence of the proposed translation. Indeed, the examples suggest that we only deal with one-agent Coalition Logic or with deterministic STIT choices/actions. Considering our translation, deterministic choices are direct consequences of a one-agent Coalition Logic ($\mathcal{A}gt = \{a\}$), since the outcome moment is completely determined when every agent of $\mathcal{A}gt$ has chosen an action. Let us correct that ambiguity by an informal discussion.

Figure 3 shows two representations of a strategic game, in the sense of Pauly. It involves two agents ($\mathcal{A}gt = \{a, b\}$), and both have two possible choices. The usual representation in STIT is a kind of combination of those two schemes. We keep the tree structure, but choices are explicitly represented

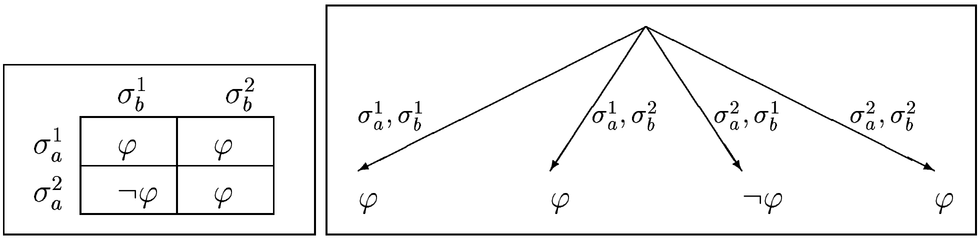

Fig. 3. A strategic game (I)

at each moment.

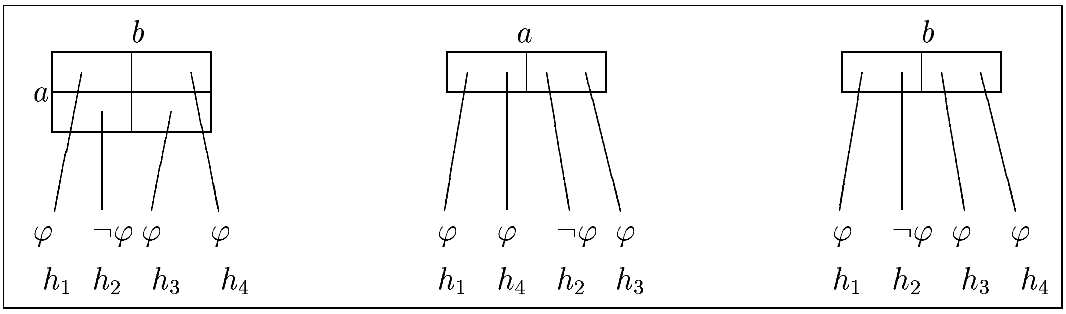

Fig. 4. A strategic game (II)

Left scheme of Figure 4 is the direct translation of the CL-model on Figure 3. Clearly, when a and b opt for a choice, the outcome is deterministic. However, actions of individual agents are not deterministic, as we can see on the two others schemes. And if both a and b can *see to it* that φ independently, they need to form a coalition in order to ensure $\neg\varphi$.

The proposed translation works well, because we have added two more constraints to the *vanilla* STIT:

(i) $<$ must be discrete

(ii) $\forall w \in W, \exists w' \in W\ (w < w'$ and $\nexists w'' \in W, w < w'' < w', \bigcap_{a \in \mathcal{A}gt} s_w(w) = H_{w'})$
Intersection of agents of $\mathcal{A}gt$'s choices is *not only* non-empty but must exactly be the set of histories passing through a next moment.

The first one permits us to define the **X** operator, and then to grasp the concept of *next* moments and outcomes. The second condition makes use of the first one, and is the direct adaptation of the CL constraint stating that when every agent in $\mathcal{A}gt$ opts for an action then the next state of the world is completely determined. We here just say that in STIT, the intersection of agents of $\mathcal{A}gt$'s choices must be exactly the set of histories passing through this very completely determined moment.

Given those extra constraints, the idea behind our translation tr, is that CL exactly matches with discrete STIT, where atomic propositions are historically necessary. As a matter of fact, CL can not support different valuations of an atom at the same state.

Another way of looking at this is to say that the states of CL do not correspond to the *moments* of STIT, but to the *moment/history pairs* of STIT. This actually leads to an alternative (and simpler) translation that, we conjecture, works equally well:

$$
\begin{aligned}
tr2(p) &:= p, \text{ for } p \in \mathcal{A}tm \\
tr2(\neg\varphi) &:= \neg tr2(\varphi) \\
tr2(\varphi \vee \psi) &:= tr2(\varphi) \vee tr2(\psi) \\
tr2([A]\varphi) &:= \Diamond[A\, cstit : tr2(\varphi)]
\end{aligned}
$$

This translation is close to a *simulation* of the weak modal operator $[A]\varphi$ in terms of the two normal S5 modal operators $\Diamond\varphi$ and $[A\ cstit : \varphi]$. Similar simulations have been given for weak deontic modal logics [7,6]. Another simulation of the CL-operator $[A]\varphi$ was recently given by van der Hoek and Wooldridge [11]. They simulate (although they do not use this terminology) $[A]\varphi$ by $\Diamond_A \Box_{\overline{A}} \varphi$, where the diamond and the box are normal modal operators for reasoning about *propositional control*.

6 Conclusion

Given that Coalition Logic is a fragment of Alternating-time Temporal Logic it would be interesting to investigate translations from ATL to STIT. We believe that this can be done, by introducing strategies into the STIT framework as done in [3,12].

A more challenging research avenue is to import deontic concepts that have been investigated in the STIT framework such as in [13,12] into CL and ATL. It seems that this can be done in a rather straightforward manner.

References

[1] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In *Proceedings of the 38th IEEE Symposium on Foundations of Computer Science*, Florida, October 1997.

[2] N. Belnap and M. Perloff. Seeing to it that: A canonical form for agentives. In H. E. Kyburg, R. P. Loui, and G. N. Carlson, editors, *Knowledge Representation and Defeasible Reasoning*, pages 167–190. Kluwer, Boston, 1990.

[3] N. Belnap, M. Perloff, and M. Xu. *Facing the future: agents and choices in our indeterminist world.* Oxford, 2001.

[4] J.F.A.K. van Benthem. Correspondence theory. In D.M. Gabbay and F. Guenthner, editors, *Handbook of philosophical logic, vol. II.* reidel, 1984.

[5] D. Elgesem. *Action Theory and Modal Logic.* PhD thesis, Department of philosophy, University of Oslo, 1993.

[6] O. Gasquet and A. Herzig. Translating non-normal modal logics into normal modal logics. In A.I.J Jones and M. Sergot, editors, *Proceedings International Workshop on Deontic Logic*, TANO, Oslo, 1993.

[7] O. Gasquet and A. Herzig. From Classical to Normal Logics. In Heinrich Wansing, editor, *Proof Theory of Modal Logics*, volume 2 of *Applied Logic Series*, pages 293–311. Kluwer, 1996.

[8] V. Goranko and W.J. Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2):241–280, 2004. ISSN=0039-7857.

[9] Valentin Goranko. Coalition games and alternating temporal logics. In *TARK '01: Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*, pages 259–272, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.

[10] H.H. Hansen and C. Kupke. A coalgebraic perspective on monotone modal logic. In *Proccedings of the 7th Workshop on Coalgebraic Methods in Computer Science (CMCS 2004)*, volume 106 of *Electronic Notes in Theoretical Computer Science*, pages 121–143, 2004.

[11] W. van der Hoek and M. Wooldridge. On the logic of cooperation and propositional control. *Artificial Intelligence*, 164(1-2):81–119, 2005.

[12] John F. Horty. *Agency and Deontic Logic*. Oxford University Press, Oxford, 2001.

[13] John F. Horty and Nuel D. Belnap, Jr. The deliberative STIT: A study of action, omission, and obligation. *Journal of Philosophical Logic*, 24(6):583–644, 1995.

[14] A. Jones and M. Sergot. A formal characterization of institutionalized power. *Journal of the IGPL*, 4(3):429–445, 1996.

[15] Marc Pauly. A modal logic for coalitional power in games. *J. Log. Comput.*, 12(1):149–166, 2002.

[16] I. Pörn. *The Logic of Power*. Oxford: Blackwell, 1970.

[17] S. Wölfl. Qualitative action theory: A comparison of the semantics of alternating time temporal logic and the kutschera-belnap appraoch to agency. In J. Alferes and J. Leite, editors, *Proceedings Ninth European Conference on Logics in Artificial Intelligence (JELIA'04)*, volume 3229 of *Lecture Notes in Artificial Intelligence*, pages 70–81. Springer, 2004.