

The logic of communication graphs

Eric Pacuit² and Rohit Parikh¹

¹ Computer Science
The Graduate Center of CUNY,
365 5th Avenue, New York City, NY 10016
epacuit@cs.gc.cuny.edu,
www.cs.gc.cuny.edu/~epacuit

² CS, Math and Philosophy
Brooklyn College** and The Graduate Center of CUNY
365 5th Avenue, New York City, NY 10016
rparikh@gc.cuny.edu
www.sci.brooklyn.cuny.edu/~rparikh

Abstract. In 1992, Moss and Parikh studied a bimodal logic of knowledge and effort called *Topologic*. In this current paper, *Topologic* is extended to the case of many agents who are assumed to have some private information at the outset, but may refine their information by acquiring information possessed by other agents, possibly via yet other agents. Each agent's information is represented by a partition over a set of possible states, and when an agent learns a new piece of information, its partition is refined. The set of possible partitions is restricted to those that can arise via communication among the agents.

Let us assume that the agents are connected by a *communication graph*. In the communication graph, an edge from agent i to agent j means that agent i can directly receive information from agent j . Agent i can then refine its information by learning information that j has, including information acquired by j from another agent, k . We introduce a multi-agent modal logic with knowledge modalities and a modality representing communication among agents. We show that the validities of *Topologic* remain valid and that the communication graph is completely determined by the validities of the resulting logic. Applications of our logic to the Rice-Clarke dilemma are obvious.

1 Introduction

In [MP], Moss and Parikh introduce a bimodal logic intended to formalize reasoning about points and sets. This new logic called *Topologic* can also be understood as an epistemic logic with an effort modality. Formally, the two modalities are: K and \Diamond . The intended interpretation of $K\phi$ is that ϕ is known; and the intended interpretation of $\Diamond\phi$ is that after some amount of effort ϕ is true. For example, the formula

$$\phi \rightarrow \Diamond K\phi$$

** 2900 Bedford Avenue, Brooklyn, NY 11210

means that if ϕ is true, then after some “work”, ϕ is known, i.e., if ϕ is true, then ϕ can be known with some effort. What exactly is meant by “effort” depends on the application. For example, we may think of effort as meaning taking a measurement, performing a calculation or observing a computation. In this paper we will think of effort as meaning consulting some agent’s database of known formulas.

There is a temptation to think that the effort modality can be understood as (only) a temporal operator, reading $\diamond\phi$ as “ ϕ is true some time in the future”. While there is a connection between the logics of knowledge and time and logics of knowledge and effort (see [H99,H00] and references therein for more on this topic), following [MP] it is assumed that such effort leaves the base facts about the world unchanged. In particular, in all topologies if ϕ does not contain any knowledge modalities, then $\phi \leftrightarrow \Box\phi$ is valid. Thus, effort will not change the base facts about the world – it can only change knowledge of these facts.

The family of logics introduced in [MP] and later studied by Dabrowski, Moss and Parikh, Georgatos, Heinemann, and Weiss ([DMP,G93,G94,G97,H99,WP]) has a semantics in which the acquisition of knowledge is explicitly represented. Familiar mathematical structures such as subset spaces, topologies, intersection spaces and complete lattices of subsets corresponding to natural notions of knowledge acquisition are attached to standard Kripke structures.

Given a set W , a subset space is a pair $\langle W, \mathcal{O} \rangle$, where \mathcal{O} is a collection of subsets of W . A point $x \in W$ represents a complete observation about the world in which *all* facts are settled, whereas a set $U \in \mathcal{O}$ represents an observation. The pair (x, U) , called a *neighborhood situation*, can be thought of as an actual situation together with an observation made about the situation. Formulas are interpreted at neighborhood situations. Thus the knowledge modality K represents movement within the current observation, while the effort modality \diamond represents a refining of the current observation. [MP] provides a sound and complete axiomatization for all subset spaces. In [G93] and [G94], Georgatos provides a sound and complete axiomatization for subset spaces that are topological spaces and complete lattices. Dabrowski, Moss, and Parikh prove the same result using an embedding into **S4** ([DMP]). [G97] provides a sound and complete axiomatization for treelike spaces, and Weiss ([WP]) has provided a sound and complete axiomatization for intersection-spaces. Interestingly, it is shown in [WP] that an infinite number of axiom schemes are necessary for any complete axiomatization of intersection spaces. More recently, Heinemann [H99,H00] has looked at subset spaces and logics of knowledge and time, and the connection between hybrid logic and subset spaces [H02,H04].

In this paper, we present a multi-agent topologic in which the effort modality \diamond is intended to mean communication among agents. In order for any communication to take place, we must assume that the agents understand a common language. Thus we assume a set Φ_0 of propositional variables, understood by all the agents, but with only specific agents knowing their actual values at states in our models. The agents will refine their information by acquiring information possessed by other agents, possibly via other agents. This implies that if

agents are restricted in whom they can or cannot communicate with, then this fact will restrict the knowledge theoretic formulas that can come to be true, i.e., knowledge theoretic formulas in the scope of the effort modality.

Consider the current situation with Bush and Tenet. If Bush wants some information from a particular CIA operative, say Bob, he must get this information through Tenet. Suppose that ϕ is a formula representing the exact whereabouts of Bin Laden and that Bob is the CIA operative in charge of maintaining this information. In particular, $K_{\text{Bob}}\phi$, and suppose that at the moment, Bush does not know the exact whereabouts of Bin Laden ($\neg K_{\text{Bush}}\phi$). Obviously Bush can find out the exact whereabouts of Bin Laden ($\diamond K_{\text{Bush}}\phi$) by going through the appropriate channels, but of course, *we* cannot find out such information ($\neg \diamond K_e\phi \wedge \neg \diamond K_r\phi$) since we do not have the appropriate security clearance. Presumably, going through the appropriate channels implies that as a *pre-requisite* for Bush learning ϕ , Tenet will also have come to know ϕ . We can represent this situation by the following formula:

$$\neg K_{\text{Bush}}\phi \wedge \Box(K_{\text{Bush}}\phi \rightarrow K_{\text{Tenet}}\phi)$$

where \Box is the dual of diamond.

Let \mathcal{A} be a set of agents. A **communication graph** is a directed graph $G_{\mathcal{A}} = (\mathcal{A}, E)$ where $E \subseteq \mathcal{A} \times \mathcal{A}$. Intuitively $(i, j) \in E$ means that i can directly receive information from agent j , *without* j knowing this fact. Thus an edge between i and j in the communication graph represents a one-sided relationship between i and j . Agent i has access to any piece of information that agent j knows. For example, during a lecture the students have access to the lecturer's information, but not vice versa. Another common situation that is helpful to keep in mind is accessing a website. When there is an edge between i and j we think of agent j as creating a website in which everything he *currently* knows is available, and agent i can access this website without j being aware that the site is being accessed. Of course, j may be able to access another agent's website and so update some of his information. Therefore, it is important to stress that when i accesses j 's website, he is accessing j 's current information. It is of course possible that another agent has no access to j 's website, or only indirectly.

The assumption that i can access all of j 's information is a significant idealization from these common situations. This idealization rests on two assumptions: 1. all the agents share a common language, and 2. the agents make public all possible pieces of information which they know and which are expressible in this language. The fact that agents are assumed to share a common language is discussed in Section 4. For the second assumption, consider the tension between paparazzi and celebrities. This tension can be understood as the celebrities simply not wanting all of their current information made public. In other words, they want to remove, or at least restrict, the connection in the communication graph from the paparazzi to themselves. Or they may threaten a lawsuit between the paparazzi and the public media. Such assumptions can be dealt with in our framework, but a more detailed discussion will be reserved for the full version.

This paper is organized as follows. Section 2 formalizes what is meant by "communication". Section 3 presents the syntax and semantics of our logic, and

Section 4 proves the main technical result that the valid formulas characterize the communication graph. Finally in Section 5 we conclude and discuss further research.

2 Partition Spaces

In this section we develop our basic update operation. We first review some relevant facts and definitions about partitions on a set and define a partition space. Given a set W , a partition \mathcal{P} on W is a collection of nonempty sets $\mathcal{P} \subseteq 2^W$ such that

1. $\cup \mathcal{P} = W$
2. For all $P_1, P_2 \in \mathcal{P}$, $P_1 \neq P_2$ implies $P_1 \cap P_2 = \emptyset$

Elements of a partition \mathcal{P} will be called **partition cells**. Given an element $w \in W$ and a partition \mathcal{P} on W , let $\mathcal{P}(w)$ be the partition cell that contains w . That is $\mathcal{P}(w) = P$, where $P \in \mathcal{P}$ and $w \in P$.

Definition 1 (Partition Space). *A partition space is a tuple $\langle W, \mathbb{P} \rangle$, where W is a set and \mathbb{P} be a (finite) collection of partitions on W .*

Analogous to the subset spaces of [MP], a partition space is a set together with the set of partitions that we are interested in. For example, \mathbb{P} could be the set of all possible partitions on W . We think of \mathbb{P} as being the set of partitions that could possibly arise in a given situation. In this case, we do not have to consider “unlikely” partitions such as the singleton partition in which the agent knows all true facts.

Let \mathcal{P} and \mathcal{Q} be two partitions on W . We say that \mathcal{P} is a **refinement** of \mathcal{Q} , denoted $\mathcal{P} \preceq \mathcal{Q}$ if

$$\forall P \in \mathcal{P}, \exists Q \in \mathcal{Q} : P \subseteq Q$$

It is easy to see that \mathcal{P} is a refinement of \mathcal{Q} ($\mathcal{P} \preceq \mathcal{Q}$) iff each partition cell in \mathcal{Q} is a union of partition cells from \mathcal{P} . It is also not difficult to see that \preceq is reflexive, transitive and antisymmetric; hence a partial order. To see that \preceq is antisymmetric. Suppose that $\mathcal{P} \preceq \mathcal{Q}$ and $\mathcal{Q} \preceq \mathcal{P}$. Suppose that $X \in \mathcal{P}$. Then there is a $Y \in \mathcal{Q}$ such that $X \subseteq Y$; and since $Y \in \mathcal{Q}$ there is a $X' \in \mathcal{P}$ such that $Y \subseteq X'$. Hence $X \subseteq Y \subseteq X'$, which implies $X = Y$ (since it must be the case that $X = X'$). Therefore, $X \in \mathcal{Q}$. Similarly we can show that if $X \in \mathcal{Q}$, then $X \in \mathcal{P}$.

We say that \mathcal{P} is **finer** than \mathcal{Q} if $\mathcal{P} \preceq \mathcal{Q}$, or that \mathcal{Q} is **coarser** than \mathcal{P} . Suppose that \mathcal{P} represents i 's current information. Then moving to a refinement $\mathcal{Q} \preceq \mathcal{P}$ represents an increase in i 's knowledge. We will not be interested in *any* increase of knowledge, but rather any increase in knowledge caused by communication among agents governed by the communication graph. This will be discussed in more detail below. The following notation will turn out to be useful.

Definition 2. Let \mathcal{P} and \mathcal{Q} be two partitions on a set W . The **least common refinement** of \mathcal{P} and \mathcal{Q} , denoted by $\mathcal{P} \sqcap \mathcal{Q}$, is the partition generated by intersecting the cells from \mathcal{P} and \mathcal{Q} . That is,

$$\mathcal{P} \sqcap \mathcal{Q} \stackrel{\text{def}}{=} \{\mathcal{P}(w) \cap \mathcal{Q}(w) \mid w \in W\}$$

Clearly, $\mathcal{P} \sqcap \mathcal{Q}$ is a partition; and $\mathcal{P} \sqcap \mathcal{Q} \preceq \mathcal{P}$ and $\mathcal{P} \sqcap \mathcal{Q} \preceq \mathcal{Q}$. Given two partitions, it is easy to see that $\mathcal{P} \sqcap \mathcal{Q}$ is the coarsest partition that refines both \mathcal{P} and \mathcal{Q} , thus we can think of the operation \sqcap as a meet between \mathcal{P} and \mathcal{Q} . We can also define a join operation:

Definition 3. The **least coarsest partition** between \mathcal{P} and \mathcal{Q} is the finest partition \mathcal{R} such that $\mathcal{P} \preceq \mathcal{R}$ and $\mathcal{Q} \preceq \mathcal{R}$. We denote \mathcal{R} by $\mathcal{P} \sqcup \mathcal{Q}$.

Given a partition space $\langle W, \mathbb{P} \rangle$, \preceq is a partial order on \mathbb{P} , and \sqcap and \sqcup give us a meet and join respectively.

A partition \mathcal{P} for an agent i represents i 's current information. Thus when i learns a new piece of information, i 's partition \mathcal{P} is refined. Since any refinement of a partition is itself a partition, we must make some assumptions about what kind of information can be learned by an agent. We therefore assume that i can only update with true information. Otherwise, updating with a false piece of information may result in an agent acquiring a false, justified belief which cannot be represented using (only) partitions.

We also point out that upon receiving the information ϕ an agent might not come to know ϕ . For example, suppose that i is told "There is a bug on i 's shoulder, but i does not know it". Then, after i updates with this proposition ϕ , i will not know ϕ , but rather the proposition "There is a bug on i 's shoulder". These propositions of the form $\phi \wedge \neg K_i \phi$ were first discussed by G. E. Moore.

We first extend our basic notions to a multi-agent setting. Let \mathcal{A} be a finite set of agents. For simplicity, we assume that the set of states is the same for all agents. A **multi-agent partition** is a n -tuple $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_n)$, where n is the number of agents and each partition \mathcal{P}_i is a partition on W .

Definition 4 (Multi-Agent Partition Space). Given a set W , a **multi-agent partitions space** is a tuple $\langle W, \mathbb{P} \rangle$, where \mathbb{P} is a set of multi-agent partitions.

We think of a multi-agent partition space $\langle W, \mathbb{P} \rangle$ as a set of states together with all the n -tuples of partitions that could possibly arise *given a communication graph*. This will be made more precise below. We write \mathcal{P}_i for the i -th projection of \mathcal{P} . We can extend our notation defined above to vectors of partitions. If $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_n)$ and $\mathcal{Q} = (\mathcal{Q}_1, \dots, \mathcal{Q}_n)$ are vectors of partitions, then we write $\mathcal{P} \preceq \mathcal{Q}$ if $\mathcal{P}_i \preceq \mathcal{Q}_i$ for all $i = 1, \dots, n$. Other operators are defined pointwise on vectors:

$$\mathcal{P} \sqcap \mathcal{Q} \stackrel{\text{def}}{=} (\mathcal{P}_1 \sqcap \mathcal{Q}_1, \mathcal{P}_2 \sqcap \mathcal{Q}_2, \dots, \mathcal{P}_n \sqcap \mathcal{Q}_n)$$

Similarly for $\mathcal{P} \sqcup \mathcal{Q}$. So, a vector $\mathcal{P} \in \mathbb{P}$ represents each agent's information. As above, \preceq is a partial order on \mathbb{P} and \sqcap and \sqcup are meet and join respectively. In fact, we can say more. Given a vector \mathcal{P} , there is a vector \mathcal{P}^I that represents the implicit information that the agents currently have. \mathcal{P}^I is obtained by replacing each agent's partition with the least common refinement, i.e., for each $i \in \mathcal{A}$, $\mathcal{P}_i^I = \sqcap_{i \in \mathcal{A}} \mathcal{P}_i$. Similarly, we can define a common knowledge partition \mathcal{P}^C in which each agent's partition is replaced by the least coarsest partition. For each $i \in \mathcal{A}$, $\mathcal{P}_i^C = \sqcup_{i \in \mathcal{A}} \mathcal{P}_i$. Given \mathcal{P} , \mathcal{P}^C is the information that is commonly known among all the agents.

In fact we will not be interested in *any* refinement of i 's partition, but rather only those refinements that can arise from "communication" between two agents. Suppose that agent i can directly communicate with agent j , i.e., there is an edge between i and j in the communication graph. In this case any piece of information that j knows can be learned by agent i . Suppose that \mathcal{Q} is agent j 's partition and \mathcal{P} is agent i 's partition. Given a set $X \subseteq W$, i can update his partition with X provided j knows X , i.e., X contains a union of cells from j 's partition \mathcal{Q} . We can define our basic operation on partitions. Given a set X which is a union of partition cells from \mathcal{Q} , i updates \mathcal{P} at state w by splitting the cell $\mathcal{P}(w)$ into two sets: one that intersects X and the other intersects $W - X$ with the other partition cells remaining fixed.

In this paper we identify a piece of information with a set of states. Our basic refinement operation accepts a partition \mathcal{P} and a set X , and returns the partition refined by X .

Definition 5. Let \mathcal{P} be a partition on W and X a subset of W . The **information refinement** of \mathcal{P} , denoted $ref(\mathcal{P}, X)$ is defined as follows:

$$ref(\mathcal{P}, X) \stackrel{def}{=} \{P \cap X, P \cap (W - X) \mid P \in \mathcal{P}\} - \{\emptyset\}$$

In the above operation X can be any subset of W . In our framework, the set X represents some information known by an agent. Therefore, we say that $ref(\mathcal{P}, X)$ is an information refinement **based on** \mathcal{Q} , if X is a union of cells from \mathcal{Q} , i.e., $X = Q_1 \cup \dots \cup Q_l$, where each $Q_i \in \mathcal{Q}$ for $i = 1, \dots, l$.

Obviously, $ref(\mathcal{P}, X) \preceq \mathcal{P}$. Intuitively, if \mathcal{Q} is j 's partition, we think of X as being some information known by j and so $ref(\mathcal{P}, X)$ is result of agent i learning X .

Notice that in the above definition, the cells of \mathcal{Q} remain fixed. Thus the type of communication that takes place between i and j is rather impersonal. Suppose that i asks j whether a certain fact ϕ is true or false. As a matter of fact, suppose that ϕ is true and that j knows this. In this situation, not only does i 's information get updated, but so does j 's information, since j learned that i now knows ϕ . We assume, that instead of asking j directly whether ϕ is true, i is able to query j 's knowledge database in complete secrecy of j .

A vector of partitions \mathcal{P} represents the current state of information of all the agents. If some admissible communication between two agents i and j takes place, then \mathcal{P} is refined. Admissible communication between i and j simply

means that i and j are directly connected in the communication graph; and we will use the information refinement function defined above to state precisely *how* \mathcal{P} is refined by the communication.

Definition 6. Suppose that $\mathcal{G} = (A, E)$ is a communication graph. Let \mathcal{P} and \mathcal{P}' be two vectors of partitions on W . We say that \mathcal{P} is a one-step refinement of \mathcal{P}' , denoted by $\mathcal{P} \preceq_1 \mathcal{P}'$, if there exist $i, j \in A$ with $i \neq j$ and a set X such that X is a union of partition cells from \mathcal{P}_j , $(i, j) \in E$, $\mathcal{P}_k = \mathcal{P}'_k$ for all $k \neq i$ and $\mathcal{P}_i = \text{ref}(\mathcal{P}'_i, X)$.

This represents the situation described above, where i learns some information from j 's database of known facts, and the other agents, including agent j , are completely ignorant of this fact. Obviously if $\mathcal{P} \preceq_1 \mathcal{P}'$, then $\mathcal{P} \preceq \mathcal{P}'$, but not conversely.

Definition 7. Let \mathcal{G} be a communication graph and \mathcal{P} and \mathcal{P}' be two vectors of partitions on W . We say that \mathcal{P} is an **information refinement** of \mathcal{P}' , denoted $\mathcal{P} \preceq_{\mathcal{G}} \mathcal{P}'$ if there exists vectors $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$ such that $\mathcal{P}_1 = \mathcal{P}$, $\mathcal{P}_m = \mathcal{P}'$ and $\mathcal{P}_i \preceq_1 \mathcal{P}_{i+1}$, for $i = 1, \dots, m$.

Thus, $\preceq_{\mathcal{G}}$ is the reflexive, transitive closure of \preceq_1 .

Let $\langle W, \mathbb{P} \rangle$ be a multi-agent partition space. We think of the elements of \mathbb{P} as being the partitions that could possibly arise. Currently, \mathbb{P} can be *any* set of vectors of partitions. However, if we are given a communication graph we are only interested in the set of partitions that can arise from communication among agents respecting the communication graph. Given a multi-agent partition space $\langle W, \mathbb{P} \rangle$, we assume that there is a vector \mathcal{P}^0 that represents the agents' knowledge before any communication has taken place. \mathcal{P}^0 will be called the **initial vector**, and multi-agent partition spaces in which an initial vector is singled out will be called pointed multi-agent partition spaces. In this paper we will always assume that multi-agent partition spaces are pointed. We can now define the partition spaces that will be of interest to us in this paper.

Definition 8. Let \mathcal{G} be a communication graph and \mathcal{P} a vector of partitions on a set W . We say that $\mathbb{P}_{\mathcal{G}}$ is generated by \mathcal{G} from the initial partition \mathcal{P} if $\mathbb{P}_{\mathcal{G}}$ is the smallest set containing \mathcal{P} and for all $\mathcal{P}' \in \mathbb{P}_{\mathcal{G}}$ and all vectors \mathcal{P}'' , if $\mathcal{P}'' \preceq_{\mathcal{G}} \mathcal{P}'$, then $\mathcal{P}'' \in \mathbb{P}_{\mathcal{G}}$.

We say that a multi-agent partition space $\langle W, \mathbb{P} \rangle$ is generated by a communication graph \mathcal{G} when $\mathbb{P} = \mathbb{P}_{\mathcal{G}}$. In this paper, we will assume that any multi-agent space is generated from some \mathcal{P} by some communication graph.

Given a multi-agent partition space $\langle W, \mathbb{P} \rangle$ we can define the *downward closure* of a vector of partitions $\mathcal{Q} \in \mathbb{P}$:

$$\downarrow_{\mathbb{P}} \mathcal{P} \stackrel{\text{def}}{=} \{ \mathcal{Q} \mid \mathcal{Q} \in \mathbb{P} \text{ and } \mathcal{Q} \preceq_{\mathcal{G}} \mathcal{P} \}$$

When \mathbb{P} is clear from context, we may write $\downarrow \mathcal{P}$ instead of $\downarrow_{\mathbb{P}} \mathcal{P}$. If \mathbb{P} is generated from a communication graph \mathcal{G} , then we write $\downarrow_{\mathcal{G}} \mathcal{P}$, since in this case if $\mathcal{Q} \preceq_{\mathcal{G}} \mathcal{P}$, then $\mathcal{Q} \in \mathbb{P}$.

The type of refinement that we have described in this section is appropriate when modelling what agents know or come to know about the physical world. A more complex semantics will be needed to deal with what agents know about other agents' knowledge. The problem is that after some communication has taken place, the standard assumption that the partition structure is commonly known must be dropped. Consider the point of view of agent i . Agent i may learn some information from agent j , but in general will be unaware of communication between other agents in the communication graph. Thus agent i will be uncertain about the exact partition structure that represents the current situation. Moreover, while i may learn that j knows X , i is uncertain about *how* j came to know X , i.e., what questions did j ask and to whom. Of course, we assume that the communication graph is common knowledge, but uncertainty remains. Given a vector \mathcal{P} , when i updates with information X from agent j , there will be many different vectors of partitions compatible with j knowing X . Furthermore, this will be true *for each* agent.

A history based semantics can be used to deal with the general situation in which agents can have knowledge about other agents. We will give a brief sketch of some of the details in the next section. A complete discussion of this more general approach is reserved for the full version.

Example: Suppose there are three agents $\mathcal{A} = \{1, 2, 3, 4\}$ and suppose that the communication graph \mathcal{G} is the tree rooted at 1, where 1 has two children: 2 and 4, and 2 has only 3 as a child. Suppose that the initial partitions of the agents are given by the vector $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3, \mathcal{P}_4)$. Since neither 3 nor 4 are connected to any other agent, their partitions cannot change, i.e., for all $\mathcal{P}' \preceq_{\mathcal{G}} \mathcal{P}$, $\mathcal{P}'_2 = \mathcal{P}_2$ and $\mathcal{P}'_4 = \mathcal{P}_4$.

Since agent 1 is connected to all of the other agents, it is possible by asking enough questions, agent 1 can generate the partition, $\mathcal{P}_1 \sqcap \mathcal{P}_2 \sqcap \mathcal{P}_3 \sqcap \mathcal{P}_4$. However, since the only connection between agent 1 and agent 3 is through agent 2, any \mathcal{P}' such that $\mathcal{P}' \preceq_{\mathcal{G}} \mathcal{P}$ must reflect this fact. That is, if X is some information known only to agent 3, then there is no one step information refinement based on X of agent 1's partition. However, there is an information refinement in which 1's partition is updated with X *after* agent 2's partition is updated with X .

3 The Logic of Communication Graphs

Let Φ_0 be a countable set of propositional variables. $\mathcal{L}_0(\Phi_0)$ is the propositional (base) language based on Φ_0 . Let $\mathcal{L}_1(\Phi_0) = \{K_i\phi \mid \phi \in \mathcal{L}_0(\Phi_0), i \in \mathcal{A}\} \cup \mathcal{L}_0(\Phi_0)$. Finally let $\mathcal{L}_2(\Phi_0)$ be $\mathcal{L}_1(\Phi_0)$ closed under boolean combinations and \diamond . So formulas in $\mathcal{L}_2(\Phi_0)$ will not contain any embedded K_i operators, but may contain K_i embedded in a \diamond operator. Note that we are ruling out formulas of the form $K_i\diamond\phi$. However, this is not a significant restriction, since in any topologic for any formula $\phi \in \mathcal{L}_0(\Phi_0)$, $\Box\phi$ is equivalent to ϕ , i.e., no amount of effort can change the base facts about the world. We will not include Φ_0 as a parameter when it is not important.

Definition 9. A **multi-agent model** is a tuple $\langle \mathcal{G}, W, \mathbb{P}, v \rangle$ where \mathcal{G} is a communication graph, $\langle W, \mathbb{P} \rangle$ is a multi-agent partition space generated by \mathcal{G} and $v : \Phi_0 \rightarrow 2^W$ is a valuation function.

We can now define truth in a model. A **truth relation** $\models_{\mathcal{M}}$, where \mathcal{M} is a multi-agent model $\langle \mathcal{G}, W, \mathbb{P}, v \rangle$ is a subset of $(W \times \mathbb{P}) \times \mathcal{L}$ defined as follows (we write $w, \mathcal{P} \models_{\mathcal{M}} \phi$ instead of $((w, \mathcal{P}), \phi) \in \models_{\mathcal{M}}$).

1. $w, \mathcal{P} \models_{\mathcal{M}} p$ iff $w \in v(p)$
2. $w, \mathcal{P} \models_{\mathcal{M}} \neg\phi$ iff $w, \mathcal{P} \not\models_{\mathcal{M}} \phi$
3. $w, \mathcal{P} \models_{\mathcal{M}} \phi \wedge \psi$ iff $w, \mathcal{P} \models_{\mathcal{M}} \phi$ and $w, \mathcal{P} \models_{\mathcal{M}} \psi$
4. $w, \mathcal{P} \models_{\mathcal{M}} K_i\phi$ iff $\forall v \in \mathcal{P}_i(w), v, \mathcal{P} \models_{\mathcal{M}} \phi$
5. $w, \mathcal{P} \models_{\mathcal{M}} \Box\phi$ iff $\forall \mathcal{Q} \in \downarrow_{\mathbb{P}} \mathcal{P}, w, \mathcal{Q} \models_{\mathcal{M}} \phi$

Other propositional connectives are defined in the standard way. We abbreviate $\neg K_i \neg \phi$ and $\neg \Box \neg \phi$ as $L_i \phi$ and $\Diamond \phi$ respectively. We say ϕ is **valid** in \mathcal{M} if for all (w, \mathcal{P}) , $w, \mathcal{P} \models \phi$, denoted by $\models_{\mathcal{M}} \phi$. Since some of axioms will be given in terms of \Diamond , we state the definition of truth for \Diamond formulas

$$w, \mathcal{P} \models_{\mathcal{M}} \Diamond\phi \text{ iff } \exists \mathcal{Q} \in \downarrow_{\mathbb{P}} \mathcal{P}, w, \mathcal{Q} \models \phi$$

Thus the formula $\Diamond K_i \phi$ is interpreted as “There is a sequence of information refinements that results in agent i knowing ϕ .”

Axioms

1. All propositional tautologies
2. $(p \rightarrow \Box p) \wedge (\neg p \rightarrow \Box \neg p)$, for $p \in \Phi_0$.
3. $\Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$
4. $\Box\phi \rightarrow \phi$
5. $\Box\phi \rightarrow \Box\Box\phi$
6. $K_i(\phi \rightarrow \psi) \rightarrow (K_i\phi \rightarrow K_i\psi)$
7. $K_i\phi \rightarrow \phi$
8. $K_i\phi \rightarrow K_i K_i \phi$
9. $\neg K_i \phi \rightarrow K_i \neg K_i \phi$
10. $K_i \Box \phi \rightarrow \Box K_i \phi$

We include the following rules: modus ponens, K_i -necessitation and \Box -necessitation. We write $\vdash \phi$ if ϕ follows from any of the above schemes and rules. These axioms and rules are known to be sound and complete with respect to the set of all subset spaces ([MP]). Thus, they represent the core set of axioms and rules for any topologic. Soundness of axioms 1-8 and the rules are easy to verify.

Of course, technically, axiom 10 is not part of our language, since it contains a \Box embedded in a K_i operator. Nonetheless, we can show that this axiom is valid in our model. In fact this axiom will remain valid in the more general semantics defined below. For an example, we show that the mix axiom $K_i \Box \phi \rightarrow \Box K_i \phi$ is sound. It is easier to consider this in its contrapositive form: $\Diamond L_i \phi \rightarrow L_i \Diamond \phi$. This can be interpreted as if it is possible that agent i thinks ϕ is possible, then i thinks that it is possible that ϕ can be true.

Proposition 1. $\diamond L_i \phi \rightarrow L_i \diamond \phi$ is valid in all multi-agent models.

Proof. Suppose that $\mathcal{M} = \langle \mathcal{G}, W, \{\mathbb{P}_i\}_{i \in \mathcal{A}}, v \rangle$ is a multi-agent model, and suppose that $w \in W$ and \mathcal{P} is an arbitrary vector of partitions. Suppose that $w, \mathcal{P} \models \diamond L_i \phi$. Then there exists $\mathcal{Q} \preceq_{\mathcal{G}} \mathcal{P}$ such that $w, \mathcal{Q} \models L_i \phi$. So, there exists $v \in \mathcal{Q}_i(w)$ such that $v, \mathcal{Q} \models \phi$. Now since $\mathcal{Q} \preceq_{\mathcal{G}} \mathcal{P}$, $\mathcal{Q}_i(w) \subseteq \mathcal{P}_i(w)$, we have $v \in \mathcal{P}_i(w)$. But since $v, \mathcal{Q} \models \phi$, $v, \mathcal{P} \models \diamond \phi$. Hence $w, \mathcal{P} \models L_i \diamond \phi$. \square

Recall that given a vector \mathcal{P} , \mathcal{P}^I represents the implicit knowledge of all the agents. We can imagine that \mathcal{P}^I arises after all the agents discuss each and every fact known to each of them. This is of course assuming that any agent can access any other agent's knowledge database. If the agents communicate according to a communication graph, then it may not be possible to generate \mathcal{P}^I . However, it will be possible to generate a coarser partition $\mathcal{P}^{I, \mathcal{G}}$ which is based on the communication graph. Given an agent $i \in \mathcal{A}$, define $reach_{\mathcal{G}}(i)$ to be the set of all $j \in \mathcal{A}$ such that there is a path from i to j in \mathcal{G} (we may write $reach(i)$ when \mathcal{G} is understood). We can then define $\mathcal{P}^{I, \mathcal{G}}$ as follows: for each $i \in \mathcal{A}$,

$$\mathcal{P}_i^{I, \mathcal{G}} = \prod_{i \in reach_{\mathcal{G}}(i)} \mathcal{P}_i$$

Thus, $\mathcal{P}^{I, \mathcal{G}}$ arises if all the agents that can communicate according to the communication graph actually do communicate. So, $\mathcal{P}^{I, \mathcal{G}}$ is the vector that results if all the communication that *can* take place does take place. It is not hard to see that $\mathcal{P}^{I, \mathcal{G}}$ is a “lower bound” of \mathcal{P} in the set $\mathbb{P}_{\mathcal{G}}$ in the following sense. The following lemma follows easily from the definitions.

Lemma 1. Let \mathcal{G} be a communication graph and $\mathbb{P}_{\mathcal{G}}$ the set of partitions generated from \mathcal{G} . Then for any $\mathcal{P} \in \mathbb{P}_{\mathcal{G}}$ and for each $\mathcal{P}' \preceq_{\mathcal{G}} \mathcal{P}$ and $\mathcal{P}'' \preceq_{\mathcal{G}} \mathcal{P}$, we have $\mathcal{P}^{I, \mathcal{G}} \preceq_{\mathcal{G}} \mathcal{P}'$ and $\mathcal{P}^{I, \mathcal{G}} \preceq_{\mathcal{G}} \mathcal{P}''$

We can now show that the following scheme is also valid in all models.

$$\diamond \square \phi \rightarrow \square \diamond \phi$$

Proposition 2. $\diamond \square \phi \rightarrow \square \diamond \phi$ is valid in all multi-agent models.

Proof. Suppose that $\mathcal{M} = \langle \mathcal{G}, W, \{\mathbb{P}_i\}_{i \in \mathcal{A}}, v \rangle$ is a multi-agent model, and suppose that $w \in W$ and \mathcal{P} is an arbitrary vector of partitions. Suppose that $w, \mathcal{P} \models \diamond \square \phi$. Then there is a refinement $\mathcal{Q} \preceq_{\mathcal{G}} \mathcal{P}$ such that $w, \mathcal{Q} \models \square \phi$. Let $\mathcal{R} \in \downarrow \mathcal{P}$. We must show $w, \mathcal{R} \models \diamond \phi$. By the lemma 1, $\mathcal{P}^{I, \mathcal{G}} \preceq_{\mathcal{G}} \mathcal{Q}$ and $\mathcal{P}^{I, \mathcal{G}} \preceq_{\mathcal{G}} \mathcal{R}$. Therefore, $w, \mathcal{P}^{I, \mathcal{G}} \models \phi$; and so $w, \mathcal{R} \models \diamond \phi$. \square

Before showing the connection between valid formulas and the communication graphs, we will discuss some of the details for a semantics for the more general case in which we can express the knowledge which agents have about other agents' knowledge. Assume that an event is a query of a database. Formally we can define an event as a tuple (ϕ, i, j) to mean that i learns information

ϕ from j , where ϕ is a base formula (an element of $\mathcal{L}_0(\Phi_0)$). Of course there must be an edge between i and j in the communication graph. A history is a finite sequence of events. Thus a history represents a particular sequence of question and answers. Assume that initially, nature informs each agent of the truth value of a particular set of propositional variables. This generates an initial vector of partitions, say \mathcal{P}^0 . Now given any history there is a vector of partitions that is generated by that sequence of questions starting from the initial partition. Let $Part(H)$ be the vector of partitions generated by history H from initial vector \mathcal{P}^0 . Truth in this model will be defined at a state w and a history H . Truth of propositional variables is independent of the history, so $w, H \models p$ iff $w \in V(p)$, where V is some valuation function. Boolean connectives are obvious. Given two histories H and H' , suppose that $H \preceq H'$ iff H' extends H , i.e., H' is H concatenated with some event. Then,

$$w, H \models \diamond\phi \text{ iff } \exists H', H \preceq H' \text{ and } w, H' \models \phi$$

Given a history H , let $\lambda_i(H)$ be i 's local history. I.e., this is a sequence of events that i can "see". Formally λ_i maps each event of the form (ϕ, i, j) to itself and other events to the null string. Then define $H \sim_i H'$ iff $\lambda_i(H) = \lambda_i(H')$. We can now define truth of a knowledge formula:

$$w, H \models K_i\phi \text{ iff } \forall H' \sim_i H, \forall v \in (Part(H')_i)(w), v, H' \models \phi$$

This definition addresses both causes of i 's uncertainty: 1. i 's uncertainty about the partition representing the information of all the agents and 2. i 's uncertainty about the current state.

4 Connection with Communication Graphs

In this section we will investigate the close connection between formulas valid in a model based on the communication graph and the communication graph. We will prove our main technical claim that the valid formulas characterize the communication graph.

Let ϕ be a formula in our language \mathcal{L} , and consider the formula $K_i\phi \rightarrow \diamond K_j\phi$. Intuitively, this formula says that if i knows ϕ then it is possible for agent j to know ϕ . One would expect that this formula will always be true provided that j is connected to i in the communication graph. However, this does not quite work for *any* formula ϕ . For example, let ϕ be the formula $p \wedge \neg K_j p$, where $p \in \Phi_0$. Suppose that j is connected to i in some communication graph \mathcal{G} . It is easy to construct a model in which $K_i(p \wedge \neg K_j p)$ is true at some pair (w, \mathcal{P}) . However, no pair (w, \mathcal{P}') with $\mathcal{P}' \preceq_{\mathcal{G}} \mathcal{P}$ can satisfy the formula $K_j(p \wedge \neg K_j p)$, since K_j is an **S5** modal operator.

Nonetheless, there is a certain class of formulas for which the above statement will hold.

Definition 10. ϕ is stable in \mathcal{M} iff $\phi \rightarrow \square\phi$ is valid in \mathcal{M} .

We say that ϕ is stable if ϕ is stable in all models. If ϕ is a ground formula, i.e., $\phi \in \mathcal{L}_0$, then ϕ is stable. This is easy to see, since using axiom 1 and 2, one can show that if $\phi \in \mathcal{L}_0$, then $\vdash \phi \leftrightarrow \Box\phi$.

At this point, it is worth pointing out that we are assuming that the all the agents share the same language. That is all of the agents are aware of the entire set Φ_0 of propositional letters, and so it is possible that any agent can learn any well-formed formula. This assumption can be relaxed in order to deal with situations in which agents only partially share a language. Technically, we need only restrict the sets X that can be used in definition 6 to show that $\mathcal{P}' \preceq_1 \mathcal{P}$. The only sets X that can be learned from agent j by agent i are the sets that are *definable* in i 's language. Even if agent i and j share the same language, agent j might not want agent i to have access to certain formulas.

Lemma 2. *Let \mathcal{G} be a communication graph and \mathcal{M} a model generated by \mathcal{G} . If ϕ is stable in \mathcal{M} and there is a path from j to i in the communication graph, then $K_i\phi \rightarrow \Diamond K_j\phi$ is valid in \mathcal{M} .*

Proof. Suppose that $\mathcal{M} = \langle \mathcal{G}, W, \{\mathbb{P}_i\}_{i \in \mathcal{A}}, v \rangle$ is a multi-agent model and ϕ is stable in \mathcal{M} . Suppose that $w, \mathcal{P} \models K_i\phi$. Then for all $v \in \mathcal{P}_i(w)$, $v, \mathcal{P} \models \phi$. We must show that there is a \mathcal{Q} such that $\mathcal{Q} \preceq_{\mathcal{G}} \mathcal{P}$ and $w, \mathcal{Q} \models K_j\phi$. For simplicity, we will first assume that there is an edge between j and i in \mathcal{G} . Let X be the union of \mathcal{P}_i cells in which ϕ is true, i.e., $X = P_1 \cup \dots \cup P_m$ where for all $k = 1, \dots, m$, $P_k \in \mathcal{P}_i$ and for all $v \in P_k$, $v, \mathcal{P} \models \phi$. Define \mathcal{Q} to be the vector which is exactly like \mathcal{P} except in the j th position, replace \mathcal{P}_j with $ref(\mathcal{P}_j, X)$. Since there is an edge between j and i , $\mathcal{Q} \preceq_1 \mathcal{P}$, and so $\mathcal{Q} \preceq_1 \mathcal{P}$. Let v be any element in $\mathcal{Q}_j(w)$, then by construction, $v \in \mathcal{P}_j(w) \cap X$. Since $v \in X$, $v, \mathcal{P} \models \phi$; and therefore since ϕ is stable, $v, \mathcal{Q} \models \phi$. Hence, $w, \mathcal{Q} \models K_j\phi$.

If j and i are connected instead of directly connected the result is an easy extension of the above proof. Suppose that the path from j to i goes through the agents i_1, \dots, i_k . Then agent using the X defined above, we can define a sequence vectors in which i_1 learns X from i , i_{m+1} learns X from i_m for $m = 1, \dots, k-1$, and j learn X from i_k . \square

In fact we can show something stronger, that the communication graph is characterized by formulas valid in models based on the graph.

Theorem 1. *Let $\mathcal{G} = (\mathcal{A}, E)$ be a communication graph. Then $(i, j) \in E$ if and only if, for all $l \in \mathcal{A}$ such that $l \neq i$ and $l \neq j$ and all stable ϕ , the scheme*

$$K_j\phi \wedge \neg K_l\phi \rightarrow \Diamond(K_i\phi \wedge \neg K_l\phi)$$

is valid in all models generated by \mathcal{G} .

We leave the details for the full version and sketch the proof. If $(i, j) \in E$ for some communication graph, then Lemma 2 shows that provided ϕ is stable, then $K_j\phi \rightarrow \Diamond K_i\phi$ will be valid in any model based on \mathcal{G} . It is not hard to see that the above proof can be adapted to show that $K_j\phi \wedge \neg K_l\phi \rightarrow \Diamond(K_i\phi \wedge \neg K_l\phi)$

is valid in all models for $l \neq i$ and $l \neq j$. If $w, \mathcal{P} \models K_i\phi \wedge \neg K_l\phi$ then there is a $v \in \mathcal{P}_i(w)$ such that $v, \mathcal{P} \models \neg\phi$, then using the fact that the refinement \mathcal{Q} defined in the proof of Lemma 2 does not change any partition other than j 's, we can show that $w, \mathcal{Q} \models K_i\phi \wedge \neg K_l\phi$. For the other direction, if $(i, j) \notin E$, then either i and j are not connected or there is a path going through some agent l that connects i and j . In the first case, it is easy to construct a model based on \mathcal{G} in which K_ip is true for some propositional variable p in some situation (w, \mathcal{P}) and also that K_jp is false in the same situation. If we assume that i and j are the only two agents, then it is easy to see that no refinement of \mathcal{P} can result in j knowing p . In the other case, if there is a path from i to j going through l , then any refinement that increases i 's knowledge must also increase l 's knowledge, and so the above formula will not be valid.

5 Conclusion

In this paper we have introduced a logic of knowledge and communication. Communication among agents is restricted by a communication graph, and idealized in the sense that the agents are unaware when their knowledge base is being accessed. We have shown that the communication graph is characterized by the validities of formulas in models based on that communication graph.

Related Work: This paper fits in with a growing body of work on social software ([Pa]). One of the main goals of the social software research program is to develop mathematical tools that can be used to study social procedures. Other work that falls into this category is [PR] which studies the semantics of messages and [PaPaC] which studies a logic of knowledge with obligation.

In this paper, we have presented a logic of multi-agent knowledge with an update operator. Similar logics have been studied starting with [PI] and more recently in [BM,K,Vd,Ge]. In chapter 4 of [K], Kooi provides an excellent overview of the current state of affairs of these dynamic epistemic logics. We do not consider general epistemic updates as is common in the literature, but rather study a specific type of epistemic update and its connection with a communication graph.

Further Work: We suspect that the logic of communication graphs has the finite model property and so is decidable. We leave the proof for further investigation. Other standard questions such as a complete axiomatization will also be studied. Another interesting extension would be to allow different types of updates, such as lying, conscious updates, updating to subgroups and so on.

Finally we remark that this logic can be seen as a demonstration for the need for cryptographic protocols. Two issues are important here. This first is that an agent may only want part of its knowledge base to be accessible by the public. This may be modeled in our framework by attaching to each agent a set of formulas that are in the public domain, and so when i is directly connected to j , i can only update by sets definable in the publicly accessible language. The second issue is that we may not know the exact structure of the communication graph. For example, if Ann accesses some information from Bob's website, but

unknown to Ann, Charles is listening in, then the communication graph does not have an edge between Ann and Bob, but only a path from Ann to Bob going through Charles. Then clearly as a condition for Ann learning some information from Bob, Charles must become informed of that same piece of information. Thus cryptographic protocols essentially ensure that there are direct edges between agents in the communication graph.

References

- [BM] Baltag, A. and Moss, L., Logics for Epistemic Programs, to appear in the *Knowledge, Rationality, and Action* section of *Synthese*.
- [DMP] Dabrowski, A, Moss, L, and Parikh, R. Topological reasoning and the logic of knowledge. *Annals of Pure and Applied Logic*, **78**, (1996), pp. 73 - 110.
- [G93] Georgatos, K, Modal Logics for Topological Spaces. PhD Dissertation. Graduate School and University Center. City University of New York, 1993.
- [G94] Georgatos, K, Knowledge Theoretic Properties of Topological Spaces. In *Knowledge Representation and Uncertainty*. M. Masuch and L. Polos, Eds. Lecture Notes in Artificial Intelligence, vol. 808, pages 147-159, Springer-Verlag, 1994.
- [G97] Georgatos, K, Knowledge on Treelike Spaces. *Studia Logica*, **59**, (1997), pp. 271 - 231.
- [Ge] Gerbrandy, J., *Bisimulations on Planet Kripke*, Ph.D. dissertation, University of Amsterdam, 1999.
- [H99] Heinemann, B., Temporal Aspects of the Modal Logic of Subset Spaces, *Theoretical Computer Science*, **224(1-2)**:135-155, 1999.
- [H00] Heinemann, B., Extending Topological Nexttime Logic. In S. D. Goodwin, A. Trudel, editors, *Temporal Representation and Reasoning*, TIME-00, Cape Breton, Nova Scotia, Canada, pages 87-94, IEEE Computer Society Press, Los Alamitos, CA, 2000.
- [H02] Heinemann, B., A Hybrid Treatment of Evolutionary Sets. In C. A. Coello Coello, A. de Albornoz, L. E. Sucar, O. Cair Battistutti, editors, MICAI'2002: *Advances in Artificial Intelligence*, Mrida, Yucatn, Mexico. Volume 2313 of Lecture Notes in Artificial Intelligence, pages 204-213, Springer, Berlin, 2002.
- [K] Kooi, B., *Knowledge, Chance, and Change*, Ph.D. dissertation, University of Groningen, 2003.
- [H04] Heinemann, B., A Hybrid Logic of Knowledge Supporting Topological Reasoning. In *Algebraic Methodology and Software Technology*, AMAST 2004, Stirling, United Kingdom. Lecture Notes in Computer Science, Springer, Berlin, 2004. *To appear*.
- [MP] Moss, L. and Parikh, Topological Reasoning and the Logic of Knowledge, *TARK IV*, Ed. Y. Moses, Morgan Kaufmann, 1992.
- [Pa] Parikh, R., Social Software, *Synthese*, **132: 3**, Sep 2002, pp. 187-211.
- [PaPaC] Parikh, R., Pacuit, E. and Cogan, E., The logic of knowledge based obligation. Revised version to be presented at DALT '04.
- [PR] Parikh, R. and Ramanujam, R., A knowledge based semantics of messages, in *J. Logic, Language, and Information*, **12**, pp. 453 - 467, 2003.
- [P] Plaza, J., Logics of public communications, *Proceedings, 4th International Symposium on Methodologies for Intelligent Systems*, 1989.
- [Vd] van Ditmarsch, H., *Knowledge Games*, Ph.D. dissertation, University of Groningen, 2000.

- [V] Vickers, S. *Topology Via Logic*, Cambridge University Press. 1989.
- [WP] Weiss, M. A. and Parikh, R., "Completeness of Certain Bimodal Logics of Subset Spaces", *Studia Logica*, **71:1**, pp. 1 - 30, 2002.