Strengthening Agents Strategic Ability with Communication

Xiaowei Huang
Jinan University, China
University of Oxford, UK

Qingliang Chen
Jinan University,
Guangzhou 510632, China

Kaile Su
Jinan University, China
IIIS, Griffith University, Australia

Abstract

The current frameworks of reasoning about agents' collective strategy are either too conservative or too liberal in terms of the sharing of local information between agents. In this paper, we argue that in many cases, a suitable amount of information is required to be communicated between agents to both enforce goals and keep privacy. Several communication operators are proposed to work with an epistemic strategy logic ATLK. The complexity of model checking resulting logics is studied, and surprisingly, we found that the additional expressiveness from the communication operators comes for free.

Introduction

Strategic reasoning is an active area in multiagent systems. An extensive set of logic frameworks, see e.g., (Alur, Henzinger, and Kupferman 2002; Horty 2001; Pauly 2002; Chatterjee, Henzinger, and Piterman 2010; Mogavero, Murano, and Vardi 2010), have been proposed to reason about agents' strategic ability. Within these logics, the alternating-time temporal logic ATL (Alur, Henzinger, and Kupferman 2002) is one of the most prominent. To work with *incomplete information systems* in which agents can only make partial observation about the underlying system states, the semantics of the logic has been re-investigated with several proposals, see e.g. (van der Hoek and Wooldridge 2002; Schobbens 2004; van Otterloo and Jonker 2005; Jamroga and Ågotnes 2007; Guelev, Dima, and Enea 2011; Huang and van der Meyden 2014c), etc. In these proposed logic frameworks, a collective strategy of a group of agents is defined as a collection of strategies, one for each agent in the group, and an agent's strategy depends on either its own local information or the group's information, which can be their distributed knowledge (Jamroga and Ågotnes 2007; Guelev, Dima, and Enea 2011) or common knowledge (Jamroga and Ågotnes 2007; Diaconu and Dima 2012). The sharing with distributed knowledge is essentially the sharing of all the available information between agents.

This paper aims to complement these semantic settings with communication between agents. The rationale is based on the following two arguments. The first argument is that, there exist cases in which enabling communication between agents in a group can strengthen the capability of agents and enable them to achieve goals that are not achievable otherwise. The second argument is that, there exist cases in which sharing agents' entire local information is undesirable. For example, an agent is willing to collaborate with other agents to achieve the group's goals, and at the same time, intends to keep some of its own privacy. This relates to, but not the same as, secure multi-party computation (Yao 1982), in which agents want to complete a computation and keep their inputs private.

The approach we suggest in the paper is to introduce communication operators into an epistemic strategy logic ATLK, to quantify over the amount of information needed to satisfy both the goals and the required privacy conditions. Two approaches of defining communication operators are explored. The first is in a style of registration and cancellation: a communication needs to register before working, and once registered, it will be effective until officially cancelled. The second is a temporary communication: the communication occurs and only occurs upon request at the current state.

We studied the complexity of model checking ATLK and its extended logics, and found that, surprisingly, the additional expressiveness of communication operators comes for free. All of them are NEXP-complete for the multiagent systems represented succinctly and symbolically.

s_0	c	d
a	1	0
b	0	0

s_1	c	d
a	1	0
b	0	1

s_2	c	d
a	1	0
b	0	1

s_3	c	d
a	0	0
b	0	1

Table 1: Utilities of different states

Illustrative Example

Example 1 *Consider a game of two players A and B. It starts by tossing two independent coins v and w. Every outcome of the coins represents a possible state. We use bit* 0 *to represent tail and bit* 1 *to represent head. E.g.,* $(1, 0)$ *denotes that coin v lands head and w lands tail. Therefore, there are four states* $s_0 = (0, 0)$, $s_1 = (0, 1)$, $s_2 = (1, 0)$, $s_3 = (1, 1)$. *On every state, the players have two actions* $\{a, b\}$ *and* $\{c, d\}$, *respectively. The utilities of joint actions of agents are given in Table 1.*

The agents are notified with partial information: player

	v	w	$A: v \wedge w$	$A: v \otimes w$	$B: w$
s_0	0	0	0	0	0
s_1	0	1	0	1	1
s_2	1	0	0	1	0
s_3	1	1	1	0	1

Table 2: Agents have partial information about the states

A learns conjunction $v \wedge w$ and exclusive disjunction $v \otimes w$ of the outcome, while player B learns the value of dice w. The details of the states and the information agents have are collected in Table 2. We can see that, agent A can not distinguish states s_1 and s_2, while agent B can not distinguish states s_0 and s_2, and states s_1 and s_3.

A query we would like to make is, do the two agents have a collective strategy to achieve utility 1? A collective strategy is a collection of strategies, one for each agent. As stated in most of the literature, a strategy needs to be uniform, i.e., each agent takes the same action on those states where it has the same local information. Expressed as a formula of ATL language[1]*, the query is*

$$\phi_1 \equiv \langle\langle \{A, B\} \rangle\rangle AX(u = 1). \tag{1}$$

Unfortunately, this is unsatisfiable. As we will elaborate later, a formula is satisfiable on a system if it is satisfiable on all initial states (in this example, s_0, s_1, s_2, s_3) of the system. To see why the formula is unsatisfiable, we notice that on state s_0, to achieve utility 1, agent B has to take action c and agent A has to take action a. Because of the uniformity of agents' strategy, agent B will also take action c on state s_2, which makes agent A take action a. Because agent A can not differentiate states s_1 and s_2, it will take action a on state s_1. To match with this, agent B has to take action c on state s_1 and therefore on state s_3. However, we notice that no matter which action taken by agent A, the utility 1 can not be reached on state s_3.

The approach we study in the paper is to introduce communication between agents so that they can collectively enforce the goal. Communication enhances the ability of agents in distinguishing states and therefore enables better strategies. A naive way of introducing communication is to allow agents in a group to share their entire local information, see e.g., (Guelev, Dima, and Enea 2011). For example, in the game of Example 1, sharing entire local information between the two agents will make them have complete information about the state, and therefore have a collective strategy to satisfy the formula ϕ_1.

However, in many situations, sharing without reservation can be undesirable. A system designer may care not only the ability of agents to enforce goals but also some *security or privacy conditions*, as exemplified in the following example.

Example 2 *For the game of Example 1, the formula*

$$\phi_2 \equiv \neg w \Rightarrow \neg(K_B v \vee K_B \neg v) \tag{2}$$

expresses that agent B does not know the outcome of coin v when coin w lands tail, and the formula

$$\phi_3 \equiv v \otimes w \Rightarrow \neg(K_A v \vee K_A \neg v \vee K_A w \vee K_A \neg w) \tag{3}$$

expresses that agent A does not know the outcome of coins when their exclusive disjunction is 1. The requirement of the game can then be

$$\phi_4 \equiv \phi_2 \wedge \phi_3 \wedge AX(u = 1) \tag{4}$$

It is not hard to see that, allowing agents to share their entire local information will not enable the existence of a collective strategy to satisfy formula ϕ_4, because the conditions ϕ_2 and ϕ_3 do not hold when the agents have complete information.

The approach we explore in the paper is to let the agents transmit a suitable amount of information.

Example 3 *To enable the existence of a collective strategy to satisfy ϕ_4, we may let agent A transmit the value of $v \wedge w$ to agent B. With this message, agent B can distinguish states s_1 and s_3, and therefore the group can have the following collective strategy to satisfy both the goal $AX(u = 1)$ and the privacy conditions ϕ_2 and ϕ_3:*

- *agent A takes action b on state s_3 and a on other states,*
- *agent B takes action d on state s_3 and c on other states.*

In the paper, knowledge formulas like ϕ_2 and ϕ_3 are used to express security or privacy conditions, which in this context mean that agents do not have some specific knowledge about the current state.

Model Checking Multi-agent Systems

Let $\mathcal{B}(Var)$ be the set of boolean formulas over variables[2] Var. For s being a truth assignment of the variables Var and $f \in \mathcal{B}(Var)$ a formula, we write $e(s, f)$ for the evaluation of f on s. We may write $e(s, f)$ (or $\neg e(s, f)$) to denote that $e(s, f) = 1$ ($e(s, f) = 0$).

A multi-agent system consists of a collection of agents running in an environment (Fagin et al. 1995). Let $Agt = \{1, ..., n\}$ be a set of agents. The environment E is a tuple $(Var_e, init_e, \{Acts_i\}_{i\in Agt}, \{OVar_i\}_{i\in Agt}, \longrightarrow_e)$. The component Var_e is a set of environment variables such that every truth assignment to Var_e is an environment state. Let L_e be the set of environment states. The component $init_e \subseteq L_e$ is a set of initial environment states, $OVar_i \subseteq Var_e$ is a subset of environment variables that agent i is able to observe, $Acts_i$ is a set of local actions for agent i such that $Acts_i \cap Acts_j = \emptyset$ if $i \neq j$ and $JActs = \Pi_{i\in Agt} Acts_i$ is a set of joint actions, and $\longrightarrow_e \subseteq L_e \times JActs \times L_e$ is a transition relation. The environment *nondeterministically* updates its state by taking into consideration the joint actions taken by the agents. Agents' observable variables may be overlapping, i.e., $OVar_i \cap OVar_j \neq \emptyset$, to simulate the case where agents have shared variables. We use $s_{e,i}$ to denote the part of an environment state s_e that can be observed by agent i, i.e., $s_{e,i} = s_e|_{OVar_i}$. This can be generalized to a set of states, e.g., $L_{e,i} = \{s_{e,i} \mid s_e \in L_e\}$.

An agent A_i, for $i \in Agt$, is a tuple $(Var_i, init_i, \longrightarrow_i)$. The component Var_i is a set of local variables such that each truth

[1]Here we take a slightly different syntax as that of (Alur, Henzinger, and Kupferman 2002), i.e., a strategy operator can be followed by a CTL formula. The semantics of the language will be given in Section 3.

[2]W.l.o.g., we assume that variables are boolean.

assignment to Var_i is a local state. Let L_i be the set of local states of agent i. The component $init_i \subseteq L_i$ is a set of initial local states, $\longrightarrow_i \subseteq L_i \times L_{e,i} \times Acts_i \times L_i$ is a transition relation: a tuple $(l_i, o_i, a_i, l'_i) \in \longrightarrow_i$ means that when agent i is at state l_i and has an observation o_i on the environment state, it may take action a_i and move into the state l'_i. If there are several a_i with the same l_i and o_i, the agent i will *nondeterministically* choose one of them to execute.

Let $Var = Var_e \cup \bigcup_{i\in Agt} Var_i$ and $Acts = \bigcup_{i\in Agt} Acts_i$. A multi-agent system is defined as $M(E, \{A_i\}_{i\in Agt}) = (S, I, \longrightarrow, \{B_i\}_{i\in Agt})$. The set $S = L_e \times \Pi_{i\in Agt} L_i$ is a set of global states. For a global state $s = (l_e, l_1, ..., l_n)$, we write $s_i \equiv l_i$ for $i \in Agt$, and $s_e \equiv l_e$. The same for joint actions. The set I is a set of initial states such that $s \in I$ if $s_e \in init_e$ and $s_i \in init_i$ for all $i \in Agt$. The transition relation $\longrightarrow \subseteq S \times JActs \times S$ is defined as $(s, a, t) \in \longrightarrow$ if $(s_i, s_{e,i}, a_i, t_i) \in \longrightarrow_i$ for all $i \in Agt$ and $(s_e, a, t_e) \in \longrightarrow_e$. The indistinguishable relation $B_i \subseteq S \times S$ is such that $B_i(s,t)$ iff $s_i = t_i$ and $s_{e,i} = t_{e,i}$. We use $N_i(s) = \{a_i \in Acts_i \mid \exists t_i \in L_i : (s_i, s_{e,i}, a_i, t_i) \in \longrightarrow_i\}$ to denote the set of local actions of agent i that are enabled on global state s. We assume that the environment transition relation $\longrightarrow_e$ is serial, i.e., for every state s and every joint action a such that $a_i \in N_i(s)$ for all $i \in Agt$, there exists a state t such that $(s, a, t) \in \longrightarrow_e$. However, we do not assume the same for agents, i.e., given a local state s_i and an observation $s_{e,i}$, a local action a_i may be disabled.

A (uniform and memoryless) strategy θ_i of agent i maps each state $s \in S$ to a nonempty set of local actions such that $\theta_i(s) \subseteq N_i(s)$ and for all states $s, t \in S$, $B_i(s,t)$ implies $\theta_i(s) = \theta_i(t)$. A strategy θ_i of agent i can be used to update the agent A_i such that all transitions are consistent with θ_i. Formally, for $A_i = (Var_i, init_i, \longrightarrow_i)$, we define $A_i[\theta_i] = (Var_i, init_i, \longrightarrow'_i)$ such that $(t_i, o_i, a_i, t'_i) \in \longrightarrow'_i$ iff $(t_i, o_i, a_i, t'_i) \in \longrightarrow_i$ and $a_i \in \theta_i(s)$ for some global state s with $s_i = t_i$ and $s_{e,i} = o_i$. Moreover, given a collective strategy $\theta_G = \{\theta_i\}_{i\in G}$ of a set G of agents, we define an updated system $M(E, \{A_i\}_{i\in Agt})[\theta_G] = M(E, \{A_i\}_{i\in Agt\setminus G} \cup \{A_i[\theta_i]\}_{i\in G})$. For any updated system M, we write M_0 for the original system where no strategy has been applied.

We use a language ATLK to describe the specifications of a multi-agent system M. Formally, ATLK has the syntax:

$$\phi ::= p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid EX\phi \mid EG\phi \mid E(\phi_1 U \phi_2) \mid \langle\langle G\rangle\rangle\phi \mid E_G\phi$$

where $p \in Var$ is an atomic proposition and $G \subseteq Agt$ is a set of agents. Intuitively, formula $\langle\langle G\rangle\rangle\phi$ means that the agents in G have a collective strategy to enforce ϕ, and $E_G\phi$ means that every agent in G knows ϕ. In particular, we have $E_{\{i\}}\phi = K_i\phi$. Formulas $EX\phi$, $EG\phi$ and $E(\phi_1 U \phi_2)$ have standard meaning as in CTL. Other operators can be obtained in the usual way, e.g., $A\phi = \neg E \neg\phi$, $F\phi = True U\phi$, etc.

A (labelled) fullpath $\rho = s^0 a^1 s^1 ...$ is an infinite sequence of states and actions such that s^0 is an initial state, and for every $k \geq 0$, $(s^k, a^{k+1}, s^{k+1}) \in \longrightarrow$. We use $\rho(m)$ to denote the state s^m. Moreover, we write $Path(M, s)$ for the set of fullpaths ρ of M such that $\rho(0) = s$, and $rch(M)$ for the set of reachable states of M, i.e., $s \in rch(M)$ if there exists a state $t \in I$ such that there exists a fullpath $\rho \in Path(M, t)$ such that $s = \rho(m)$ for some $m \geq 0$.

The semantics of the language on a system M is described as a relation $M, s \models \phi$, which is defined recursively as follows for state $s \in S$ and formula ϕ.

- $M, s \models p$ if $e(s, p)$.
- $M, s \models \neg\phi$ if not $M, s \models \phi$.
- $M, s \models \phi_1 \vee \phi_2$ if $M, s \models \phi_1$ or $M, s \models \phi_2$.
- $M, s \models \langle\langle G\rangle\rangle\phi$ if there exists a collective strategy θ_G such that for all agents $i \in G$, there is $M_0[\theta_G], t \models \phi$ for all states $t \in rch(M_0)$ with $B_i(s,t)$.
- $M, s \models E_G\phi$ if for all agents $i \in G$ and all states $t \in rch(M_0)$ with $B_i(s,t)$, there is $M, t \models \phi$.
- $M, s \models EX\phi$ if $M, t \models \phi$ for some state t with $(s,t) \in \longrightarrow$.
- $M, s \models E(\phi_1 U \phi_2)$ if there exist a path $\rho \in Path(M, s)$ and a number $m \geq 0$ such that $M, \rho(k) \models \phi_1$, for all $0 \leq k \leq m-1$, and $M, \rho(m) \models \phi_2$.
- $M, s \models EG\phi$ if there exist a path $\rho \in Path(M, s)$ such that $M, \rho(k) \models \phi$ for all $k \geq 0$.

Note that, when dealing with formula $\langle\langle G\rangle\rangle\phi$, the strategy θ_G is applied on the original system M_0, instead of the current system M. Moreover, M_0 is used in computing indistinguishable states when interpreting formula $E_G\phi$; agents are incapable of observing the strategies that are currently applied, including its own strategy.

Given a multi-agent system M and a formula ϕ, the model checking problem, written as $M \models \phi$, is to decide whether $M, s \models \phi$ for all $s \in I$.

Adding Fixed Communication

A communication is held by an agent sending a message to another agent via an instantaneous lossless channel. An agent's local state represents the maximal information it has. It is reasonable to assume that a message sent by an agent does not contain more information than its local state[3]. In the paper we assume that agents do not have memory. Therefore, agent i's local state contains current valuation of local variables Var_i and observable environment variables $OVar_i$.

Every communication is associated with a directed channel. A communication is *fixed* if, once the channel is established, the message will be transmitted from the sender to the receiver on every state that follows.

A multiagent system M needs to maintain a set C_M of existing communication. During the establishment stage, a tuple (k, i, j, φ) such that $k = |C_M|$, $i, j \in Agt$ and $\varphi \in \mathcal{B}(OVar_i \cup Var_i)$ is registered/added into C_M. Intuitively, k denotes the index of the new communication, i and j are sender and receiver respectively, and φ is a formula that will be used to interpret future messages. In the following, we use c_k to denote the kth communication that has been registered in C_M, $sndr_k$ and $rcvr_k$ to denote the sender and receiver, and φ_k to denote the message formula.

Once a communication has been registered, the actual message will be transmitted from the sender to the receiver on every state since then (including the current state). For a communication c_k and a state s, agent $sndr_k$ transmits a

[3]We assume that an agent always tells some truth about its current local state.

pair $m_{k,s} = (k, e(s, \varphi_k))$ to agent $rcvr_k$. Recall that $e(s, \varphi_k)$ is the evaluation of formula φ_k on the state s. When there is no confusion, we may simply use formula φ_k to denote a message, as in Example 3.

A message $m_{k,s}$ provides agent $rcvr_k$ with information on the current state s. This additional information may or may not increase the information agent $rcvr_k$ has about the current state. The impact of the transmission of a message $m_{k,s}$ can be expressed as an update to agents' indistinguishable relations. Formally, we define $B_i[c_k]$ as follows: for any two states s and t,

$$B_i[c_k](s,t) = \begin{cases} B_i(s,t) \wedge (e(s,\varphi_k) \Leftrightarrow e(t,\varphi_k)) & \text{if } i = rcvr_k \\ B_i(s,t) & \text{if } i \neq rcvr_k \end{cases}$$

Intuitively, the updated relation has an extra condition that the related states have the same evaluation on the formula φ_k. The update on the indistinguishable relations can be generalized to work with a set C_M of communication. For any two states s and t,

$$B_i[C_M](s,t) \equiv B_i[c_0]...[c_k](s,t) \text{ for } k = |C_M| - 1.$$

The following proposition says that the ordering of applying these communication does not matter.

Proposition 1 *$B_i[C_M](s,t)$ iff $B_i[c](s,t)$ for all $c \in C_M$.*

In this section, to simplify notations, we assume that a received message can not be re-sent directly or sent as a component of a new message, from its receiver to another agent. The case of nested communication will be formally treated in Section 6.

A fixed communication may be cancelled upon request. The cancellation of the communication with agents G as receivers is defined as follows.

$$B_i[C_M \setminus G](s,t) = \begin{cases} B_i(s,t) & \text{if } i \in G \\ B_i(s,t)[C_M] & \text{if } i \notin G \end{cases}$$

where $B_i(s,t)$ is the original definition in the system M. The necessity of cancelling a communication will be elaborated in the next section with an example.

Now we are ready to define the way how a system $M = (S, I, \longrightarrow, \{B_i\}_{i \in Agt})$ maintains a set C_M of communication. This is done by constructing a system $M[C_M] = (S, I, \longrightarrow, \{B'_i\}_{i \in Agt})$ such that $B'_i = B_i[C_M]$ for all $i \in Agt$.

To specifying the registration and cancellation of communication, we introduce into the language ATLK two new operators $\mathtt{S}_G$ and $\mathtt{N}_G$ for $G \subseteq Agt$ being a set of agents. The new language ATLK^{fc} has the syntax as follows.

$$\begin{aligned} \phi \quad ::= \quad & p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid EX\phi \mid EG\phi \mid E(\phi_1 U \phi_2) \mid \\ & \langle\langle G \rangle\rangle \phi \mid E_G\phi \mid \mathtt{S}_G\phi \mid \mathtt{N}_G\phi \end{aligned}$$

The semantics of the operators is defined as follows.

- $M[C], s \models \mathtt{S}_G\phi$ if there exists a set C' of communication between agents in G such that $M[C \cup C'], s \models \phi$.
- $M[C], s \models \mathtt{N}_G\phi$ if $M[C \setminus G], s \models \phi$.
- $M[C], s \models \langle\langle G \rangle\rangle\phi$ if there exists a collective strategy θ_G such that for all agents $i \in G$, there is $M_0[C][\theta_G], t \models \phi$ for all states $t \in rch(M_0[C])$ with $B_i[C](s,t)$.

s_0	c	d
a	1	0
b	0	0

s_1	c	d
a	0	1
b	0	0

s_2	c	d
a	0	0
b	1	0

s_3	c	d
a	0	0
b	0	1

Table 3: Utilities of different states, second round

- $M[C], s \models E_G\phi$ if for all agents $i \in G$ and all states $t \in rch(M_0[C])$ with $B_i[C](s,t)$, there is $M[C], t \models \phi$.

The semantics of other constructs follows the similar pattern as that of Section 3.

With fixed communication, the model checking problem is, given a multiagent system M and a formula ϕ of the language ATLK^{fc}, to decide whether $M[\emptyset] \models \phi$.

Example 4 *The specification stated in Example 2 can now be expressed with the following formula.*

$$\phi_5 \equiv \mathtt{S}_{\{A,B\}}\langle\langle\{A,B\}\rangle\rangle\phi_4 \tag{5}$$

which says that there exist a set of communication for agents A and B such that, the agents have a strategy to enforce ϕ_4. The satisfiability of ϕ_5 on the game M, i.e., $M \models \phi_5$, can be witnessed by the communication given in Example 3.

Necessity of Cancelling a Communication

Fixed communication enhance the system with updated indistinguishable relations of agents. If not cancelled, these updates occur in the entire system execution. This may result in unnecessary communication or undesirable result.

Example 5 *The following is a variant of Example 1 and 2. The game has two rounds. The first round is the same as described before. In the second round, the two coins are flipped again and the agents are notified with the same information. Assume that agents can distinguish states of different rounds. Technically, this can be implemented by relating each state with a round number, e.g., $(1, s_1)$ denotes the state s_1 of the first round and $(2, s_3)$ denotes the state s_3 of the second round, and enhancing agents' indistinguishable relations with a constraint: $n \neq m \Rightarrow \neg B_x((n,s),(m,t))$ for all $x \in \{A,B\}$, $s,t \in \{s_0,s_1,s_2,s_3\}$.*

The utilities of the second round are different with those of the first round, and are given in Table 3. Consider the following formula for the second round

$$\phi_6 \equiv \neg(K_B v \vee K_B \neg v) \wedge AX(u = 1) \tag{6}$$

which says that the required utility is 1 and agent B does not know the outcome of coin v. We notice that the formula $AX\ \mathtt{S}_{\{A,B\}}\langle\langle\{A,B\}\rangle\rangle\phi_6$ can be satisfied by agent B sending a message w to agent A in the second round. With that message, agent A has complete information about the state but agent B's ability is kept the same. Then the group has a collective strategy as follows:

- *agent B takes action c on state s_0 and s_2 and action d on state s_1 and s_3,*
- *agent A takes action a on state s_0 and s_1 and action b on state s_2 and s_3.*

However, the following formula, which combines the requirement of the first round and the second round, is not satisfiable.

$$\phi_7 \equiv \mathsf{S}_{\{A,B\}}\langle\langle\{A,B\}\rangle\rangle(\phi_4 \wedge AX\ \mathsf{S}_{\{A,B\}}\langle\langle\{A,B\}\rangle\rangle\phi_6) \quad (7)$$

Taking both the messages designated for ϕ_4 and for ϕ_5 will cause the agents have complete information about the system state and therefore do not satisfy the security conditions.

This can be handled by the capability of cancelling a communication. Let M_2 be the new game of two rounds, and

$$\phi_8 \equiv \mathsf{S}_{\{A,B\}}\langle\langle\{A,B\}\rangle\rangle(\phi_4 \wedge AX\ \mathtt{N}_{\{A,B\}}\mathsf{S}_{\{A,B\}}\langle\langle\{A,B\}\rangle\rangle\phi_6). \quad (8)$$

We have that $M_2 \models \phi_8$.

Handling Nested Communication

In this section, we consider nested communication which is not covered in the definitions of Section 4. A communication from an agent is nested if it contains information that is received from other agents.

Example 6 *Consider a variant of Example 1 with an additional player C who is not notified with any information about the outcome but wants to discover the outcome of coin v. It is assumed that C can only communicate with B. Intuitively, the following formula*

$$\phi_9 \equiv \mathsf{S}_{\{A,B\}}\mathsf{S}_{\{B,C\}}\langle\langle\{A,B\}\rangle\rangle(\phi_4 \wedge (w \Rightarrow K_C v \vee K_C\neg v)) \quad (9)$$

says that after the communication between A and B, and B and C, along with the previous requirement ϕ_4, an extra condition that C can discover the outcome of v when coin w lands head holds. If no nested communication is allowed, the formula is not satisfiable, because without the information from A, agent B does not know any information about v and thus can not help agent C.

However, the following nested communication may be applied to satisfy the formula ϕ_9. The communication between A and B are the same as suggested in Example 3. The communication between B and C can be done by letting agent B send both $v \wedge w$ and w to C. Note that the message $v \wedge w$ is received from agent A.

According to the definition, a communication c_k from an agent i has to satisfy the condition that $\varphi_k \in \mathcal{B}(OVar_i \cup Var_i)$. To handle nested communication, for every communication c_k, we introduce an extra local variable v_k for agent $rcvr_k$, i.e., let

$$Var_i[c_k] = \begin{cases} Var_i \cup \{v_k\} & \text{if } i = rcvr_k \\ Var_i & \text{if } i \neq rcvr_k \end{cases}$$

This can be generalised to work with a set C_k of communication as $Var_i[C_k] = Var_i[c_0]...[c_k]$ for $k = |C_M| - 1$. The following expression

$$rel \equiv \bigwedge_{j=0}^{|C_M|-1} v_j \Leftrightarrow \varphi_j$$

records the equivalence relations between newly-introduced variables and their corresponding messages. Let Var_{C_M} be the set of newly introduced variables, we can define

$$B_i[c_k](s,t) = \begin{cases} \exists Var_{C_M} : B_i(s,t) \wedge (e(s,\varphi_k) \Leftrightarrow e(t,\varphi_k)) \wedge rel & \\ & \text{if } i = rcvr_k \\ B_i(s,t) & \text{if } i \neq rcvr_k \end{cases}$$

and generalise it to $B_i[C_M]$ and $M[C_M]$.

Example 7 *Continue with Example 6. There are three communication: $c_1 = (1, A, B, v \wedge w)$ will introduce a local variable v_1 for agent B such that v_1 remembers the value of $v \wedge w$; $c_2 = (2, B, C, v_1)$ and $c_3 = (3, B, C, w)$ will introduce two local variables v_2 and v_3 for agent C such that v_2 remembers the value of v_1 and v_3 remembers the values of w.*

Then for states s_1 and s_3 where $w = 1$, the updated indistinguishable relation for agent C is $B_C[c_1, c_2, c_3](s_1, s_3) \equiv \exists v_1, v_2, v_3 : B_C(s_1, s_3) \wedge (e(s_1, v_1) = e(s_3, v_1)) \wedge (e(s_1, w) = e(s_3, w)) \wedge (v_1 \Leftrightarrow v \wedge w) \wedge (v_2 \Leftrightarrow v_1) \wedge (v_3 \Leftrightarrow w)$, and thus we have that $\neg B_C[c_1, c_2, c_3](s_1, s_3)$, which means that agent C can distinguish state s_1 and s_3 with the messages from agent B. Therefore, after the nested communication, we have that $w \Rightarrow K_C v \vee K_C \neg v$, and therefore ϕ_9 is satisfiable.

Adding Temporary Communication

In this section, we suggest another way of handling the case of Example 5. A communication is *temporary* if the corresponding message is transmitted only at the current state. A temporary communication can be used in combination with strategy operator or knowledge operator to reasoning about agents' enhanced capability at the current state.

Instead of introducing another communication operator, we update the strategy operator and knowledge operator to allow additional communication. The new language ATLK^{tc} has the syntax as follows.

$$\phi ::= p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid EX\phi \mid EG\phi \mid E(\phi_1 U \phi_2) \mid \langle\langle G\rangle\rangle^{tc}\phi \mid E_G^{tc}\phi$$

The semantics of the updated operators is defined as follows.

1. $M, s \models \langle\langle G\rangle\rangle^{tc}\phi$ if there exist a set C of communication between agents in G such that with the additional communication, there exists a collective strategy θ_G such that for all agents $i \in G$ and all states $t \in rch(M)$ such that $B_i[C](s,t)$, there is $M_0[\theta_G], t \models \phi$.
2. $M, s \models E_G^{tc}\varphi$ if there exist a set C of communication between agents in G such that with the additional communication, for all agents $i \in Agt$ and all states $t \in rch(M_0)$ such that $B_i[C](s,t)$, there is $M, t \models \varphi$

Example 8 *We can specify the requirement of Example 5 with the following formula*

$$\phi_{10} \equiv \langle\langle\{A,B\}\rangle\rangle^{tc}(\phi_4 \wedge AX\ \langle\langle\{A,B\}\rangle\rangle^{tc}\phi_6) \quad (10)$$

By the previous analysis, the formula is satisfiable.

Unlike ATLK^{fc} in which nested communication is set up by nested operators, no nested communication can be obtained for ATLK^{tc} as the communication occurs upon request at the current state and is automatically dropped afterwards. The model checking problem is, given a multiagent system M and a formula ϕ of the language ATLK^{tc}, to decide whether $M, t \models \phi$ for all initial states $t \in I$.

The logics ATLK^{tc} and ATLK^{fc} represent two different, yet natural, ways of communication between agents. In ATLK^{fc}, the nested communication is established with several nested operators, that is, the nesting is done explicitly. The same approach cannot be applied to ATLK^{tc}, whose communication will be lost after the application of the operator. It is unnatural, and probably overly-complicated, to have a nested communication within a single operator.

Complexity

We analyse the complexity of model checking logics presented in the previous sections. The model checking problem is measured over the number of variables Var, the number of local actions $Acts$, and the number of operators in ϕ. Let $m = |Var|$, $k = |Acts|$ and $l = |\phi|$. The set Ags of agents is fixed.

The system is in *succinct representation*: the number of system states is $|S| = O(2^m)$. To be consistent with the communication in which messages contain boolean formulas, we assume that initial states $init_e, init_i$ and transition relations $\longrightarrow_e, \longrightarrow_i$ are all represented as boolean formulas such that $init_e \in \mathcal{B}(Var_e)$, $init_i \in \mathcal{B}(Var_i)$, $\longrightarrow_e \in \mathcal{B}(Var_e \cup Acts \cup Var'_e)$, and $\longrightarrow_i \in \mathcal{B}(Var_i \cup Var_e \cup Acts_i \cup Var'_i)$, where Var'_e and Var'_i are next-time variables for the environment and the agent i, respectively. The representations of a set of states and a transition relation as boolean formulas is a standard technique in symbolic model checking, see (Clarke, Grumberg, and Peled 1999) for the details. For instance, a formula $f \in \mathcal{B}(Var_e)$ represents a set of environment states as follows:

$$\{s \in L_e \mid e(s, f) = 1\}$$

Recall that an environment state is a truth assignment over the variables Var_e. Therefore, the expression represents the set of truth assignments that make the formula f true. For the transition relation, a formula $t \in \mathcal{B}(Var_e \cup Acts \cup Var'_e)$ represents a transition relation as follows:

$$\{(s, a, s') \in L_e \times JActs \times L_e \mid e(s \cup a \cup s', t) = 1\}$$

Moreover, the relation B_i is also represented as a boolean formula such that $B_i \in \mathcal{B}(Var_e \cup Var'_e)$. We also call this *symbolic representation*. Note that, formulas can be of exponential size with respect to m and k.

We have the following proposition.

Proposition 2 *Both communication and strategy can be represented as a truth table of exponential size.*

Proof: (Sketch) Given a set $Var_i \cup OVar_i$ of local variables, a formula $\varphi \in \mathcal{B}(Var_i \cup OVar_i)$ carries a single-bit information about the current local state of agent i. It can be alternatively represented as a boolean function $f_\varphi : L_i \times L_{e,i} \to \{0, 1\}$. Each local state $(s_i, s_{e,i}) \in L_i \times L_{e,i}$ is a truth assignment to the variables $Var_i \cup OVar_i$ and therefore the function f_φ can be represented as a truth table over the variables $Var_i \cup OVar_i$. The truth table has an exponential number of lines, each of which represents a mapping from a state $(s_i, s_{e,i})$ to a boolean value $f((s_i, s_{e,i}))$.

A strategy θ_i is a mapping from $L_i \times L_{e,i}$ to a nonempty subset of local actions $Acts_i$. For agent i, we introduce a set of $|Acts_i|$ boolean variables $BActs_i$, such that each local nondeterministic choice corresponds to a truth assignment to $BActs_i$. Therefore, a strategy θ_i can be represented as a truth table over the variables $Var_i \cup OVar_i \cup BActs_i$. The table has an exponential number of lines, each of which expresses that a subset of local actions are selected in a state $(s_i, s_{e,i})$. □

First of all, we have the following conclusion for ATLK by extending a result from (Huang, Chen, and Su 2015).

Theorem 1 *Model checking ATLK is NEXP-complete for a multi-agent system of succinct and symbolic representation.*

Now we can have the following two conclusions.

Theorem 2 *Model checking ATLK^{tc} is NEXP-complete.*

Proof: (Sketch) The lower bound comes from Theorem 1. For the upper bound, the algorithm returns the *reversed* result of the following procedure:

1. guesses an initial states s_0 of the model M and
2. returns the *reversed* result of $sat(M[\emptyset], s_0, \phi)$.

The function $sat(M[C], s, \phi)$ is computed inductively by the structure of the formula ϕ. For the space limit, we only give details for strategy operator.

1. $sat(M[C], s, \langle\langle G\rangle\rangle^{tc}\phi)$ if we can
   (a) guess a set of communication C' and obtain $M_0[C \cup C']$,
   (b) guess a strategy θ_G based on relation $B_i[C \cup C']$ and obtain the system $M_0[C \cup C'][\theta_G]$, and
   (c) verify $sat(M_0[C \cup C'][\theta_G], t, \phi)$ on all states t such that $t \in rch(M_0)$ and $B_i[C \cup C'](s, t)$,

   where by $M_0[C \cup C']$, we remove all strategies from the system but keep those fixed communication.

Note that, for the point (a) of the case $sat(M, s, \langle\langle G\rangle\rangle^{tc}\phi)$, we do not need to explicitly guess a truth table for C and then apply it on M_0. Because the number of communication can be exponential, this will not enable us to achieve the complexity bound. Instead, we do the following:

1. guess the resulting updated relation $B_i[C \cup C']$, and
2. check whether $B_i[C] \Rightarrow B_i[C \cup C']$ and $B_i[C \cup C'] \Rightarrow \bigvee_{j\in G} B_j[C]$.

The second step is to make sure that the relation $B_i[C \cup C']$ is an enhanced relation of $B_i[C]$ by accepting certain information from other agents in G. It can be shown that the guessing can be done in exponential time with respect to m, and the verification of $B_i \Rightarrow B_i[C]$ and $B_i[C] \Rightarrow \bigvee_{j\in G} B_j$ can also be done in exponential time with respect to m. Therefore we have the NEXP upper bound. □

Theorem 3 *Model checking ATLK^{fc} is NEXP-complete.*

Proof: (Sketch) The lower bound comes from Theorem 1. For the upper bound, the algorithm largely resembles that of Theorem 2. For the nested communication, we do not explicitly use the additional variables. Instead, we treat these additional variables as intermediate variables that are existentially quantified away when generating the new relation. See the expression $B_i[c_k](s, t)$ for such an expression. □

With the above theorems, we can conclude that the logics ATLK^{fc} and ATLK^{tc} have the same model checking complexity as that of ATLK, if the system is in succinct or symbolic representation.

Related Work

The studying of communication in a logic framework can be related to the logic of public announcements (Baltag, Moss, and Solecki 1998), which treats the effects of communication of agents in a way of updating the system by a production with another system representing a (communication) action. For such an update, one needs to fully specify the communication action and the communication leads to changes on the system. Because of the changes on the system, the semantics works with perfect recall. Our framework is more flexible. First, the details of the communication are decided by model checking algorithm and the communication does not lead to changes on the system (communication is registered in a set instead). Second, although the presented semantics works with imperfect recall, it can be adapted for perfect recall by standard techniques in temporal epistemic logic. Note that, the computational complexity can be higher for perfect recall semantics. Moreover, (Agotnes and van Ditmarsch 2008) extends the public announcement logic with strategy operator, and (de Lima 2011) considers an alternative semantics for strategy formula $\langle\langle G \rangle\rangle \phi$: there is an announcement by group G after which ϕ, where the other agents remain silent.

As for the model checking complexity relevant to the current framework, (Huang, Chen, and Su 2015) presents a set of complexity results for model checking concurrent and symbolic representations of multiagent systems.

Conclusions and Future Work

The paper presents our first step towards considering the impact of communication on agents' collaboration and coordination. It is somewhat surprising that, the additional communication between agents does not incur increased complexity for the important model checking problem.

As the next step, we may consider incorporating communication into e.g, more expressive strategy logics such as (Huang and van der Meyden 2014c; 2014a) or logics with probabilistic knowledge or strategy operators such as (Huang, Su, and Zhang 2012; Huang and Luo 2013; Huang, Luo, and van der Meyden 2011), etc. We are also interested in extending the current symbolic model checking algorithm (Huang and van der Meyden 2014b) to work with communication operators.

Acknowledgement This paper is supported by National Natural Science Foundation of China (No.61572234 and No.61272415) and Fundamental Research Funds for the Central Universities of China (No.21615441).

References

Agotnes, T., and van Ditmarsch, H. 2008. Coalitions and announcements. In *AAMAS 2008*, 673–680.

Alur, R.; Henzinger, T. A.; and Kupferman, O. 2002. Alternating-Time Temporal Logic. *Journal of the ACM* 49(5):672–713.

Baltag, A.; Moss, L. S.; and Solecki, S. 1998. The logic of public announcements, common knowledge, and private suspicions. In *TARK 1998*, 43–56.

Chatterjee, K.; Henzinger, T. A.; and Piterman, N. 2010. Strategy logic. *Inf. Comput.* 208(6):677–693.

Clarke, E.; Grumberg, O.; and Peled, D. 1999. *Model Checking*. The MIT Press.

de Lima, T. 2011. Alternating-time temporal announcement logic. In *CLIMA XII*, 105–121.

Diaconu, R., and Dima, C. 2012. Model-checking alternating-time temporal logic with strategies based on common knowledge is undecidable. *Applied Artificial Intelligence* 26(4):331–348.

Fagin, R.; Halpern, J.; Moses, Y.; and Vardi, M. 1995. *Reasoning About Knowledge*. The MIT Press.

Guelev, D. P.; Dima, C.; and Enea, C. 2011. An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. *Journal of Applied Non-Classical Logics* 21(1):93–131.

Horty, J. F. 2001. *Agency and deontic logic*. Oxford University Press.

Huang, X., and Luo, C. 2013. A Logic of Probabilistic Knowledge and Strategy. In *AAMAS 2013*, 845–852.

Huang, X., and van der Meyden, R. 2014a. An epistemic strategy logic. In *SR 2014*, 35–41.

Huang, X., and van der Meyden, R. 2014b. Symbolic model checking epistemic strategy logic. In *AAAI 2014*, 1426–1432.

Huang, X., and van der Meyden, R. 2014c. A temporal logic of strategic knowledge. In *the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR2014)*.

Huang, X.; Chen, Q.; and Su, K. 2015. The complexity of model checking succinct multiagent systems. In *IJCAI 2015*, 1076–1082.

Huang, X.; Luo, C.; and van der Meyden, R. 2011. Symbolic Model Checking of Probabilistic Knowledge. In *13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK XII)*, 177–186.

Huang, X.; Su, K.; and Zhang, C. 2012. Probabilistic Alternating-time Temporal Logic of Incomplete information and Synchronous Perfect Recall. In *AAAI-12*, 765–771.

Jamroga, W., and Ågotnes, T. 2007. Constructive knowledge: what agents can achieve under imperfect information. *Journal of Applied Non-Classical Logics* 17(4):423–475.

Mogavero, F.; Murano, A.; and Vardi, M. Y. 2010. Reasoning about strategies. In *FSTTCS*, 133–144.

Pauly, M. 2002. A modal logic for coalitional power in games. *Journal of Logic and Computation* 12(1):149–166.

Schobbens, P.-Y. 2004. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science* 85(2):82–93.

van der Hoek, W., and Wooldridge, M. 2002. Tractable multiagent planning for epistemic goals. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'02)*, 1167–1174.

van Otterloo, S., and Jonker, G. 2005. On Epistemic Temporal Strategic Logic. *Electronic Notes in Theoretical Computer Science (ENTCS)* 126:77–92.

Yao, A. C.-C. 1982. Protocols for secure computations (extended abstract). In *FOCS1982*, 160–164.