Reasoning About Common Knowledge with Infinitely Many Agents

Joseph Y. Halpern*
Computer Science Department
Cornell University
halpern@cs.cornell.edu

Richard A. Shore§
Mathematics Department
Cornell University
shore@math.cornell.edu

Abstract

Complete axiomatizations and exponential-time decision procedures are provided for reasoning about knowledge and common knowledge when there are infinitely many agents. The results show that reasoning about knowledge and common knowledge with infinitely many agents is no harder than when there are finitely many agents, provided that we can check the cardinality of certain set differences $G - G'$, where G and G' are sets of agents. Since our complexity results are independent of the cardinality of the sets G involved, they represent improvements over the previous results even with the sets of agents involved are finite. Moreover, our results make clear the extent to which issues of complexity and completeness depend on how the sets of agents involved are represented.

1 Introduction

Reasoning about knowledge and common knowledge has been shown to be widely applicable in distributed computing, AI, and game theory. (See [2] for numerous examples.) Complete axioms for reasoning about knowledge and common knowledge are well known in the case of a fixed finite set of agents. However, in many applications, the set of agents is not known in advance and has no *a priori* upper bound (think of software agents on the web or nodes on the Internet, for example); it is often easiest to model the set of agents as an infinite set. And in some applications in game theory and economics (where reasoning about knowledge and common knowledge is quite standard; see, for example, [1, 4]) economies are often modeled as consisting of infinitely many (even uncountably many) agents. For example, when analyzing a game played with two teams, we may well want to say that everyone on team 1 knows that everyone on team 2 knows some fact p, or that it is common knowledge among the agents on team 1 that p is common knowledge among the agents on team 2. We would want to say this even if the teams consist of infinitely many agents.

The logics for reasoning about the knowledge of groups of agents contain modal operators K_i (where $K_i\varphi$ is read "agent i knows φ"), E_G (where $E_G\varphi$ is read "everyone in group G knows φ"), and C_G (where $C_G\varphi$ is read "φ is common knowledge among group G"). The operators E_G and C_G make perfect sense even if we allow the sets G to be infinite—their semantic definitions remain unchanged. If the set of agents is finite, so that, in particular, G is finite, there is a simple axiom connecting $E_G\varphi$ to $K_i\varphi$, namely, $E_G\varphi \Leftrightarrow \wedge_{i\in G} K_i\varphi$. Once we allow infinite groups G of agents, there is no obvious analogue for this axiom. Nevertheless, in this paper, we show that there exist natural sound and complete axiomatizations for reasoning about knowledge and common knowledge even if there are infinitely many agents.

It is also well known that if there are finitely many agents, then there is a decision procedure that decides if a formula φ is satisfiable (or valid) that runs in time exponential in $|\varphi|$, where φ is the length of the formula viewed as a string of symbols. We prove a similar result for a language with infinitely many agents. However, two issues arise (that, in fact, are also relevant even if there are only finitely many agents, although they have not been considered before):

1. In the statement of the complexity result in [2], E_G and C_G are both viewed as having length $2 + 2|G|$ (where $|G|$ is the cardinality of G). Clearly we cannot use this definition here if we want to get interesting complexity results, since $|G|$ may be infinite. Even if we restrict our attention to finite sets G, we would like a decision procedure that treats these sets in a uniform way, independently of their cardinality. Here we view E_G as having length 1 and C_G as having length 3, independent of the cardinality of G. (See, for example, the proof of Proposition 3.5 for the role of independence and the definition of $Sub(\varphi)$ in the proof of Theorem 4.5 for an indication as to why C_G has length 3

*Supported in part by NSF under grant IRI-96-25901.
§Supported in part by NSF Grant DMS-9802843 and DOD MURI Grant SA1515PG.

rather than 1.) Even with this definition of length, we prove that the complexity of the satisfiability problem is still essentially exponential time. (We discuss below what "essentially" means.) Thus our results improve previously-known results even if there are only finitely many agents.

2. In the earlier proofs, it is implicitly presumed that the sets G are presented in such a way that there is no difficulty in testing membership in G. As we show here, in order to decide if certain formulas are satisfiable, we need to be able to test if certain subsets of agents of the form $G_0 - (G_1 \cup \ldots \cup G_k)$ are empty, where $G_0, \ldots, G_k$ are sets of agents. In fact, if we are interested in a notion of knowledge that satisfies *positive introspection*—that is, if agent i knows φ, then she knows that she knows it—then we also need to check whether such subsets are singletons. And if we are interested in a notion of knowledge that satisfies *negative introspection*—that is, if agent i does not know φ, then she knows that she does not know it—then we may also need to check whether such subsets have cardinality m, for certain finite m. The difficulty of deciding these questions depends in part on how $G_0, \ldots, G_k$ are presented and which sets of agents we can talk about in the language. For example, if $G_0, \ldots, G_k$ are recursive sets, deciding if $G_0 - (G_1 \cup \ldots \cup G_k)$ is nonempty may not even be recursive. Here, we provide a decision procedure for satisfiability that runs in time exponential in $|\varphi|$ provided that we have oracles for testing appropriate properties of sets of the form $G_0 - (G_1 \cup \ldots \cup G_k)$. Moreover, we show that any decision procedure must be able to answer the questions we ask. In the full paper, we prove a stronger result, providing a tight bound on the complexity of deciding satisfiability that takes into account the complexity of answering questions about the cardinality of $G_0 - (G_1 \cup \ldots \cup G_k)$.

Again, this issue is of significance even if there are only finitely many agents. For example, in the SDSI approach to security [8], there are names, which can be viewed as representing sets of agents. SDSI provides a (nondeterministic) algorithm for computing the set of agents represented by a name. If we want to make statements such as "every agent represented by name $\mathtt{n}$ knows φ" (statements that we believe will be useful in reasoning about security [5, 6]) then the results of this paper show that to decide validity in the resulting logic, we need more than just an algorithm for resolving the agents represented by a given name. We also need algorithms for resolving which agents are represented by one name and not another. More generally, if we assume that we have a separate language for representing sets of agents, our results characterize the properties of sets that we need to be able to decide in order to reason about the group knowledge of these agents.

In the next section, we briefly review the syntax and semantics of the logic of common knowledge. In Section 3 we state the main results and prove them under some simplifying assumptions that allow us to bring out the main ideas of the proof. We drop these assumptions in Section 4, where we provide some of the details of the proofs of the main results (the remainder can be found in the full paper).

2 Syntax and Semantics: A Brief Review

Syntax: We start with a (possibly infinite) set $\mathcal{A}$ of agents. Let $\mathcal{G}$ be a set of nonempty subsets of $\mathcal{A}$. (Note that we do not require $\mathcal{G}$ to be closed under union, intersection, or complementation; it can be an arbitrary collection of subsets.) We get the language $\mathcal{L}_{\mathcal{G}}^{C}(\Phi)$ by starting with a set Φ of primitive propositions, and closing under $\wedge$, $\neg$, and the modal operators K_i, for $i \in \mathcal{A}$, and E_G, C_G, for $G \in \mathcal{G}$. Thus, if $p, q \in \Phi$, $i \in \mathcal{A}$, and $G, G' \in \mathcal{G}$, then $K_i C_G(p \wedge E_{G'} q) \in \mathcal{L}_{\mathcal{G}}^{C}(\Phi)$. We read $K_i\varphi$ as "agent i knows φ", $E_G\varphi$ as "everyone in group G knows φ", and $C_G\varphi$ as "φ is common knowledge among the agents in G". Let $\mathcal{L}_{\mathcal{G}}^{E}$ be the sublanguage of $\mathcal{L}_{\mathcal{G}}^{C}$ that does not include the C_G operators. Let $|\varphi|$ be the length of the formula viewed as a string of symbols, where the modal operators K_i and E_G are counted as having length 1 and C_G is counted as having length 3 (even if G is an infinite set of agents) and all primitive propositions are counted as having length 1.

In [2, 7], $\mathcal{A}$ is taken to be the set $\{1, \ldots, n\}$; in [7], $\mathcal{G}$ is taken to be the singleton $\{\{1, \ldots, n\}\}$ (so that we can only talk about every agent in $\mathcal{A}$ knowing φ and common knowledge among the agents in $\mathcal{A}$), while in [2], $\mathcal{G}$ is taken to consist of all nonempty subsets of $\mathcal{A}$.

Semantics: As usual, formulas in $\mathcal{L}_{\mathcal{G}}^{C}$ are either true or false at a world in a Kripke structure. Formally, a Kripke structure M over $\mathcal{A}$ and Φ is a tuple $(S, \pi, \{\mathcal{K}_i : i \in \mathcal{A}\})$, where S is a set of states or possible worlds, π associates with each state in S a truth assignment to the primitive propositions in Φ (so that $\pi(s) : \Phi \rightarrow \{\mathbf{true}, \mathbf{false}\}$), and $\mathcal{K}_i$ is a binary relation on S for each agent $i \in \mathcal{A}$. We occasionally write $\mathcal{K}_i(s)$ for $\{t : (s,t) \in \mathcal{K}_i\}$.

We define the truth relation $\models$ as follows:

$(M,s) \models p$ (for $p \in \Phi$) iff $\pi(s)(p) = \mathbf{true}$

$(M,s) \models \varphi \wedge \psi$ iff both $(M,s) \models \varphi$ and $(M,s) \models \psi$

$(M,s) \models \neg\varphi$ iff $(M,s) \not\models \varphi$

$(M,s) \models K_i\varphi$ iff $(M,t) \models \varphi$ for all $t \in \mathcal{K}_i(s)$

$(M,s) \models E_G\varphi$ iff $(M,s) \models K_i\varphi$ for all $i \in G$

$(M,s) \models C_G\varphi$ iff $(M,s) \models E_G^k\varphi$ for $k = 1, 2, 3, \ldots$, where E_G^k is defined inductively by taking $E_G^1\varphi =_{\mathrm{def}} E_G\varphi$ and $E_G^{k+1}\varphi =_{\mathrm{def}} E_G E_G^k\varphi$.

We say that t is *G-reachable from s in M* if there exist $s_0, \ldots, s_k$ with $s = s_0$, $t = s_k$, and $(s_i, s_{i+1}) \in \cup_{i \in G} \mathcal{K}_i$. For later use, we extend this definition so that if $S' \subseteq S$, we say that t is *G-reachable from s in S'* if $s_0, \ldots, s_k \in S'$. The following characterization of common knowledge is well known [2].

Lemma 2.1: *$(M, s) \models C_G \varphi$ iff $(M, t) \models \varphi$ for all t that are G-reachable from s in M.*

Let $\mathcal{M}_{\mathcal{A}}(\Phi)$ be the class of all Kripke structures over $\mathcal{A}$ and Φ (with no restrictions on the $\mathcal{K}_i$ relations). We are also interested in various subclasses of $\mathcal{M}_{\mathcal{A}}(\Phi)$, obtained by restricting the $\mathcal{K}_i$ relations. In particular, we consider $\mathcal{M}^r_{\mathcal{A}}(\Phi)$, $\mathcal{M}^{rt}_{\mathcal{A}}(\Phi)$, $\mathcal{M}^{rst}_{\mathcal{A}}(\Phi)$, and $\mathcal{M}^{elt}_{\mathcal{A}}(\Phi)$, the class of all structures over $\mathcal{A}$ and Φ where the $\mathcal{K}_i$ relations are reflexive (resp., reflexive and transitive; reflexive, symmetric, and transitive; Euclidean, serial, and transitive).[1] For the remainder of this paper, we take Φ to be fixed, and do not mention it, writing, for example $\mathcal{L}^C_{\mathcal{G}}$ and $\mathcal{M}_{\mathcal{A}}$ rather than $\mathcal{L}^C_{\mathcal{G}}(\Phi)$ and $\mathcal{M}_{\mathcal{A}}(\Phi)$.

Axioms: The following are the standard axioms and rules that have been considered for knowledge.

Prop. All substitution instances of tautologies of propositional calculus.

K1. $(K_i \varphi \wedge K_i(\varphi \Rightarrow \psi)) \Rightarrow K_i \psi, i \in \mathcal{A}$.

K2. $K_i \varphi \Rightarrow \varphi,\ i \in \mathcal{A}$.

K3. $\neg K_i \mathit{false},\ i \in \mathcal{A}$.

K4. $K_i \varphi \Rightarrow K_i K_i \varphi,\ i \in \mathcal{A}$.

K5. $\neg K_i \varphi \Rightarrow K_i \neg K_i \varphi,\ i \in \mathcal{A}$.

MP. From φ and $\varphi \Rightarrow \psi$ infer ψ.

KGen. From φ infer $K_i \varphi$, $i \in \mathcal{A}$.

Technically, Prop and K1–K5 are axiom schemes, rather than single axioms. K1, for example, holds for all formulas φ and ψ. A formula such as $K_1 q \vee \neg K_1 q$ is an instance of axiom Prop (since it is a substitution instance of the propositional tautology $p \vee \neg p$, obtained by substituting $K_1 q$ for p).

We will be interested in the following axioms and rule for reasoning about everyone knows:

E1. $E_G \varphi \Rightarrow K_i \varphi$ if $i \in G$, $G \in \mathcal{G}$.

E2. $(\wedge_{i \in \mathcal{A}'} K_i \varphi \wedge \wedge_{G' \in \mathcal{G}'} E_{G'} \varphi) \Rightarrow E_G \varphi$ if $\mathcal{A}'$ is a finite subset of $\mathcal{A}$, $\mathcal{G}'$ is a finite subset of $\mathcal{G}$, and $G \subseteq (\mathcal{A}' \cup (\cup \mathcal{G}'))$.

[1] Recall that a relation R is Euclidean if $(s,t), (s,u) \in R$ implies that $(t,u) \in R$.

E3. $(E_G \varphi \wedge E_G(\varphi \Rightarrow \psi)) \Rightarrow E_G \psi$.

E4. $E_G(E_G \varphi \Rightarrow \varphi)$.

E5. $E_G \varphi \Rightarrow \varphi$.

E6. $\neg \varphi \Rightarrow E_G \neg E_G \varphi$.

EGen. From φ infer $E_G \varphi$.

If $\mathcal{A}$ is finite (so that all the sets in $\mathcal{G}$ are finite) we can simplify E1 and E2 to

E. $E_G \varphi \Leftrightarrow \wedge_{i \in G} K_i \varphi$.

It is easy to see that E follows from E1 and E2 (in the presence of Prop and MP) and every instance of E1 and E2 follows from E if $\mathcal{A}$ is finite. E is used instead of E1 and E2 in [2, 7]. Note that E2 is recursive iff deciding if $G - (\mathcal{A}' \cup (\cup \mathcal{G}')) = \emptyset$ is recursive. (We determine precisely which such questions we must be able to answer in Proposition 3.3.)

E3 and EGen are the obvious analogues of K1 and KGen for E_G. We do not need them in the case that $\mathcal{A}$ is finite; it is easy to see that they follow from K1, KGen, and E. In the case that $\mathcal{A}$ is infinite, however, they are necessary.

Axiom E4 is sound in $\mathcal{M}^r_{\mathcal{A}}$, $\mathcal{M}^{rt}_{\mathcal{A}}$, $\mathcal{M}^{rst}_{\mathcal{A}}$, and $\mathcal{M}^{elt}_{\mathcal{A}}$. It is easy to see that E4 follows from K2, E1, and EGen, so will not be needed in systems that contain these axioms. Moreover, it is not hard to show that E4 follows from E1, E2, and K5 if the set of agents is finite. However, it does not follow from these axioms if the set of agents is infinite.

Axiom E5 follows from K2 and E1. Moreover, we use it only in systems that already include K2 and E1. Nevertheless, for technical reasons, it is useful to list it separately.

Axiom E6 (with E_G replaced by K_i) is the standard axiom used to characterize symmetric $\mathcal{K}_i$ relations [2]. It follows easily from K2, K5, E1, and E2 if $\mathcal{A}$ is finite. However, like E4, it must be specifically included if $\mathcal{A}$ is infinite.

Finally, we have the following well-known axiom and inference rule for common knowledge:

C1. $C_G \varphi \Rightarrow E_G(\varphi \wedge C_G \varphi)$.

RC1. From $\varphi \Rightarrow E_G(\psi \wedge \varphi)$ infer $\varphi \Rightarrow C_G \psi$.

Historically, in the case of one agent, the system with axioms and rules Prop, K1, MP, and KGen has been called K; adding K2 to K gives us T; adding K4 to T gives us S4; adding K5 to S4 gives us S5; replacing K2 by K3 in S5 gives us KD45. We use the subscript $\mathcal{G}$ to emphasize the fact that we are considering systems with sets of agents coming from $\mathcal{G}$ rather than only one agent and the superscript C to emphasize that we add E1–E3, EGen, C1, and RC1 to the system. In this way, we get the systems $\mathrm{K}^C_{\mathcal{G}}$, $\mathrm{T}^C_{\mathcal{G}}$, and $\mathrm{S4}^C_{\mathcal{G}}$. Thus, $\mathrm{K}^C_{\mathcal{G}}$ consists of Prop, K1, MP, KGen, E1, E2, E3, EGen, C1, and RC1; we get $\mathrm{S4}_{\mathcal{G}}^{C}$ by adding K2 and K4 to $\mathrm{K}_{\mathcal{G}}^{C}$. We get $\mathrm{KD45}_{\mathcal{G}}^{C}$ by adding K3–K5 and E4 to $\mathrm{K}_{\mathcal{G}}^{C}$ and we get $\mathrm{S5}_{\mathcal{G}}^{C}$ by adding K2, K4, K5 and E6 to $\mathrm{K}_{\mathcal{G}}^{C}$.

One of the two main results of this paper shows that each of these axiom systems is sound and complete with respect to an appropriate class of structures. For example, $\mathrm{K}_{\mathcal{G}}^{C}$ is a sound and complete axiomatization with respect to $\mathcal{M}_{\mathcal{A}}$ and $\mathrm{S5}_{\mathcal{G}}^{C}$ is a sound and complete axiomatization with respect to $\mathcal{M}_{\mathcal{A}}^{rst}$. In the case that $\mathcal{A}$ is finite, this result is well known (see [2, 7]—as mentioned earlier, E is used in the axiomatization instead of E1–E3 and EGen). What is perhaps surprising is that E1–E3 and EGen suffice even if $\mathcal{A}$ is infinite. For example, suppose $\mathcal{G}$ just consists of the singleton $\mathcal{A}$. In that case, E2 becomes vacuous. Thus, while the axioms force $E_{\mathcal{A}}\varphi$ to imply that each agent in $\mathcal{A}$ knows φ, we have no way of expressing the converse. Indeed, as we show in the full paper, it is easy to construct a structure for the axioms with the standard interpretations of all the K_i relations but a nonstandard one of $E_{\mathcal{A}}$, where all the agents in $\mathcal{A}$ know φ and yet $E_{\mathcal{A}}\varphi$ does not hold.

3 The Main Results and a Proof in a Simplified Setting

In this section, we state the two main results of this paper—complete axiomatizations and decision procedures. We then provide a proof of a simpler version of these results that illustrates some of the main ideas. We first state the completeness results.

Theorem 3.1: *For formulas in the language $\mathcal{L}_{\mathcal{G}}^{C}$:*

(a) $\mathrm{K}_{\mathcal{G}}^{C}$ is a sound and complete axiomatization with respect to $\mathcal{M}_{\mathcal{A}}$,

(b) $\mathrm{T}_{\mathcal{G}}^{C}$ is a sound and complete axiomatization with respect to $\mathcal{M}_{\mathcal{A}}^{r}$,

(c) $\mathrm{S4}_{\mathcal{G}}^{C}$ is a sound and complete axiomatization with respect to $\mathcal{M}_{\mathcal{A}}^{rt}$,

(d) $\mathrm{S5}_{\mathcal{G}}^{C}$ is a sound and complete axiomatization with respect to $\mathcal{M}_{\mathcal{A}}^{rst}$,

(e) $\mathrm{KD45}_{\mathcal{G}}^{C}$ is a sound and complete axiomatization with respect to $\mathcal{M}_{\mathcal{A}}^{elt}$.

Before stating the results regarding complexity, we first show that questions about certain facts regarding sets of the form $G_0 - (G_1 \cup \ldots \cup G_k)$ are reducible to satisfiability.

Definition 3.2: Given a set $\mathcal{J}$ of subsets of $\mathcal{A}$ and an integer $m \geq 1$, define a sequence $\mathcal{J}_0^m, \mathcal{J}_1^m, \ldots$ of sets of subsets of $\mathcal{A}$ inductively as follows. Let $\mathcal{J}_0^m = \mathcal{J}$. Suppose we have defined $\mathcal{J}_0^m, \ldots, \mathcal{J}_k^m$. Then $\mathcal{J}_{k+1}^m = \mathcal{J} \cup \{G - \cup\mathcal{H} : \mathcal{H} \subseteq \mathcal{J}_k^m, \mathcal{H}$ finite$, |G - \cup\mathcal{H}| \leq m\}$. Let $\mathcal{J}^m = \cup_i \mathcal{J}_i^m$; let $\widehat{\mathcal{J}}^m = \{G - \cup\mathcal{H} : G \in \mathcal{J}, \mathcal{H} \subseteq \mathcal{J}^m, \mathcal{H}$ finite$\}$. For uniformity, we take $\widehat{\mathcal{J}}^0 = \{G - \cup\mathcal{H} : G \in \mathcal{J}, \mathcal{H} \subseteq \mathcal{J}, \mathcal{H}$ finite$\}$. ∎

Let $\mathcal{J}^*$ be the algebra generated by $\mathcal{J}$ (that is, the Boolean combinations of sets in $\mathcal{J}$). It is useful to talk about the length of a description of various sets in $\mathcal{J}^*$ (particularly those in $\widehat{\mathcal{J}}^m$ for some m). Formally, we assume we have a language whose primitive objects consist of the elements of $\mathcal{J}$ and the symbols $\cup$ and $-$ (for set difference). The length of a description is then the number of symbols of $\mathcal{J}$ that appear in it. Notice that, in general, an element of $\mathcal{J}^*$ may have several different descriptions. We are not always careful to distinguish a set from its description. (We hope that the reader will be able to tell which is intended from context.) We use $l(G)$ to denote the length of the description of $G \in \mathcal{J}^*$.

Let $\mathcal{G}_{\mathcal{A}} = \mathcal{G} \cup \{\{i\} : i \in \mathcal{A}\}$. Throughout the paper (and, in particular, in the proof of the next proposition), for ease of exposition, we identify $E_{\{i\}}$ with K_i, for $i \in \mathcal{A}$ (which allows us to write E_G for each $G \in \mathcal{G}_{\mathcal{A}}$).

Proposition 3.3:

(a) The question of whether $|G| > 0$ for $G \in \widehat{\mathcal{G}}_{\mathcal{A}}^0$ is reducible (in time linear in $l(G)$) to the satisfiability problem for the language $\mathcal{L}_{\mathcal{G}}^{E}$ with respect to all of $\mathcal{M}_{\mathcal{A}}$, $\mathcal{M}_{\mathcal{A}}^{r}$, $\mathcal{M}_{\mathcal{A}}^{rt}$, $\mathcal{M}_{\mathcal{A}}^{rst}$, and $\mathcal{M}_{\mathcal{A}}^{elt}$.

(b) The questions of whether $|G| > 0$ and $|G| > 1$ for $G \in \widehat{\mathcal{G}}_{\mathcal{A}}^1$ are reducible (in time linear in $l(G)$) to the satisfiability problem for the language $\mathcal{L}_{\mathcal{G}}^{E}$ with respect to all of $\mathcal{M}_{\mathcal{A}}^{rt}$, $\mathcal{M}_{\mathcal{A}}^{rst}$, and $\mathcal{M}_{\mathcal{A}}^{elt}$.

(c) For all $m \geq 1$, the question of whether $|G| > m$ for $G \in \widehat{\mathcal{G}}_{\mathcal{A}}^m$ is reducible (in time linear in $l(G) + m$) to the satisfiability problem for $\mathcal{L}_{\mathcal{G}}^{E}$ with respect to $\mathcal{M}_{\mathcal{A}}^{rst}$ and $\mathcal{M}_{\mathcal{A}}^{elt}$.

(d) The question of whether $|G_1 \cap \ldots \cap G_k| > 0$, for $G_1, \ldots, G_k \in \mathcal{G}_{\mathcal{A}}$ is reducible (in time linear in k) to the satisfiability problem for $\mathcal{L}_{\mathcal{G}}^{E}$ with respect to $\mathcal{M}_{\mathcal{A}}^{elt}$.

Proof: For part (a), suppose $G \in \widehat{\mathcal{G}}_{\mathcal{A}}^0$. Thus, $G = G_0 - (G_1 \cup \ldots \cup G_k)$, for some $G_0, \ldots, G_k \in \mathcal{G}_{\mathcal{A}}$. Consider the formula $\varphi_a =_{\mathrm{def}} \neg E_{G_0} p \wedge E_{G_1} p \wedge \ldots \wedge E_{G_k} p$, where p is a primitive proposition. Clearly φ_a is satisfiable in $\mathcal{M}_{\mathcal{A}}$, $\mathcal{M}_{\mathcal{A}}^{r}$, $\mathcal{M}_{\mathcal{A}}^{rt}$, $\mathcal{M}_{\mathcal{A}}^{rst}$, or $\mathcal{M}_{\mathcal{A}}^{elt}$ iff $|G_0 - (G_1 \cup \ldots \cup G_k)| > 0$.

For part (b), given G, we construct two formulas $\varphi_{G,p}$ and ψ_G with the following properties.

- $\varphi_{G,p}$ is satisfiable iff $|G| > 0$.
- If $(M, s) \models \varphi_{G,p}$, then $(M, s) \models \neg K_j p$ for some $j \in G$.

- ψ_G is satisfiable iff $|G| > 1$.
- $|\varphi_{G,p}|$ and $|\psi_G|$ are both linear in $l(G)$.

This, of course, suffices to prove the result.

We construct the formulas $\varphi_{G,p}$ by induction on the least h such that $G = G' - \cup\mathcal{H}$ and $\mathcal{H} \subseteq (\mathcal{G}_\mathcal{A})^1_h$. (We are here thinking of G as specified by its description.) If $\mathcal{H} \subseteq (\mathcal{G}_\mathcal{A})^1_0 = \mathcal{G}_A$, suppose $\mathcal{H} = \{G_1, \ldots, G_k\}$. Then we take $\varphi_{G,p}$ to be $\neg E_{G'}p \wedge E_{G_1}p \wedge \ldots \wedge E_{G_k}p$. This clearly has the desired properties.

Now suppose $\mathcal{H} \subseteq (\mathcal{G}_\mathcal{A})^1_h$ for $h \geq 1$. Without loss of generality, we can assume that $\mathcal{H} = \{G_1, \ldots, G_{k'}, G_{k'+1}, \ldots, G_k\}$, where $G_1, \ldots, G_{k'} \in \mathcal{G}_\mathcal{A}$ and, for $j = k'+1, \ldots, k$, $G_j \in (\mathcal{G}_\mathcal{A})^1_h - \mathcal{G}_\mathcal{A}$ is of the form $G'_j - \cup\mathcal{H}_j$ with $G'_j \in \mathcal{G}_\mathcal{A}$, $\mathcal{H}_j \subseteq (\mathcal{G}_\mathcal{A})^1_{h-1}$, and $|G_j| = 1$. Define $\varphi_{G,p}$ as

$$\neg E_{G'}\neg(\neg p \wedge \bigwedge_{j=k'+1}^{k} \varphi_{G_j,p_j}) \wedge E_{G_1}p \wedge \ldots \wedge E_{G_{k'}}p \wedge \bigwedge_{j=k'+1}^{k} E_{G'_j}p_j,$$

where we assume that the sets of primitive propositions that appear in φ_{G_j,p_j}, $j = k'+1, \ldots, k$, are mutually exclusive and do not include p.

Finally, define ψ_G to be $\varphi_{G,p} \wedge E_{G'}(q \wedge (\neg p \Rightarrow \varphi_{G,q}))$, where we assume that the primitive propositions that appear in $\varphi_{G,p}$ and $\varphi_{G,q}$ are disjoint. In the full paper, we show that $\varphi_{G,p}$ and ψ_G have the required properties. This clearly suffices to prove (b).

The proof of part (c) is similar in spirit to that of (b); we leave the details to the full paper.

Finally, for part (d), consider the formula φ_d defined as

$$E_{G_1}p_1 \wedge \ldots \wedge E_{G_{k-1}}p_{k-1} \wedge E_{G_k}(\neg p_1 \vee \ldots \vee \neg p_{k-1}).$$

We leave it to the reader to check that φ_d is satisfiable in $\mathcal{M}^{elt}_\mathcal{A}$ iff $G_1 \cap \ldots \cap G_k = \emptyset$. ∎

We already saw that for axiom E2 to be recursive, we need to be able to decide whether $|G_0 - (G_1 \cup \ldots \cup G_k)| \geq 1$ (or, equivalently, whether $G_0 \subseteq G_1 \cup \ldots \cup G_k$) for $G_0, \ldots, G_k \in \mathcal{G}_\mathcal{A}$. Proposition 3.3 shows that if there is no recursive algorithm for answering such questions, the satisfiability problem for the logic (even without C_G operators) is also not decidable. For simplicity here, we assume we have oracles that can answer the questions that we need to answer (according to Proposition 3.3) in unit time; we consider the complexity of querying the oracle in more detail in the full paper. More precisely, let O_m be an oracle that, for a set $G \in \widehat{\mathcal{G}}^m_\mathcal{A}$, tells us whether $|G| > k$, for any $k < m$. (Thus, queries to oracle O_m have the form (G, k).) Let O' be an oracle that tells us whether $G_1 \cap \ldots \cap G_k = \emptyset$, for $G_1, \ldots, G_k \in \mathcal{G}_\mathcal{A}$.

Theorem 3.4: *There is a constant $c > 0$ and an algorithm that decides if a formula $\varphi \in \mathcal{L}^C_\mathcal{G}$ is satisfiable in a structure in $\mathcal{M}_\mathcal{A}$ (resp., $\mathcal{M}^r_\mathcal{A}$, $\mathcal{M}^{rt}_\mathcal{A}$, $\mathcal{M}^{rst}_\mathcal{A}$, $\mathcal{M}^{elt}_\mathcal{A}$) that runs in time $2^{c|\varphi|}$ for some constant $c > 0$ given oracle O_0 (resp., O_0, O_1, $O_{|\varphi|}$, both $O_{|\varphi|}$ and O'), where queries to the oracle take unit time. Moreover, if $\mathcal{G}$ contains a subset with at least two elements, then there exists a constant $d > 0$ such that every algorithm for deciding the satisfiability of formulas in $\mathcal{M}_\mathcal{A}$ (resp., $\mathcal{M}^r_\mathcal{A}$, $\mathcal{M}^{rt}_\mathcal{A}$, $\mathcal{M}^{rst}_\mathcal{A}$, $\mathcal{M}^{elt}_\mathcal{A}$) runs in time at least $2^{d|\varphi|}$, even given access to oracle O_0 (resp., O_0, O_1, $O_{|\varphi|}$, both $O_{|\varphi|}$ and O'), for infinitely many formulas φ.*

Before proving Theorems 3.1 and 3.4, we prove a somewhat simpler theorem that allows us to both explain intuitively why the results are true and point out some of the difficulties in proving them.

Proposition 3.5: *If there is an oracle that decides whether $G = \emptyset$ for each Boolean combination G of elements in $\mathcal{G}_\mathcal{A}$, then, for every formula $\varphi \in \mathcal{L}^C_\mathcal{G}$, we can effectively find a formula φ^σ in a language $\mathcal{L}^C_{\mathcal{G}'}$, where $\mathcal{G}'$ consists of all subsets of a set $\mathcal{A}'$ of at most $2^{|\varphi|}$ agents, such that $|\varphi^\sigma| = |\varphi|$ and φ is satisfiable in $\mathcal{M}_\mathcal{A}$ iff φ^σ is satisfiable in $\mathcal{M}_{\mathcal{A}'}$.*

Proof: Given φ, let $\mathcal{G}_\varphi$ be the set of subsets G of agents such that E_G or C_G appears in φ. (Recall that we are identifying K_i with $E_{\{i\}}$.) Note that $|\mathcal{G}_\varphi| \leq |\varphi|$.

Suppose $\mathcal{G} = \{G_1, \ldots, G_N\}$. An *atom over* $\mathcal{G}$ is a nonempty set of the form $G'_1 \cap \ldots \cap G'_N$, where $G'_i = G_i$ or $G'_i = \overline{G_i}$. Clearly there are at most 2^N atoms over $\mathcal{G}$. Let $\mathcal{A}'$ consist of the nonempty atoms over $\mathcal{G}_\varphi$. Note that $|\mathcal{A}'| \leq 2^{|\varphi|}$. Define $\sigma : \mathcal{A} \to \mathcal{A}'$ by taking $\sigma(i)$ to be the unique atom over $\mathcal{G}_\varphi$ containing i. We extend σ to a map from $2^\mathcal{A} \to 2^{\mathcal{A}'}$ by taking $\sigma(G) = \{\sigma(i) : i \in G\}$. Translate φ to φ^σ by replacing all occurrences of E_G and C_G in φ by $E_{\sigma(G)}$, and $C_{\sigma(G)}$, respectively. Clearly $|\varphi| = |\varphi^\sigma|$. (Note that it is important here that we take the length of E_G and C_G to be independent of G.)

If φ is satisfiable let (M, s) witness that fact. Convert M into a structure M^σ over $\mathcal{A}'$ with the same state space by setting $(s, t) \in \mathcal{K}_A$ iff $(s, t) \in \cup_{j \in A}\mathcal{K}_j$ for each $A \in \mathcal{A}'$. An easy induction shows that for every formula ψ with sets (of agents) chosen from $\mathcal{G}_\varphi$, we have $(M, s) \models \psi$ if and only if $(M^\sigma, s) \models \psi^\sigma$. The only point that needs any comment is that E_G (and so also C_G) has the same meaning in M (in terms of reachability) as $E_{\sigma(G)}$ ($C_{\sigma(G)}$) in M^σ, by the definition of $\sigma(G)$ and the $\mathcal{K}_A$ relations. Thus $(M^\sigma, s) \models \varphi^\sigma$ as required.

For the other direction, suppose $(M', s) \models \varphi^\sigma$ for some structure M' over $\mathcal{A}'$. We define a structure M over $\mathcal{A}$ by defining $\mathcal{K}_i = \mathcal{K}_{\sigma(i)}$. Again an easy induction shows that for every formula ψ with sets chosen from $\mathcal{G}_\varphi$, $(M', s) \models \psi$ if and only if $(M, s) \models \psi^\sigma$. Once again, the only point to notice is that E_G (and so also C_G) has the same meaning in M' (in terms of reachability) as $E_{\sigma(G)}$ ($C_{\sigma(G)}$) in M by the definition of $\sigma(G)$ and the relations $\mathcal{K}_j$. Thus $(M, s) \models \varphi$ as required. ∎

Corollary 3.6: *Given an oracle that decides, for each Boolean combination G of elements in $\mathcal{G}_{\mathcal{A}}$, whether $G = \emptyset$, there is an algorithm that decides whether a formula $\varphi \in \mathcal{L}_{\mathcal{G}}^C$ is satisfiable in $\mathcal{M}_{\mathcal{A}}$ that runs in time $2^{c2^{|\varphi|}}$ for some $c > 0$.*

Proof: Clearly, to check if φ is satisfiable, it suffices to check if φ^σ is satisfiable. In [7], there is an exponential time algorithm for checking satisfiability. However, this algorithm presumes that the set of agents is fixed. A close look at the algorithm actually shows that it runs in time $2^{cm|\varphi|}$, where m is the number of agents. In our translation, the set of agents is exponential in $|\varphi|$, giving us a double-exponential time algorithm. ∎

Corollary 3.7: *If $\mathcal{G}$ is closed under intersection and complementation, then $\mathrm{K}_{\mathcal{G}}^C$ is a sound and complete axiomatization for the language $\mathcal{L}_{\mathcal{G}}^C$ with respect to $\mathcal{M}_{\mathcal{A}}$.*

Proof: Soundness is straightforward, so we focus on completeness. Suppose φ is valid. By Proposition 3.5, so is φ^σ. Since $\mathcal{A}'$ is finite, $\mathrm{K}_{\mathcal{G}'}^C$ is a complete axiomatization for $\mathcal{L}_{\mathcal{G}'}^C$ with respect to $\mathcal{M}_{\mathcal{A}'}$. Thus, $\mathrm{K}_{\mathcal{G}'}^C \vdash \varphi^\sigma$. We can translate this proof step by step to a proof of φ in $\mathrm{K}_{\mathcal{G}}^C$. We simply replace every formula ψ that appears in the proof of φ^σ by ψ^τ, where ψ^τ is obtained by replacing each occurrence of K_A in ψ by E_A unless $A = \{i\}$ is a singleton, in which case we replace K_A by K_i, and replacing each occurrence of E_G, and C_G in ψ by $E_{\cup G}$, and $C_{\cup G}$, respectively. Since we have assumed $\mathcal{G}$ is closed under complementation and intersection, it is also closed under union, and hence ψ^τ is a formula in $\mathcal{L}_{\mathcal{G}}^C$.

It is easy to check that the translated proof is still a proof over the language $\mathcal{L}_{\mathcal{G}}^C$: Tautologies become tautologies as $(\varphi \vee \psi)^\tau = \varphi^\tau \vee \psi^\tau$ and similarly for negations. For K1 and MP it is just because $(\varphi \rightarrow \psi)^\tau = \varphi^\tau \rightarrow \psi^\tau$. Instances of KGen in the proof of φ^σ become instances of EGen or KGen in the proof of φ; similarly, instances of K1 are converted to instances of K1 or E1. It is easy to see that instances of E1, E2, E3, EGen, C1, and RC1 are converted to legitimate instances of the same axiom. ∎

While Corollaries 3.6 and 3.7 are close to our desired theorems, they also make clear the difficulties we need to overcome in order to prove Theorems 3.1 and 3.4. Specifically,

- we need to cut the complexity down from double-exponential to single exponential;
- we need to prove completeness without assuming that $\mathcal{G}$ is closed under complementation and intersection;
- we want to use an oracle that tests only whether a set of the form $G_0 - (G_1 \cup \ldots \cup G_k)$ is nonempty, rather than one that applies to arbitrary Boolean combinations;
- we want to extend these results to the case that the $\mathcal{K}_i$ relations satisfy properties like transitivity.

With regard to the last point, while in general it is relatively straightforward to extend completeness and complexity results to deal with relations that have properties like transitivity, it is not so straightforward in this case. For example, even if $M \in \mathcal{M}_{\mathcal{A}}^{rt}$, the relations in the structure M^σ constructed in Proposition 3.5 are not necessarily transitive. As shown in Proposition 3.3, we need a different oracle to deal with transitivity.

4 Proving the Main Results

In this section, we sketch the proof of Theorems 3.1 and 3.4 for $\mathcal{M}_{\mathcal{A}}$, $\mathcal{M}_{\mathcal{A}}^r$, and $\mathcal{M}_{\mathcal{A}}^{rt}$; we leave the proof in the case of $\mathcal{M}_{\mathcal{A}}^{rst}$ and $\mathcal{M}_{\mathcal{A}}^{elt}$ to the full paper. The structure of the proof is similar to that of Corollaries 3.6 and 3.7; we describe step by step the modifications required to deal with the problems raised in the previous section.

4.1 The Proof for $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}_{\mathcal{A}}^r$

In Proposition 3.5 we showed that we could translate a formula φ to a formula φ^σ such that φ was satisfiable in $\mathcal{M}_{\mathcal{A}}$ iff φ^σ was satisfiable in $\mathcal{M}_{\mathcal{A}'}$, where $\mathcal{A}'$ consisted of the atoms over $\mathcal{G}_\varphi$. Our goal is to maintain the translation idea, but use as our target set of agents a set whose elements we can determine with the oracles at our disposal (for testing the nonemptiness of certain set differences). As a first step, we try to abstract the key ingredients of Proposition 3.5. Suppose we have a set $\mathcal{A}'$ of agents and a partial map $\sigma : \mathcal{A} \rightarrow \mathcal{A}'$. Again, we can extend σ to a map from $2^{\mathcal{A}}$ to $2^{\mathcal{A}'}$: $\sigma(G) = \{\sigma(i) : i \in \mathcal{G}\}$. Given a formula φ, let φ^σ be the formula that results by replacing all the occurrences of G in φ by $\sigma(G)$. In Proposition 3.5, $\mathcal{A}'$ is the set of atoms over $\mathcal{G}_\varphi$ and $\sigma(i)$ is the unique atom containing i. We were able to show that, for that choice of $\mathcal{A}'$ and σ, the formulas φ and φ^σ were equisatisfiable. What does it take to obtain such a result in general? The following result shows that we need to be able to find a mapping $\tau : \mathcal{A}' \rightarrow 2^{\mathcal{A}} - \{\emptyset\}$ with one key property.

Proposition 4.1: *Given a formula φ and a partial map $\sigma : \mathcal{A} \rightarrow \mathcal{A}'$ such that $\sigma(G) \neq \emptyset$ for all $G \in \mathcal{G}_\varphi$, suppose that there is a mapping $\tau : \mathcal{A}' \rightarrow 2^{\mathcal{A}} - \{\emptyset\}$ such that for all $G \in \mathcal{G}_\varphi$, we have $\cup\{\tau(A) : A \in \sigma(G)\} = G$. Then φ is satisfiable in $\mathcal{M}_{\mathcal{A}}$ (resp., $\mathcal{M}_{\mathcal{A}}^r$) iff φ^σ is satisfiable in $\mathcal{M}_{\mathcal{A}'}$ (resp., $\mathcal{M}_{\mathcal{A}'}^r$).*

Proof: Left to the full paper. ∎

For the mapping σ of Proposition 3.5 we can take τ to be the identity, but this requires an oracle for nonemptiness of atoms. We now show how to choose $\mathcal{A}'$ and define maps σ and τ in a way that requires only information about whether sets of the form $G_0 - (G_1 \cup \ldots \cup G_k)$ are empty.

Definition 4.2: Given a set $\mathcal{G}$ of sets of agents and $G \in \mathcal{G}$, a set $\mathcal{H} \subseteq \mathcal{G}$ is a G-*maximal* subset of $\mathcal{G}$ if $G - \cup\mathcal{H} \neq \emptyset$ and $G - ((\cup\mathcal{H}) \cup G') = \emptyset$ for all $G' \in \mathcal{G} - \mathcal{H}$. Let $\mathcal{R}(\mathcal{G}) = \{(G, \mathcal{H}) | G \in \mathcal{G}, \mathcal{H}$ is a G-maximal subset of $\mathcal{G}\}$. ∎

Note that we can check whether $\mathcal{H}$ is a G-maximal subset of $\mathcal{G}$ by doing at most $|\mathcal{G}|$ tests of the form $(G - \cup\mathcal{H}') = \emptyset$, and we can find all pairs $(G, \mathcal{H})$ in $\mathcal{R}(\mathcal{G}_\varphi)$ by doing at most $|\mathcal{G}|2^{|\mathcal{G}|-1}$ such tests.

Lemma 4.3: *Suppose* $(G, \mathcal{H}) \in \mathcal{R}(\mathcal{G})$ *for some set* $\mathcal{G}$ *of subsets of* $\mathcal{A}$.

(a) $G - \cup\mathcal{H}$ *is an atom over* $\mathcal{G}$ *and, in fact,* $G - \cup\mathcal{H} = \cap(\mathcal{G} - \mathcal{H}) \cap (\cap_{H \in \mathcal{H}} \overline{H})$.

(b) *If* $(G', \mathcal{H}) \in \mathcal{R}(\mathcal{G})$*, then* $(G - \bigcup\mathcal{H}) = (G' - \bigcup\mathcal{H})$.

(c) *If* $(G', \mathcal{H}') \in \mathcal{R}(\mathcal{G})$ *and* $\mathcal{H} \neq \mathcal{H}'$*, then* $(G - \bigcup\mathcal{H}) \cap (G' - \bigcup\mathcal{H}') = \emptyset$.

Proof: Left to the full paper. ∎

If $(G, \mathcal{H}) \in \mathcal{R}(\mathcal{G})$, let $A^{\mathcal{G}}_{\mathcal{H}}$ denote the atom associated with $\mathcal{H}$ defined in Lemma 4.3(a). It is independent of G by Lemma 4.3(b). We omit $\mathcal{G}$ when it is clear from the context which set we have in mind.

We now show how to define a translation satisfying the hypotheses of Proposition 4.1 using the elements of $\mathcal{R}(\mathcal{G}_\varphi)$ identified according to the second coordinate alone.

Given a formula φ, let $\mathcal{A}^\varphi = \{\mathcal{H} : \exists G[(G, \mathcal{H}) \in \mathcal{R}(\mathcal{G}_\varphi)]\}$. Define $\sigma_1 : \mathcal{A} \to \mathcal{A}^\varphi$ by setting $\sigma_1(i) = \mathcal{H}$ if $i \in A_{\mathcal{H}}$ (as defined after Lemma 4.3) and undefined otherwise. As before, we extend σ_1 to $2^{\mathcal{A}}$ by defining $\sigma_1(G) = \{\sigma_1(i) : i \in G\}$.

Lemma 4.4: *Define* $\tau : \mathcal{A}^\varphi \to 2^{\mathcal{A}}$ *by setting* $\tau(\mathcal{H}) = \cap(\mathcal{G}_\varphi - \mathcal{H})$*. Then*

(a) $\sigma_1(G) = \{\mathcal{H} : \exists G'((G', \mathcal{H}) \in \mathcal{R}(\mathcal{G}_\varphi)),\, G \notin \mathcal{H}\}$.

(b) $\sigma_1(G) \neq \emptyset$ *for* $G \in \mathcal{G}_\varphi$*,*

(c) $\tau(\mathcal{H}) \neq \emptyset$ *for* $\mathcal{H} \in \mathcal{A}^\varphi$*,*

(d) $\cup\{\tau(\mathcal{H}) : \mathcal{H} \in \sigma_1(G)\} = G$.

Proof: Left to the full paper. ∎

Since $|\mathcal{A}| \leq 2^{|\varphi|}$, we have now reduced satisfiability with infinitely many agents to satisfiability with finitely many agents, at least for $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}^r_{\mathcal{A}}$, using only tests that we know we will need to be able to perform in any case. We next must deal with the problem we observed in the proof of Corollary 3.6, that is, there may be exponentially many agents in the subgroups mentioned in φ^{σ_1}. This is done in the following result. In this result, we assume that the complexity of checking whether $i \in G$ is no worse than linear in $|\mathcal{A}|$. While we do not assume this in general, it is true for the $\mathcal{A}'$ and sets G that arise in the translation of Proposition 4.1, which suffices for our application of the result to the proof of Theorem 3.4.

Theorem 4.5: *If* $\mathcal{A}$ *is finite and there is an algorithm for deciding if* $i \in G$ *for* $G \in \mathcal{G}$ *that runs in time linear in* $|\mathcal{A}|$*, then there is a procedure that, given a formula* φ *of* $\mathcal{L}^C_{\mathcal{G}}$*, decides whether* φ *is satisfiable with respect to* $\mathcal{M}_{\mathcal{A}}$ *(resp.,* $\mathcal{M}^r_{\mathcal{A}}$*) and runs in time* $O(|\mathcal{A}|2^{c|\varphi|})$ *for some* $c > 0$.

Proof: We first present an algorithm that decides if φ is satisfiable in $\mathcal{M}_{\mathcal{A}}$; we then show how to modify it to deal with $\mathcal{M}^r_{\mathcal{A}}$. The algorithm is just a slight modification of standard decision procedures [2, 7]. (Far more serious modifications are needed to prove the analogous result for the $\mathcal{M}^{rt}_{\mathcal{A}}$, $\mathcal{M}^{rst}_{\mathcal{A}}$, and $\mathcal{M}^{elt}_{\mathcal{A}}$; see Theorem 4.9.)

Let *Sub*(φ) be the set of subformulas of φ together with $E_G(\psi \wedge C_G\psi)$ and $\psi \wedge C_G\psi$ for each subformula $C_G\psi$ of φ. *Sub*$^+(\varphi)$ consists of the formulas in *Sub*(φ) and their negations. An easy induction on $|\varphi|$ shows that $|Sub(\varphi)| \leq |\varphi|$, so $|Sub^+(\varphi)| \leq 2|\varphi|$. (Here we need to use the fact that we take the length of C_G to be 3.)

Let S^1 consist of all subsets s of *Sub*$^+(\varphi)$ that are *maximal*, i.e., (a) for each formula $\psi \in Sub(\varphi)$, either $\psi \in s$ or $\neg\psi \in s$, (b) propositionally consistent (in that, for example, we cannot have all of $\psi \wedge \psi'$, $\neg\psi$, and $\neg\psi'$ in s), and (c) contain $E_G(\psi \wedge C_G\psi)$ iff they contain $C_G\psi$. Note that there are at most $2^{|\varphi|}$ sets in $S^1(\varphi)$.

For $s \in S^1$ and $G \in \mathcal{G}_{\mathcal{A}}$, we define $s/E_G = \{\psi | E_G\psi \in s\}$ (again, we identify K_i with $E_{\{i\}}$). Define $s/\overline{K_i} = \cup_{i \in G} s/E_G$. Define a binary relation $\mathcal{K}_i$ on S^1 for each $i \in \mathcal{A}$ by taking $(s, t) \in \mathcal{K}_i$ iff $s/\overline{K_i} \subseteq t$. We now define a sequence S^j of subsets of S^1. Suppose we have defined $S^1, \ldots, S^j$. S^{j+1} consists of all states in S^j that *seem consistent*, in that the following two conditions hold.

1. If $\neg E_G\psi \in s$, then there is some $t \in S^j$ such that $(s, t) \in \cup_{i \in G}\mathcal{K}_i$ and $\neg\psi \in t$.
2. If $\neg C_G\psi \in s$, then there is some $t \in S^j$ such that t is G-reachable from s in S^j and $\neg\psi \in t$.

If $S^j \neq S^{j+1}$ then we continue the construction. Otherwise the construction terminates; in this case, the algorithm returns "φ is satisfiable" if $\varphi \in s$ for some state $s \in S^{j+1}$ and returns "φ is unsatisfiable" otherwise.

Since $S^j \supseteq S^{j+1}$, S^1 has at most $2^{|\varphi|}$ elements, and there are $|\mathcal{A}|$ relations, it is easy to see that the whole procedure can be carried out in time $O(|\mathcal{A}|2^{c|\varphi|})$ for some $c > 0$.

It remains to show that the algorithm is correct. The ideas are relatively straightforward and left to the full paper.

To deal with $\mathcal{M}^r_{\mathcal{A}}$, the only change necessary is that in going from S^1 to S^2 in the construction, we also eliminate $s \in S^1$ if $(s, s) \notin \mathcal{K}_i$ for some $i \in \mathcal{A}$. This guarantees that the $\mathcal{K}_i$ relations are reflexive. The remainder of the proof goes through unchanged. ∎

Proof of Theorem 3.4 for $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}^r_{\mathcal{A}}$: The deterministic exponential time lower bound in Theorem 3.4 follows from the lower bound in the case where $\mathcal{A}$ is finite, which is proved in [7, Theorem 6.19] using techniques developed by Fischer and Ladner [3] for PDL. The sets G that arise in the lower bound proof have cardinality 2, so oracles are of no help here.

For the upper bound, suppose we are given a formula φ. We first compute the set $\mathcal{R}(\mathcal{G}_\varphi)$. This can be done with at most $|\varphi|2^{|\varphi|}$ calls to oracle O_0, since $|\mathcal{G}_\varphi| \leq |\varphi|$ and we need only check, for each $G \in \mathcal{G}_\varphi$ and $\mathcal{H} \subseteq \mathcal{G}_\varphi$, whether $G - \mathcal{H} = \emptyset$.

Consider the mapping σ_1 of Lemma 4.4. By part (a) of Lemma 4.4, we can compute the formula φ^{σ_1} using $\leq |\varphi|2^{|\varphi|}$ calls to oracle O_0. By Proposition 4.1 and Lemma 4.4, the formulas φ and φ^{σ_1} are equisatisfiable. By Theorem 4.5, we can decide if φ^{σ_1} is satisfiable in time $O(2^{c|\varphi|})$ for some $c > 0$ (since $|\varphi^{\sigma_1}| = |\varphi|$ and the set $\mathcal{A}$ of agents that appear in φ^{σ_1} has size at most $2^{|\varphi|}$). ∎

We now want to prove Theorem 3.1 in the case of $\mathcal{M}_{\mathcal{A}}$ and $\mathcal{M}^r_{\mathcal{A}}$. The idea is the same as that of Corollary 3.7. If φ is valid, then so is φ^{σ_1}. We can then appeal to completeness in the case of finitely many agents to get a proof of φ^{σ_1} that we can then "pull back" to a proof of φ. There is only one difficulty that we encounter when trying to put this idea into practice. Exactly how do we pull back the proof? For example, suppose the proof of φ^{σ_1} involves a formula ψ with an operator $K_{\mathcal{H}}$. In general, there will be many agents $i \in \mathcal{A}$ such that $\sigma_1(i) = \mathcal{H}$. One option is to replace $K_{\mathcal{H}}$ by $E_{\sigma_1^{-1}(\mathcal{H})}$, that is, replace $\mathcal{H}$ by all i such that $\sigma_1(i) = \mathcal{H}$. (This is what was done in the proof of Corollary 3.7.) The problem with this is that there is no guarantee that the resulting set is in $\mathcal{G}$. Alternatively, we could replace $K_{\mathcal{H}}$ by K_i for some i such that $\sigma_1(i) = \mathcal{H}$. But if so, which one?

We actually take the latter course here. We solve the problem of which i to choose by showing that there is a proof of φ^{σ_1} in which the only modal operators that arise in any formula used in the proof are modal operators that appear in φ^{σ_1} (Lemma 4.7). For these operators, there is a canonical way to do the replacement (Lemma 4.6). While it may seem almost trivial that the only operators that should be needed in the proof of φ^{σ_1} are ones that already appear in the formula, this is not the case for the standard completeness proof [2, 7], since in the proof of the validity of a formula of the form $E_G\psi$, the modal operators K_i are used for $i \in G$, although these operators may not appear in ψ. It is important that we use the axioms E1 and E2 in doing the proof, rather than the axiom E; otherwise the result would not hold. Indeed, the result does not quite hold in the case of $\mathrm{T}^C_{\mathcal{G}}$; we need to augment it with E5.

Lemma 4.6: *The mapping σ_1 (when viewed as a map with domain $2^{\mathcal{A}}$) is injective on $\mathcal{G}_\varphi$.*

Proof: Suppose $G \neq G'$. Without loss of generality, suppose that $i \in G - G'$. Then there is a G-maximal set $\mathcal{H}$ that includes G'. By Lemma 4.4(a), we have $\mathcal{H} \in \sigma_1(G)$. Since $G' \in \mathcal{H}$, it follows from Lemma 4.4(a) that $\mathcal{H} \notin \sigma_1(G')$. Thus, $\sigma_1(G) \neq \sigma_1(G')$. ∎

For the next lemma, we write $AX \vdash_\varphi \psi$ if there is a proof of φ in AX that involves only modal operators that appear in φ. Let $(\mathrm{T}^C_{\mathcal{G}})^+$ consist of $\mathrm{T}^C_{\mathcal{G}}$ augmented with the axiom E5. Although E5 follows from E1 and K2, using E5 allows us to be able to write proofs of φ that use only the modal operators in φ.

Lemma 4.7: *If $\mathcal{A}$ is finite and $\varphi \in \mathcal{L}^C_{\mathcal{G}}$ is valid with respect to $\mathcal{M}_{\mathcal{A}}$ (resp., $\mathcal{M}^r_{\mathcal{A}}$), then $\mathrm{K}^C_{\mathcal{G}} \vdash_\varphi \varphi$ (resp., $(\mathrm{T}^C_{\mathcal{G}})^+ \vdash_\varphi \varphi$).*

Proof: We first consider the case of $\mathcal{M}_{\mathcal{A}}$. Since φ is valid, $\neg\varphi$ is not satisfiable. That means, when we apply the construction in the proof of Theorem 4.5 to $\neg\varphi$, all the sets containing $\neg\varphi$ are eliminated. For each state $s \in S^1$, let φ_s be the conjunction of all the formulas in s.

We prove the result by showing, by induction on j, that if a state $s \in S^j$ does not seem consistent, then φ_s is $\mathrm{K}^C_{\mathcal{G}}$-inconsistent, i.e., $\mathrm{K}^C_{\mathcal{G}} \vdash_\varphi \neg\varphi_s$. This suffices to prove the lemma, since, by C1 and RC1, nothing is lost by our assumption that $C_G\psi \in s$ iff $E_G(\psi \wedge C_G\psi)$, standard propositional reasoning (using Prop and MP) shows that, for any formula $\psi \in \mathit{Sub}(\neg\varphi)$,

$$\mathrm{K}^C_{\mathcal{G}} \vdash_\varphi \psi \Leftrightarrow \vee_{\{s \in S^1 : \psi \in s\}} \varphi_s.$$

Negating both sides of $\Leftrightarrow$, we get

$$\mathrm{K}^C_{\mathcal{G}} \vdash_\varphi \neg\psi \Leftrightarrow \wedge_{\{s \in S^1 : \psi \in s\}} \neg\varphi_s.$$

Thus, if $\mathrm{K}^C_{\mathcal{G}} \vdash_\varphi \neg\varphi_s$ for each set s containing $\neg\varphi$, it follows by standard propositional reasoning that $\mathrm{K}^C_{\mathcal{G}} \vdash_\varphi \varphi$, as desired.

While this general approach to proving completeness is quite standard, we must take extra care because of our insistence on restricting to symbols that appear in φ, particularly when dealing with the case when a state seems inconsistent due to a formula of the form $\neg E_G\psi$ or $\neg C_G\psi$ not being satisfied. This is where the axioms E1 and E2 come into play. We leave details to the full paper. ∎

Proof of Theorem 3.1 for $\mathcal{M}_\mathcal{A}$ and $\mathcal{M}^r_\mathcal{A}$: We have already observed that the axioms are sound. For completeness, suppose that φ is valid with respect to $\mathcal{M}_\mathcal{A}$. By Proposition 4.1, so is φ^{σ_1}. By Lemma 4.7, there is a proof of φ^{σ_1} in $\mathrm{K}^C_{\mathcal{G}_\varphi}$ that mentions only the modal operators in φ^{σ_1}. Given a formula ψ in which the only modal operators that appear are modal operators that appear in φ^{σ_1} (and thus have the form $E_{\sigma_1(G)}$, $C_{\sigma_1(G)}$, and $K_{\sigma_1(i)}$, for sets G and $\{i\}$ in $\mathcal{G}_\varphi$) let ψ^{τ_1} be the unique formula all of whose modal operators appear in φ such that $(\psi^{\tau_1})^{\sigma_1} = \psi$. Lemma 4.6 assures us that ψ^{τ_1} is well defined. We can pull the proof of φ^{σ_1} back to a proof of φ, by replacing each occurrence of a formula ψ in the proof by ψ^{τ_1}.

The argument for $\mathcal{M}^r_\mathcal{A}$ is identical, except that the proof uses instances of the axiom E5. These can be eliminated by using E1 and K2, as we oberved earlier (although now the proof of φ may use modal operators K_i that do not appear in φ). ∎

4.2 Dealing with $\mathcal{M}^{rt}_\mathcal{A}$

Proposition 4.1 as it stands does not hold for $\mathcal{M}^{rt}_\mathcal{A}$. There is no guarantee that the translated formula is satisfiable in a structure in $\mathcal{M}^{rt}_\mathcal{A}$, even if φ is. Indeed, suppose that $\mathcal{G}$ is closed under intersection and complementation, so that we can use the function σ of Proposition 3.5. Suppose that φ is the formula $E_G p \wedge \neg E_G E_G p$, where $|G| \geq 2$. The formula φ^σ looks syntactically identical, except that $\sigma(G)$ is a single agent in $\mathcal{A}'$. We cannot make the $\mathcal{K}_G$ relation transitive and still satisfy φ^σ. More generally, to deal with $\mathcal{M}^{rt}_\mathcal{A}$, we must be careful in how we deal with singleton sets.

As a first step, we define *mixed* structures. Let $\mathcal{M}^{rt}_{\mathcal{A}_1+\mathcal{A}_2}$ consist of structures $M = (S, \pi, \{\mathcal{K}_i : i \in \mathcal{A}_1 \cup \mathcal{A}_2\})$ where the relations $\mathcal{K}_i$ for $i \in \mathcal{A}_1$ are reflexive and transitive and the relation $\mathcal{K}_i$ for $i \in \mathcal{A}_2$ are reflexive.

We can now define our translation in the case of $\mathcal{M}^{rt}_\mathcal{A}$. It turns out to be easier to provide a translation that combines Proposition 4.1 and Lemma 4.4, rather than separating them. As suggested by Proposition 3.3, the translation involves $\mathcal{R}(\mathcal{G}^1_\varphi)$, rather than $\mathcal{R}(\mathcal{G}_\varphi)$. Given a formula φ, let $\mathcal{A}^{\varphi,rt} = \{\mathcal{H} : \exists G[(G,\mathcal{H}) \in \mathcal{R}(\mathcal{G}^1_\varphi)]\}$. Let $\mathcal{A}_1 = \{\mathcal{H} : \exists G[(G,\mathcal{H}) \in \mathcal{R}(G^1_\varphi), |G - \cup\mathcal{H}| = 1]\}$; let $\mathcal{A}_2 = \mathcal{A}^{\varphi,rt} - \mathcal{A}_1$. Define $\sigma_2 : \mathcal{A} \to \mathcal{A}^{\varphi,rt}$ as before: $\sigma_2(i) = \mathcal{H}$ if $i \in A_\mathcal{H}$ and $\sigma_2(i)$ is undefined otherwise. Given $\mathcal{H} \in \mathcal{A}^{\varphi,rt}$, we define $\tau_2(\mathcal{H}) = \cap(\mathcal{G}^1_\varphi - \mathcal{H})$. Since it is easy to see that $\mathcal{R}(\mathcal{G}^1_\varphi) = \mathcal{R}(\mathcal{G}_\psi)$ for some appropriate ψ, it is immediate that Lemma 4.4 applies to σ_2 and τ_2.

Proposition 4.8: *φ is satisfiable in $\mathcal{M}^{rt}_\mathcal{A}$ iff φ^{σ_2} is satisfiable in $\mathcal{M}^{rt}_{\mathcal{A}_1+\mathcal{A}_2}$.*

Proof: First suppose $(M,s) \models \varphi$, where $M \in \mathcal{M}^{rt}_\mathcal{A}$. We convert $M = (S, \pi, \{\mathcal{K}_i : i \in \mathcal{A}\})$ into a structure $M' = (S, \pi, \{\mathcal{K}_\mathcal{H} : \mathcal{H} \in \mathcal{A}^{\varphi,rt}\})$ as before, by defining $\mathcal{K}_\mathcal{H} = \cup\{\mathcal{K}_i : i \in \tau_2(A)\}$. Since Lemma 4.4 applies, the proof that $(M',s) \models \varphi$ is identical to that in Proposition 4.1. We must only show that $M' \in \mathcal{M}^{rt}_{\mathcal{A}_1+\mathcal{A}_2}$. Since the union of reflexive relations is reflexive, it is immediate that $\mathcal{K}_\mathcal{H}$ is reflexive for $\mathcal{H} \in \mathcal{A}_2$. If $\mathcal{H} \in \mathcal{A}_1$, then $|A_\mathcal{H}| = 1$. Suppose that $A_\mathcal{H} = \{i\}$. We claim that $\tau_2(\mathcal{H}) = \{i\}$. By construction, $\{i\} \in \mathcal{G}^1_\varphi$. We cannot have $\{i\} \in \mathcal{H}$, since $i \notin \cup\mathcal{H}$. Thus, $\tau_2(\mathcal{H}) = \cap(\mathcal{G}^1_\varphi - \mathcal{H}) \subseteq \{i\}$. Since $\tau_2(\mathcal{H}) \neq \emptyset$ by Lemma 4.4(c), we must have $\tau_2(\mathcal{H}) = \{i\}$. Thus, $\mathcal{K}_\mathcal{H} = \mathcal{K}_i$, so $\mathcal{K}_\mathcal{H}$ is reflexive and transitive.

For the opposite direction we need to work a little harder than before, because we must ensure that all the $\mathcal{K}_i$ relations are reflexive and transitive for all $i \in \mathcal{A}$. Supppose $(M,s) \models \varphi^{\sigma_2}$ for some $M = (S, \pi, \{\mathcal{K}_\mathcal{H} : \mathcal{H} \in \mathcal{A}^{\varphi,rt}\}) \in \mathcal{M}^{rt}_{\mathcal{A}_1+\mathcal{A}_2}$. Let S_0 and S_1 be two disjoint copies of S. For a state $s \in S$, let s_i be the copy of s in S_i, $i = 0,1$. Let $M' = (S', \pi', \{\mathcal{K}_i : i \in \mathcal{A}\})$ be defined as follows:

- $S' = S_0 \cup S_1$.
- $\pi'(s_i) = \pi(s)$ for $i = 0,1$.
- If $\sigma_2(i) \in \mathcal{A}_1$, define $\mathcal{K}_i = \{(s_i, t_j) : (s,t) \in \mathcal{K}_{\sigma_2(i)}, i,j \in \{0,1\}\}$. $\mathcal{K}_i$ is clearly reflexive and transitive in this case, since $\mathcal{K}_{\sigma_2(i)}$ is.
- If $\sigma_2(i) = \mathcal{H} \in \mathcal{A}_2$, note that $|A_\mathcal{H}| \geq 2$. It is immediate from the definition that $\sigma_2(i) = \mathcal{H}$ for all $i \in A_\mathcal{H}$. Pick some $i_\mathcal{H} \in A_\mathcal{H}$. If $i = i_\mathcal{H}$, then define $\mathcal{K}_i = \{(s_0, t_1) : (s,t) \in \mathcal{K}_\mathcal{H}\} \cup \{(s_j, s_j) : j \in \{0,1\}\}$; if $i \neq i_\mathcal{H}$, define $\mathcal{K}_i = \{(s_1, t_0) : (s,t) \in \mathcal{K}_\mathcal{H}\} \cup \{(s_j, s_j) : j \in \{0,1\}\}$. Clearly $\mathcal{K}_i$ is reflexive and transitive.

This construction guarantees that

$$(s,t) \in \mathcal{K}_\mathcal{H} \text{ iff } (s_0,t_1),(s_1,t_0) \in \cup_{\{i:\sigma_2(i)=\mathcal{H}\}}\mathcal{K}_i \qquad (1)$$

and

$$(s_1,t_0) \in \cup_{\{i:\sigma_2(i)=\mathcal{H}\}}\mathcal{K}_i \text{ iff } (s_0,t_1) \in \cup_{\{i:\sigma_2(i)=\mathcal{H}\}}\mathcal{K}_i. \qquad (2)$$

A straightforward argument by induction on structure now shows that if $\psi \in \mathcal{L}^C_{\mathcal{G}_\varphi}$, then the following are equivalent for all $t \in S$:

- $(M,t) \models \psi^{\sigma_2}$
- both $(M',t_0) \models \psi$ and $(M',t_1) \models \psi$
- $(M',t_0) \models \psi$ or $(M',t_1) \models \psi$.

Of course, the interesting cases are if ψ is of the form $K_i\psi'$, $E_G\psi'$, or $C_G\psi'$. These follow immediately from observations (1) and (2). ∎

The next step is to get an analogue of Theorem 4.5 for $\mathcal{M}^{rt}_{\mathcal{A}_1+\mathcal{A}_2}$. The basic idea of the proof is the same as that of Theorem 4.5. However, in our construction, we need to make the $\mathcal{K}_i$ relations transitive. To see the difficulty, suppose that φ is $K_1 p \wedge E_G q$, where G is a set of agents containing 1. Recall that in Theorem 4.5, states are consistent subsets of $Sub^+(\varphi)$. Let s, t, and u be states such that $s = \{K_1 p, E_G q, p, q\}$, $t = \{K_1 p, \neg E_G q, p, q\}$, and $u = \{K_1 p, \neg E_G q, p, \neg q\}$. With our previous construction, we would have both $(s,t) \in \mathcal{K}_1$ and $(t,u) \in \mathcal{K}_1$. But by transitivity, we should also have $(s,u) \in \mathcal{K}_1$. But since $E_G q \in s$ and $\neg q \in u$, we have $(s,u) \notin \mathcal{K}_1$. Nevertheless, each of s, t, and u individually seems consistent. Which state should we eliminate in order to preserve transitivity?

To deal with this problem, we need to put more information (i.e., more formulas) into each state. Intuitively, if $(s,t) \in \mathcal{K}_i$, then we should have $K_i q \in t$, because if $E_G q \in s$, then $K_i q$ should also be in s, as should $K_i K_i q$ by K4. It would then follow that $K_i q$ should be in t. This, in turn, would guarantee that $(t,u) \notin \mathcal{K}_i$, since $q \notin u$.

What we would like to do now is to augment $Sub(\varphi)$ by including all formulas $K_i\psi$ such that $E_G\psi \in Sub(\varphi)$ and $i \in G \cap \mathcal{A}_1$. (We restrict to $\mathcal{A}_1$ since these are the only relations that are required to be transitive.) While doing this will allow us to force the $\mathcal{K}_i$'s to be transitive, the resulting set of formulas can have size $O(|\mathcal{A}_1||\varphi|)$, which means the resulting state space (the analogue of S^1) could then have size $2^{|\mathcal{A}_1||\varphi|}$. This would not give us the desired complexity bounds. Thus, we must proceed a little more cautiously.

Theorem 4.9: *If $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$ is finite and there is an algorithm for deciding if $i \in G$ for $G \in \mathcal{G}$ that runs in time linear in $|\mathcal{A}|$, then there is a procedure that, given a formula φ of $\mathcal{L}^C_{\mathcal{G}}$, decides whether φ is satisfiable with respect to $\mathcal{M}^{rt}_{\mathcal{A}_1+\mathcal{A}_2}$ and runs in time $O(|\mathcal{A}|2^{c|\varphi|})$ for some $c > 0$.*

Proof: We assume for ease of exposition that $\mathcal{A}_1 \neq \emptyset$; we leave the straightforward modification in case $\mathcal{A}_1 = \emptyset$ to the reader. For each $i \in \mathcal{A}_1$, let $ESub_i(\varphi)$ be the least set containing $Sub(\varphi)$ such that if $E_G\varphi \in Sub(\varphi)$ and $i \in G$, then $K_i\varphi \in ESub_i(\varphi)$. It is easy to see that $|ESub_i(\varphi)| \leq 2|Sub(\varphi)|$, since we add at most one formula for each formula in $|Sub(\varphi)|$. Let S^1_i consist of all the subsets of $ESub^+_i(\varphi)$ that are maximal, and now let $S^1 = \cup_{i \in \mathcal{A}_1} S^1_i$. Note that, as modified, $|S^1| \leq 2^{2|\varphi|}$. Thus, this modification keeps us safely within the desired exponential time bounds.

We keep the definition of $\mathcal{K}_i$ unchanged for $i \in \mathcal{A}_2$, but we need to modify it for $i \in \mathcal{A}_1$. We redefine $\mathcal{K}_i$ for $i \in \mathcal{A}_1$ by defining $(s,t) \in \mathcal{K}_i$ iff $s/\overline{K_i} \cup \{K_i\psi : K_i\psi \in s\} \subseteq t \cap (t/\overline{K_i} \cup \{K_i\psi : K_i\psi \in t\})$. It is easy to check that this modification forces the $\mathcal{K}_i$ relations to be transitive. We force all the $\mathcal{K}_i$ relations to be reflexive just as with $\mathcal{M}^r_{\mathcal{A}}$, by eliminating $s \in S^1$ if $(s,s) \notin \mathcal{K}_i$ for some $i \in \mathcal{A}_1 \cup \mathcal{A}_2$. The remainder of the construction—eliminating the states that do not seem consistent—is unchanged. We leave the proof of correctness of the algorithm to the full paper. ∎

Modulo the changes required because we are dealing with mixed structures, the proofs of Theorems 3.4 and 3.1 in the case of $\mathcal{M}^{rt}_{\mathcal{A}}$ now proceed in much the same spirit as those for $\mathcal{M}_{\mathcal{A}}$. The details can be found in the full paper. Dealing with $\mathcal{M}^{rst}_{\mathcal{A}}$ and $\mathcal{M}^{elt}_{\mathcal{A}}$ introduces yet new subtleties; these too can be found in the full paper.

5 Conclusions

We have characterized the complexity of satisfiability for epistemic logics when the set of agents is infinite. Our results emphasize the importance of how the sets of agents are described and provide new information even in the case where the sets involved are finite. We have focused here on a logic with a common knowledge operator. What happens if we have only the E_G operators? Although we have not checked details, it seems relatively straightforward to combine the techniques of [7] with those presented here to get PSPACE completeness for $\mathcal{L}^E_{\mathcal{G}}$ (even with infinitely many agents) using the same types of oracle calls as in Theorem 3.4. We are currently exploring what happens if we add the distributed knowledge operator D_G to the language [2, 7]. The interaction between E_G and D_G seems to add new subtleties.

References

[1] R. J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.

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

[3] M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.

[4] J. Geanakoplos. Common knowledge. In R. Aumann and S. Hart, editors, *Handbook of Game Theory, Vol. 2*. Elsevier Science, 1994.

[5] J. Y. Halpern and R. van der Meyden. A logic for SDSI's linked local name spaces. In *12th IEEE Computer Security Foundations Workshop*, 1999.

[6] J. Halpern, R. van der Meyden, and F. Schneider. A logical framework for trust management. Manuscript, 1999.

[7] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

[8] R. Rivest and B. Lampson. SDSI—a simple distributed security infrastructure. http://theory.lcs.mit.edu/∼cis/sdsi.html, 1996.