Known Unknowns:
Time bounds and Knowledge of Ignorance

Ido Ben-Zvi
Department of Computer Science
Technion—Israel Institute of Technology
Haifa 32000, Israel
idobz@cs.technion.ac.il

Yoram Moses
Department of Electrical Engineering
Technion—Israel Institute of Technology
Haifa 32000, Israel
moses@ee.technion.ac.il

ABSTRACT

This paper studies the role that known bounds on message transmission times in a computer network play on the evolution of the epistemic state over time. A connection to cones of causal influence analogous to, and more general than, light cones is presented. Focusing on lower bounds on message transmission times, an analysis is presented of how knowledge about when others are guaranteed to be ignorant about an event of interest ("knowing that they don't know") can arise. This has implications in competitive settings, in which knowing about another's ignorance can provide an advantage.

The Happiness of Fish

Chuangtse and Hueitse had strolled on to the bridge over the Hao, when the former observed:
"See how the small fish are darting about!
That is the happiness of the fish."
"You not being a fish yourself, said Hueitse,
"how can you know the happiness of the fish?"
"And you not being I," retorted Chuangtse,
"how can you know that I do not know?"

Chuangtse, circa 300 B.C.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Distributed Networks; F.4.1 [**Mathematical Logic**]: Modal Logic; H.1.1 [**Systems and Information Theory**]: General Systems Theory

General Terms

Distributed Systems, Knowledge Theory

Keywords

Knowledge, Ignorance, Time bounds, Synchronous communication

1. INTRODUCTION

Knowledge about whether others know particular facts plays an important role in strategic situations [1] as well as in the design of effective programs and plans in multi-agent systems [6]. In his book *The Rothschilds* [13], Frederic Morton describes an account of the events in the London Stock exchange at the time of the Battle of Waterloo, in which interactive knowledge, and specifically knowledge about ignorance, supposedly played a major role. Morton's (disputed) account can be summarized as follows: On the night of June 15, 1815 Nathan Rothschild, one of London's most prominent financiers at the time, was informed by his special private couriers that the Battle of Waterloo was won by the British. Official word by Wellington's men could only arrive on the next day. On the next morning, Rothschild went to the London Stock exchange, and signaled his agents to furiously sell consuls (government bonds). "*He knows...[who won]*" was the word among traders. The market crashed, and just before Wellington's men arrived with the news of victory, Rothschild signaled his agents to buy all available consuls, at a fraction of their original price. He is said to have made a fortune on that day.

This story is a classical example in which asymmetric information provides an essential advantage in real-life situations. For the course of events described by Morton to be plausible, not only was it necessary for Rothschild to know about the battle's outcome before everyone else. He also knew that the others were *ignorant* of the outcome. Otherwise, he would fear that one of his rivals could out-smart him, gradually buy his shares and make out with a huge gain at Rothschild's expense. The impact of the asymmetric information was supposedly further magnified by the fact that, since Rothschild was known to have a separate courier service, the others knew that he *did* know who won.

An important element in determining the epistemic circumstances in this example comes from the known properties of Rothschild's courier system, and those of Wellington's communication lines.

In this paper we study the interaction between known bounds on the transmission times—the time it takes messages to traverse a channel—and interactive knowledge. Our focus here will be on the impact that lower bounds on transmission times have on when one player can know that another is ignorant of a fact of interest. Our analysis will start by presenting a novel view of how such bounds in a communication network induce *causal cones* of information flow among events in the system, in analogy with the light cones in Einstein-Minkowski spacetime considered in physics [8, 15].[1] Lamport draws an analogy between communication in a computer network and light cones in his seminal paper [11], based on his notion of potential causality among events. We now briefly review Lamport's setting. For ease of exposition, however, we will state potential causality as a relation on agent-time nodes $\langle i,t\rangle$ rather than events. We think of such a node as referring to agent i's site at time t. Nodes play a role analogous to that of a point in the Einstein-Minkowski spacetime. We typically denote nodes by Greek letters $\theta,\theta',\ldots$. Lamport's relation is defined with respect to a given run (i.e., history) r:

Definition 1. **(Lamport causality).** Fix a run r. *Lamport's causality* is the smallest relation $\twoheadrightarrow$ over nodes satisfying the following three conditions:

1. If $\theta=\langle i,t\rangle$ and $\theta'=\langle i,t'\rangle$ for the same agent i, and $t\le t'$, then $\theta\twoheadrightarrow\theta'$;
2. If some message is sent at θ in r and received at θ' then $\theta\twoheadrightarrow\theta'$; and
3. If $\theta\twoheadrightarrow\theta'$ and $\theta'\twoheadrightarrow\theta''$, then $\theta\twoheadrightarrow\theta''$.

Intuitively, $\theta\twoheadrightarrow\theta'$ holds between two nodes of different processes if there is a message chain starting at the first process at or after θ and reaching the second process no later than θ'. In an asynchronous system, in which there are no timing guarantees and agents do not have access to clocks, an event at θ can causally affect events at θ' only if $\theta\twoheadrightarrow\theta'$. Thus intuitively, in a run r, the *cone of causal influence* of a node θ is defined by the set $\mathsf{fut}(\theta)=\{\theta':\theta\twoheadrightarrow\theta'\}$.

In an influential paper [4], Chandy and Misra showed that information flow and knowledge gain in asynchronous systems are governed by Lamport's '$\twoheadrightarrow$' relation. Theorem 2 below is a variant of their Knowledge Gain theorem. The simplest implication of their analysis can intuitively be restated as follows:

THEOREM 1 ([4]). *In an asynchronous system, if event e occurs at $\theta=\langle i,t\rangle$ in r and j knows at $\theta'=\langle j,t'\rangle$ that e occurred, then $\theta\twoheadrightarrow\theta'$.*

As the paper's title suggests, we are interested in knowledge about ignorance. In fact, we offer a partial answer to Chaungtse's question *"how do you know that I do not know?"*, albeit in a different setting. We observe that in the presence of clocks, lower bounds on communication time can bring about knowledge about ignorance. Essentially, such knowledge is obtained by considering the contrapositive form of Theorem 1.

Adding Clocks.

In many settings agents have access to clocks, and there are guarantees about how long various activities may take. In particular, guarantees are often available for the *transmission times* in the communication network—the time it takes for messages to be transmitted over each given channel. In recent work we studied the role that upper bounds on transmission times play in knowledge gain and coordination [2, 3]. Upper bounds were shown to play a central role in determining how and when nested knowledge and common knowledge can be attained. In this paper we extend the analysis and consider lower bounds on message transmission times. These provide an *a priori* guarantee on the nodes that will *not* be causally affected by θ. This is initially a co-cone (the complement of a cone), which can grow when messages are actually delivered. Specifically, if a message carrying information about θ arrives at some node later than the earliest time allowed by the lower bound, then the co-cone may grow. Nodes in the co-cone for θ will be ignorant about any spontaneous events that may have occurred at θ. As we show in this paper, by analyzing the structure of this co-cone, it is possible to determine when an agent can know about other agent's ignorance.

The Battle of Waterloo example above illustrates the importance of knowledge about other's ignorance in particular circumstances. For another example, consider a sealed-bid first-price auction for mining rights. Suppose that near the auction closing a potential bidder learns of a relevant event e, say that gold was found in an adjacent site. The bidder's valuation of the auctioned rights may have changed. But the decision regarding if, and by what amount, to alter her bid would depend on her knowledge about whether her competitor knows about e. In particular, if she knows that he is ignorant of e, then she should not increase her bid by a significant amount. Our analysis will serve to show how our favored bidder can use her information about transmission times to figure out whether her competitor is ignorant of e.

The main contributions of this paper are:

- We consider synchronous settings in which both upper bounds, lower bounds, or both types of bounds are available for each channel in the network.
- A rich structure of cones of influence and of non-influence is discovered. Whereas in the classical light-cone view the causal cone of a point in spacetime depends only on the point itself, upper bounds provide a region of points that are guaranteed *a priori* to be affected by a given point and lower bounds yield a region guaranteed to be unaffected by it. Both regions grow with time, reducing the initial indeterminacy. The actual causal map that is ultimately realized in a given run is obtained based on the realized communications.
- Notions underlying the information that agents have about the region that is guaranteed to be causally unaffected by a given event are developed, and used to characterize when knowledge of ignorance is attained.
- This work is another illustration of how assumptions about the communication medium determine the interaction between knowledge and the flow of information in computer networks.

This paper is organized as follows:
The next section presents our model of computation which is based on the interpreted systems approach of Fagin et al. [6]. Section 3 presents the interaction between transmission bounds and causal cones, showing how upper bounds affect causality, and how lower bounds do. In light of Section 3, Section 4 presents a characterization of knowledge of ignorance in the full-information protocol, given lower bounds on communication times. Finally, Section 5 provides a summary and conclusions.

[1] Our setting can be thought of as consisting of a single inertial system, in which there is a single, non-relativistic, notion of time for all sites.

2. PRELIMINARY DEFINITIONS

2.1 The Model

The paper follows the *interpreted systems* approach of Fagin *et al.* [6] to modeling multi-agent systems. Two essential building blocks are used to define the formal model. The *context* in which the agents operate, and the agents' *protocols*, which determine their behavior. A context γ and a protocol profile $P = (P_1, \ldots, P_n)$ for the agents define a unique system $\mathcal{R} = \mathcal{R}(P, \gamma)$, consisting of all possible *runs*, or histories, of P in the context γ. A run is defined as a function $r : \mathbb{N} \to \mathcal{G}$ from time (taken here to be the natural numbers) to the set of global states. Every global state $r(t)$ is identified with a tuple of local agent states, as well as a state for the *environment*, which accounts for all aspects of the system that are not a part of some agent's local state, e.g., messages in transit.

We focus our attention on three particular variants of *synchronous contexts* γ^{min}, γ^{max} and γ^{b}. Their main properties are highlighted below.

- The set of agents is denoted by $\mathbb{P} = \{1, \ldots, n\}$. In any given execution, the network is a fixed edge-labeled graph $\mathsf{Net} = (\mathbb{P}, E, w)$, where E is a set of directed edges determining the channels in the network, and w_{ij} is a label defined for every channel $(i, j) \in E$. In γ^{min} the labels are natural numbers, denoted by $min_{ij} > 0$ for every channel $(i, j) \in E$. In γ^{max} the labels are similar; they are denoted by max_{ij} and specify upper bounds on transmission times. Finally, in γ^{b} each label is a pair (min_{ij}, max_{ij}) specifying both upper *and* lower bounds. The network does not change throughout an infinite run, and a copy of the labelled network is part of every agent's local state at all times. As a result, the properties of communication are common knowledge at all times.

- Each agent's local state also contains a correct current reading of the time. As a result, the agents are assumed to share a global clock.

- Events are message sends and receives, external inputs, as well as internal computations performed by the agents. All events in a run are distinct, and each event, whether it is realized or not in a run, is associated with a particular agent.

- The environment agent is in charge of choosing the external inputs, and of determining message transmission times. External inputs are determined in a genuinely nondeterministic fashion, and are not correlated with anything that comes before in the execution, or with external inputs of other agents. Message deliveries are also nondeterministic, but respect (i) exactly one delivery per message, and (ii) transmission times do not go lower than the network lower bounds, or higher than the upper bounds.

- An event is considered *spontaneous* if it is either an external input, or if it is a message delivery, except when delivery is at the specified *upper bound* on transmission time for the channel.[2] (In γ^{min} all deliveries are spontaneous events.)

- For simplicity, the agents follow deterministic protocols. Hence, a given protocol P for the agents and a given behavior of the environment completely determine the run. A special role will be played by the *full-information protocol*, denoted by fip. In this protocol, every agent has perfect recall, and it sends a message describing its whole history to all neighbors at every time instant.

2.2 Syntax and Semantics

Our logical language is very simple, yet quite unusual in its focus on time stamping. The set Φ of primitive propositions consists of the propositions $\mathsf{occurred}_t(e)$ for all events e and times t, and the propositions $\theta \mapsto \theta'$ for all pairs of agent-time nodes. The logical language $\mathcal{L}$ is obtained by closing Φ under propositional connectives and knowledge formulas. We write $\theta \not\mapsto \theta'$ instead of $\neg(\theta \mapsto \theta')$. Our knowledge operators are indexed by a node $\theta = \langle i, t \rangle$, and so are time stamped. Thus, $\Phi \subset \mathcal{L}$, and if $\varphi \in \mathcal{L}$ and $\theta = \langle i, t \rangle$, then $K_\theta \varphi \in \mathcal{L}$. In this case, the formula $K_\theta \varphi = K_{\langle i,t \rangle} \varphi$ is read *agent i at time t knows φ*.

Since all formulas are time stamped, their truth is time-invariant, and they become properties of the run. Consequently, we define the truth of a formula with respect to a pair (R, r). We write $(R, r) \vDash \varphi$ to state that φ holds in the run r, with respect to system R. We write $r \sim_\theta r'$ for $\theta = \langle i, t \rangle$ if agent i's local state at time t in r is identical to its local state at time t in run r', and inductively define

$$(R, r) \vDash \theta \mapsto \theta' \quad \text{iff} \ \ \theta \twoheadrightarrow \theta' \text{ in the run } r;$$

$$(R, r) \vDash \mathsf{occurred}_t(e) \quad \text{iff the event } e \text{ occurs in } r \text{ by time } t;$$

$$(R, r) \vDash K_\theta \varphi \quad \text{iff } (R, r') \vDash \varphi \text{ for every run } r' \sim_\theta r.$$

Propositional connectives are handled in the standard way, and their clauses are omitted above. Despite the slight variance in nomenclature, $K_{\langle i,t \rangle} \varphi$ still conforms with standard usage as advocated in [6], by being satisfied if φ holds at all points at which i has the same local state as it does at time t. Thus, given R, an agent's knowledge at any given instant is determined by its local state.

3. BOUNDED COMMUNICATION AND CONES OF INFLUENCE

Consider a fixed inertial system in which all sites are at rest with respect to each other. In such a setting, light rays carry information at a constant speed c in Euclidean space. In terms of Einstein-Minkowski spacetime, the light rays outgoing from an event (or a 4-dimensional point p) form a surface in spacetime that is called the event's *future light cone*. The light rays converging on an event form a surface called the event's *past light cone*. The spacetime points within p's future light cone make up its *absolute future* and those within its past light cone make up its *absolute past*: the former are spacetime points that events at p can influence and the latter are the points that can influence p. Events at points outside both light cones of p can neither influence nor be influenced by events at p. Such events are considered *independent of*, or *concurrent with*, events at p.

[2]Consider a message sent from i to j at time t with upper bound b_{ij}. At time $t + b_{ij} - 1$, if the message is not delivered, then the combined (distributed) knowledge of i and j implies that it will be delivered in the next round. The delivery is thus not spontaneous with respect to the agents' knowledge.

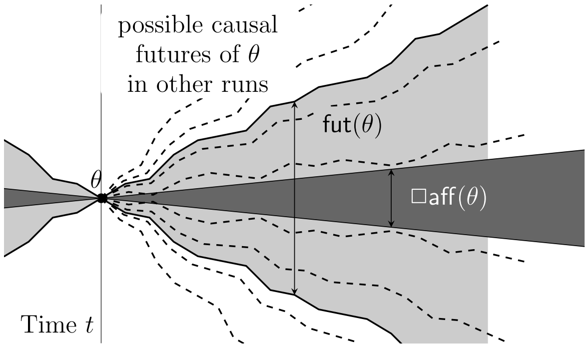

Figure 1: The future causal cones of θ in asynchronous settings and in γ^{max}

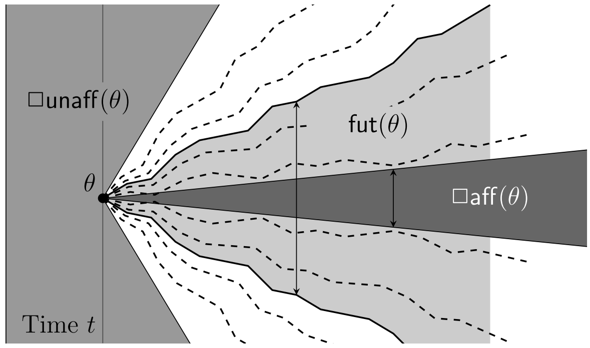

Figure 2: The future causal cones of θ in γ^{b}

Observe that the absolute future and absolute past cones of a point p are fixed and depend only on the coordinates of p. In analogy, consider a computer network in which, for every channel $(i, j) \in E$, there is a fixed transmission time trans_{ij}. Moreover, assume that the agents follow a *full-information* protocol fip in which, at every instant, they send a message describing their whole history to all neighbors. Just as in the case of light traveling in Einstein-Minkowski spacetime, in this setting every node θ would define a future cone $\mathsf{fut}(\theta)$ and a past cone $\mathsf{past}(\theta)$, as well as a region of nodes that are causally independent, or concurrent, with respect to θ. In this section we focus primarily upon future causality, where an intricate dynamics transforms potentiality into necessity, as we shall soon see.

What happens when transmission times are *not* fixed? In purely asynchronous settings, in which messages can take arbitrarily long to be delivered, a node θ' can be influenced by $\theta = \langle i, t\rangle$ only if $\theta \twoheadrightarrow \theta'$. Thus, Lamport's $\twoheadrightarrow$ relation defines a future cone (and a past cone) for every given node. In contrast to the fixed-transmission system described above, however, here the cone may differ significantly between different runs due to the varying transmission times. Figure 1 shows $\mathsf{fut}(\theta)$, the future cone of node θ in a specific run, for an outside observer with complete information about the full execution. The alternative futures that remain unrealized in the current run are shown in outline. Observe that a "core" cone can be made out in the center of $\mathsf{fut}(\theta)$, of nodes that are guaranteed *a priori* to be within $\mathsf{fut}(\theta)$, and will thus necessarily be affected by θ. We denote this cone by $\Box\mathsf{aff}(\theta)$. In an asynchronous context, the only nodes guaranteed to be in $\mathsf{fut}\langle i, t\rangle$ are $\Box\mathsf{aff}(\langle i, t\rangle) = \{\langle i, t'\rangle : t' \geq t\}$; i.e., this cone consists of all nodes on i's timeline that follow θ.

The picture becomes more interesting in the presence of upper bounds max_{ij} on message transmission times. Denote by D_{ih} the shortest distance between vertices i and h in the max-weighted network graph. Under the fip described above, we are guaranteed to have $\langle i, t\rangle \twoheadrightarrow \langle h, t'\rangle$ whenever $t' \geq t + D_{ih}$. Thus, maximal transmission times extend the inner cone into $\Box\mathsf{aff}(\langle i, t\rangle) = \{\langle j, t'\rangle : t' \geq t + D_{ij}\}$. As in the asynchronous case, for every run $r \in \mathcal{R}(\mathsf{fip}, \gamma^{max})$ and node θ, necessarily $\mathsf{fut}(\theta) \supseteq \Box\mathsf{aff}(\theta)$, as messages that are delivered earlier than at the upper bounds on a channel introduce into $\mathsf{fut}(\theta)$ nodes that were not guaranteed *a priori* to be in $\Box\mathsf{aff}(\theta)$.

We may also consider the set $\Box\mathsf{unaff}(\theta)$, counterbalancing $\Box\mathsf{aff}(\theta)$, and consisting of nodes that are necessarily unaffected causally by $\theta = \langle i, t\rangle$. As long as no lower bounds are defined, this set consists of all nodes in θ's *temporal* past, as well the nodes $\langle j, t\rangle$ where $j \neq i$.[3]

When we move to the contexts γ^{min} or γ^{b}, in which there are lower bounds on transmission times, $\Box\mathsf{unaff}(\theta)$ gets a richer structure. Lower bounds on transmission play a related, albeit somewhat different role than that of upper bounds. Suppose that a spontaneous event e takes place at $\theta = \langle i, t\rangle$ and that, based on the lower bounds, the fastest that communication from i can reach j is d_{ij}.[4] If $\theta' = \langle j, t'\rangle$ where $t' < t + d_{ij}$, then events at θ' cannot be causally influenced by e. It follows that the $\Box\mathsf{unaff}(\langle i, t\rangle)$ region is now defined as the set $\{\langle j, t'\rangle | t' < t + d_{ij}\}$. Figure 2 shows the causal cones of θ in γ^{b}.

We have considered the sets $\mathsf{fut}(\theta), \Box\mathsf{aff}(\theta)$ and $\Box\mathsf{unaff}(\theta)$, which are all easily determined given complete information regarding the run's infinite execution. To be of practical use however, we should consider whatever it is that can be made known about causal influence, given the execution up to a specific "present" point in time t' and assuming that future events in the run are as yet undetermined. We define $\mathsf{fut}(\theta, t')$ as the set $\{\langle j, t''\rangle | \theta \twoheadrightarrow \langle j, t''\rangle$ and $t'' \leq t'\}$, the set of nodes that have, by time t', already been realized as a part of θ's future.

The realized portion of the run by time t' determines the sets of necessarily affected and unaffected nodes relative to the current time, in a way that extends them beyond $\Box\mathsf{aff}(\theta)$ and $\Box\mathsf{unaff}(\theta)$, respectively. The set $\Box\mathsf{aff}(\theta, t')$ of all nodes that are guaranteed to be causally affected by θ given $\mathsf{fut}(\theta, t')$, is the union of the $\Box\mathsf{aff}(\theta')$ cones of all $\theta' \in \mathsf{fut}(\theta, t')$. We denote by $\Diamond\mathsf{unaff}(\theta, t')$ the set of nodes that are potentially unaffected by θ relative to current time t'. This set is the complement of the set $\Box\mathsf{aff}(\theta, t')$.

A more challenging definition is that of the set $\Diamond\mathsf{aff}(\theta, t')$ of nodes that, at time t', are potentially affected by θ. A node θ' is potentially affected if it is possible, given $\mathsf{fut}(\theta, t')$, that the current run will evolve so as to include θ' in $\mathsf{fut}(\theta)$. This set is inductively defined:
If $\theta = \langle i, t\rangle$ then $\Diamond\mathsf{aff}(\theta, t) = \{\langle j, t''\rangle | t'' \geq t + d_{ij}\}$, and for

[3] We assume that messages between agents are never instantaneous. In practice, the systems we set up in [2, 3] are such that minimal transmission time *per channel* is 1, so that $\Box\mathsf{unaff}(\theta)$ gets a richer structure, as described below for the context γ^{b}.

[4] In analogy to the definition of the D_{ij} values, d_{ij} is defined as the shortest distance between i and j in the min-weighted network graph.

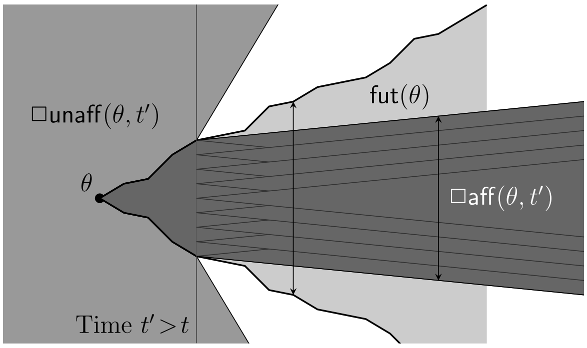

Figure 3: The necessarily affected and unaffected regions by θ in γ^{b}, w.r.t. time $t' > t$

every $t' > t$

$$\Diamond\mathsf{aff}(\theta, t') = \bigcup_{(k,\bar{t}) \in \Diamond\mathsf{aff}(\theta, t'-1)} \{\langle j, t''\rangle | t'' \geq \bar{t} + d_{kj}\}.$$

The set of necessarily unaffected nodes $\Box\mathsf{unaff}(\theta, t')$ is the complement of $\Diamond\mathsf{aff}(\theta, t')$.

Observe that with time, as longer prefixes of the run are realized, the set $\Diamond\mathsf{aff}(\theta, t') \cap \Diamond\mathsf{unaff}(\theta, t')$ of nodes that are neither necessarily affected by θ nor necessarily unaffected by it shrinks monotonically. This can be visualized by comparing the state of the cones in Figure 2 with that in Figure 3, which displays the same run at a later point in time. By time t', every node in the time interval $[t, t']$ is either in $\Box\mathsf{aff}(\theta, t')$ or in $\Box\mathsf{unaff}(\theta, t')$. Moreover, the $\Box\mathsf{aff}(\theta, t')$ cone and $\Box\mathsf{unaff}(\theta, t')$ regions grow monotonically with t'.

In summary, while light cones define fixed regions of influence and concurrency, communication dynamically determines the cones of influence and their complements. Our work on the role of transmission time upper bounds has led to a crisp characterization of when nested knowledge and common knowledge can be attained [2, 3]. The fact that the set $\Box\mathsf{aff}(\theta)$ of definitely influenced nodes is fixed *a priori*, played a crucial role in capturing how the indirect knowledge about knowledge is attained.[5] In the next section we will make use of the $\Diamond\mathsf{aff}(\theta, t')$ cone guaranteed by lower bounds, to study when an agent can be sure that another agent is (or will be) ignorant of nodes and events of interest.

4. TRANSMISSION GUARANTEES AND KNOWLEDGE OF IGNORANCE

Cones of influence and information flow as discussed in the previous section are clearly closely related to knowledge about knowledge and to knowledge about ignorance. In this section we build on the cones interpretation to analyze the dynamics of epistemic states. Our analysis here will be performed within the context γ^{min}, in which lower bounds on message transmission times are available. There are no upper bounds on message transmission times in γ^{min}; messages can take an arbitrarily long amount of time to be delivered.

In such a setting, we define

$$\mathsf{past}(r, \theta) = \{\theta' :\ \theta' \twoheadrightarrow \theta\}.$$

[5]This observation is made in hindsight, as $\Box\mathsf{aff}(\theta)$ is defined here for the first time.

A basic connection between causality and knowledge in this context is given by:

Lemma 1 ([2]). *Let $r, r' \in \mathcal{R}(P, \gamma^{min})$. If $\mathsf{past}(r, \theta) = \mathsf{past}(r', \theta)$ and both runs agree on the initial states and external inputs at all nodes of $\mathsf{past}(r, \theta)$, then $r \sim_\theta r'$.*

When the agents follow fip, a stronger condition holds.

Lemma 2. *Fix $r \in \mathcal{R}(\mathsf{fip}, \gamma^{min})$. Let θ_0, θ_1 be nodes of distinct agents $i_0 \neq i_1$. Then for all nodes θ_2,
$(\mathcal{R}, r) \models K_{\theta_2}(\theta_0 \mapsto \theta_1)$ iff $\theta_0 \twoheadrightarrow \theta_1 \twoheadrightarrow \theta_2$ in r.*

Proof.

$\Rightarrow$ If $\theta_0 \not\twoheadrightarrow \theta_1$ then $(\mathcal{R}, r) \not\models K_{\theta_2}(\theta_0 \mapsto \theta_1)$ by the Knowledge Axiom.

Suppose now that $\theta_0 \twoheadrightarrow \theta_1$ but that $\theta_1 \not\twoheadrightarrow \theta_2$. Since $i_0 \neq i_1$, agent i_1 must receive a message at some point $t_1' \in (t_0, t_1]$. Since in γ^{min} every message receive is a spontaneous event, using Theorem 2 again we get that $(\mathcal{R}, r) \not\models K_{\theta_2}(\theta_0 \mapsto \theta_1)$.

$\Leftarrow$ As $\theta_0 \twoheadrightarrow \theta_1$ we have that $(\mathcal{R}, r) \models K_{\theta_1}(\theta_0 \mapsto \theta_1)$. Using this, from $\theta_1 \twoheadrightarrow \theta_2$ and fip, we get not only that $(\mathcal{R}, r) \models K_{\theta_2}(\theta_1 \mapsto \theta_2)$, but also $(\mathcal{R}, r) \models K_{\theta_2}(K_{\theta_1}(\theta_0 \mapsto \theta_1))$. This can be weakened to the required $(\mathcal{R}, r) \not\models K_{\theta_2}(\theta_0 \mapsto \theta_1)$.

□

Based on these connections, the techniques of [4, 6, 2] can be used to show a theorem in the style of Chandy and Misra [4] stating that Lamport causality is a necessary requirement for knowledge gain in γ^{min}:

Theorem 2. ***[Knowledge Gain Theorem]*** *Suppose that e is a spontaneous event occurring at $\langle i_0, t\rangle$ in the run $r \in \mathcal{R} = \mathcal{R}(P, \gamma^{min})$. If $(\mathcal{R}, r) \models K_{\theta_k} K_{\theta_{k-1}} \cdots K_{\theta_1} \mathsf{occurred}_t(e)$, then there is a chain $\langle i_0, t\rangle \twoheadrightarrow \theta_1' \twoheadrightarrow \cdots \twoheadrightarrow \theta_k'$ in r, such that for each $h \leq k$, if $\theta_h = \langle i_h, t_h\rangle$ then $\theta_h' = \langle i_h, t_h'\rangle$ for some $t_h' \leq t_h$. Moreover, if $P = \mathsf{fip}$ then the converse holds as well.*

Proof. We shall use the convention that $\theta_h = \langle i_h, t_h\rangle$ and $\theta_h' = \langle i_h, t_h'\rangle$ throughout the proof.

Let $\mathcal{R} = \mathcal{R}(P, \gamma^{min})$, for an arbitrary protocol P. We use induction on k to prove the counter-positive form: if there is no chain $\langle i_0, t\rangle \twoheadrightarrow \theta_1' \twoheadrightarrow \cdots \twoheadrightarrow \theta_k'$ with the θ_i''s as stated in the claim, then

$$(\mathcal{R}, r) \not\models K_{\theta_k} K_{\theta_{k-1}} \cdots K_{\theta_1} \mathsf{occurred}_t(e).$$

$k = 1$: Suppose that, for all $t_1' \leq t_1$, we have $\langle i_0, t\rangle \not\twoheadrightarrow \theta_1'$. In particular $\langle i_0, t\rangle \not\twoheadrightarrow \theta_1$. Let $r' \in \mathcal{R}$ be a run such that

- $\mathsf{past}(r, \theta_1) = \mathsf{past}(r', \theta_1)$,
- all messages that are received in r outside of $\mathsf{past}(r, \theta_1)$ are delayed beyond time t_1, and finally
- e does not occur in r'.

That such a run exists is follows from the fact that in γ^{min} the adversary has the freedom to delay any mesasge indefinitely without affecting the local states of processes that the message receipt does not causally precede. By Lemma 1, since $\mathsf{past}(r, \theta_1) = \mathsf{past}(r', \theta_1)$, we have that $r \sim_{\theta_1} r'$. Since by definition of r' $(\mathcal{R}, r') \not\models \mathsf{occurred}_t(e)$, we get $(\mathcal{R}, r) \not\models K_{\theta_1} \mathsf{occurred}_t(e)$.

$k > 1$: Since we assume that

$$(\mathcal{R}, r) \vDash K_{\theta_k} K_{\theta_{k-1}} \cdots K_{\theta_1} \mathsf{occurred}_t(e),$$

we get that $(\mathcal{R}, r') \vDash K_{\theta_{k-1}} \cdots K_{\theta_1} \mathsf{occurred}_t(e)$ for every $r' \sim_{\theta_k} r$. For simplicity, assume further that the agent of node θ_h is distinct from that of θ_{h-1}, for all h.

Let T be the maximal values of all of the t_i's. We choose r' such that (i) $\mathsf{past}(r, \theta_k) = \mathsf{past}(r', \theta_k)$ and (ii) all messages received in r outside of $\mathsf{past}(r, \theta_k)$ are delayed in r' beyond time $T + 1$. The existence of r' is ensured for the same reasons as before. By Lemma 1 we obtain that $r \sim_{\theta_k} r'$, and hence that $(\mathcal{R}, r') \vDash K_{\theta_{k-1}} \cdots K_{\theta_1} \mathsf{occurred}_t(e)$. By the inductive assumption we obtain that $\langle i_0, t\rangle \twoheadrightarrow \theta'_1 \twoheadrightarrow \cdots \twoheadrightarrow \theta'_{k-1}$ in r', for some $\theta'_1, \theta'_2, .., \theta'_{k-1}$. Assume wlog that θ'_{k-1} denotes the earliest time t'_{k-1} for which such a chain exists.

If there exists some θ'_k such that $t'_k \leq t_k$ and $\theta'_{k-1} \twoheadrightarrow \theta'_k$ then we are done; so suppose not. As $\theta'_{k-1} \notin \mathsf{past}(r', \theta_k)$, we get that there are no message deliveries at the node θ'_{k-1}, contradicting the assumption that $\langle i_0, t\rangle \twoheadrightarrow \theta'_1 \twoheadrightarrow \cdots \twoheadrightarrow \theta'_{k-1}$ where θ'_{k-1} is as early as possible. So it must be that $\theta'_{k-1} \twoheadrightarrow \theta_k$ in r, and we are done.

For the claim regarding the other direction, we now assume that $P = \mathsf{fip}$ and that $\langle i_0, t\rangle \twoheadrightarrow \theta'_1 \twoheadrightarrow \cdots \twoheadrightarrow \theta'_k$. As $\langle i_0, t\rangle \twoheadrightarrow \theta'_1$, and as agents convey their full history in each sent message, we get that $(\mathcal{R}, r) \vDash K_{\theta'_1} \mathsf{occurred}_t(e)$. From perfect recall we obtain $(\mathcal{R}, r) \vDash K_{\theta_1} \mathsf{occurred}_t(e)$. Similar considerations give us that $(\mathcal{R}, r') \vDash K_{\theta_h} \cdots K_{\theta_1} \mathsf{occurred}_t(e)$ for all $h \leq k$, by induction on h. $\square$

In the rest of this section, we will focus on how different cones of influence combine to determine when an agent knows that another agent is ignorant about an event of interest. We will give a complete characterization of this question for the fip and draw implications from this to the general case of arbitrary protocols.

Recall the sealed-bid first-price auction described in the Introduction. Our bidder is named i_2, her competitor is i_1, and the bids need to be in by time t_1. Moreover, i_2 must decide on her bid at time t_2. Finally, the event e in which information about a newly found gold mine was disclosed occurred at $\theta_0 = \langle i_0, t_0\rangle$. Writing $\theta_1 = \langle i_1, t_1\rangle$ and $\theta_2 = \langle i_2, t_2\rangle$, the goal is to determine whether $K_{\theta_2} \neg K_{\theta_1} \mathsf{occurred}_{t_0}(e)$.

As discussed in Section 3, the meaning of the lower bounds min_{ij} in the network description Net is that message chains in which messages travel faster than the lower bounds allow are impossible. We say that a sequence $\theta_0, \theta_1, \ldots, \theta_m$ of nodes is a *legal message chain* with respect to Net if for every $h < m$ we have (i) $t_h < t_{h+1}$ and (ii) if $i_h \neq i_{h+1}$ then $(i_h, i_{h+1}) \in E$ and $(t_{h+1} - t_h) \geq \min_{i_h i_{h+1}}$. Clearly, for every legal message chain, there is a run of γ^{min} with network Net in which this message chain is realized, and $\theta_0 \twoheadrightarrow \theta_1 \twoheadrightarrow \cdots \twoheadrightarrow \theta_m$. Conversely, if $\theta \twoheadrightarrow \theta'$ in a run r, then there is a legal message chain starting at θ and ending at θ', that is a causal chain in r.

We are assuming that agents follow the full-information protocol fip. Thus, by Theorem 2 we have that $K_{\theta_2} \neg K_{\theta_1} \mathsf{occurred}_{t_0}(e)$ will hold if θ_2 knows that $\theta_0 \not\twoheadrightarrow \theta_1$ in the *current run*. We now formalize the required conditions, based on causal cones and legal message chains.

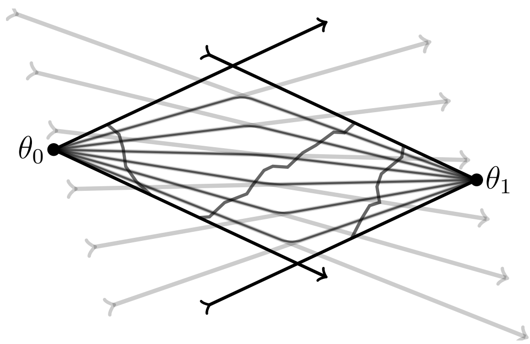

Figure 4: The set $\mathsf{Legal}_{\theta_0\theta_1}$ of possible chains from θ_0 to θ_1

Definition 2. **[Legal paths]** We denote by $\mathsf{Legal}_{\theta_0\theta_1}$ the set of legal message chains starting at θ_0 and ending at θ_1.

$\mathsf{Legal}_{\theta_0\theta_1}$ consists of all message chains that are within both $\Diamond\mathsf{aff}(\theta_0)$, which is the region of nodes that can possibly *be affected* by θ_0, and also within the analogous region of nodes that can possibly *affect* θ_1. See Figure 4.

Definition 3. **[Cut]** A $\theta_0\theta_1$-cut is a set of nodes C that appear in the paths of $\mathsf{Legal}_{\theta_0\theta_1}$ that intersects every path in $\mathsf{Legal}_{\theta_0\theta_1}$.
The cut C is called θ_0*-clean* in run r if $\theta_0 \not\twoheadrightarrow c$ in r, for every $c \in C$.

Figure 4 depicts three different $\theta_0\theta_1$ cuts (in gray).

Lemma 2 tells us that in order for an agent to be informed of any communication link between two distinct sites θ_0 and θ_1, the site on the receiving side must be in the agent's past causal cone. The above discussion suggests that the existence of a clean cut on the set of legal paths $\mathsf{Legal}_{\theta_0\theta_1}$ is of importance. Moreover, we should be looking for cuts that are somehow more "recent". The following definition formalizes this intuition.

Definition 4. **[Causal front]** Fix nodes $\theta_0, \theta_1, \theta_2$. The causal front of $\theta_0\theta_1$ with respect to θ_2 in run r, denoted by $\mathtt{Front}_{\theta_2}(r, \theta_0, \theta_1)$, is the set of nodes

$$\left\{ \phi \;\middle|\; \begin{array}{l} \phi \text{ is on some chain } \Psi \in \mathsf{Legal}_{\theta_0\theta_1} \text{ and} \\ \exists \Psi' \in \mathsf{Legal}_{\phi\theta_1} \text{ s.t. } \Psi' \cap \mathsf{past}(\theta_2) = \{\phi\} \end{array} \right\}$$

Let Ψ be a legal message chain connecting between θ_0 and θ_1 that is also, at least in part, within the scope of $\mathsf{past}(\theta_2)$. By definition of $\mathtt{Front}_{\theta_2}(r, \theta_0, \theta_1)$, it will contain a "latest contact point" ϕ, of θ_2 with the nodes of Ψ. So, as far as i_2 knows at time t_2, it is possible that $\phi \twoheadrightarrow \theta_1$. Now if it is also the case that $\theta_0 \twoheadrightarrow \phi$, then a communication path between θ_0 and θ_1 has been established. There is a certain subtlety involved in the definition. The fact that $\langle i, t\rangle$ is in $\mathtt{Front}_{\theta_2}(r, \theta_0, \theta_1)$ does not mean that $\langle i, t'\rangle$ is not in the front for $t' > t$. We can still have $\langle i, t'\rangle \in \mathtt{Front}_{\theta_2}(r, \theta_0, \theta_1)$ for some $t' > t$, if each of the nodes $\langle i, t\rangle$ and $\langle i, t'\rangle$ constitutes a latest contact point for some potential path to θ_1.

We are now ready to characterize knowledge of ignorance in γ^{min}, by showing that it reduces to existence of a "θ_0-clean" cut in the causal front:

THEOREM 3. *Let $r \in \mathcal{R}(\mathsf{fip}, \gamma^{min})$ and denote $F = \mathtt{Front}_{\theta_2}(r, \theta_0, \theta_1)$. Then $(\mathcal{R}, r) \vDash K_{\theta_2}(\theta_0 \not\twoheadrightarrow \theta_1)$ iff both (a) F is θ_0-clean, and (b) F is a $\theta_0\theta_1$-cut.*

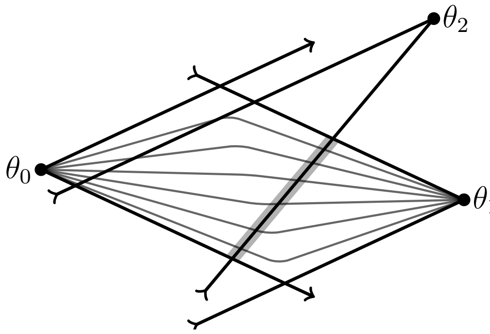

Figure 5: Thick marking gives a schematic view of $\texttt{Front}_{\theta_2}(r, \theta_0, \theta_1)$

PROOF.

$\Rightarrow$: Suppose, to the contrary, that F is either not θ_0-clean, or is not a $\theta_0\theta_1$-cut. Choose a run $r' \in \mathcal{R}$ such that

- $\mathsf{past}(r', \theta_2) = \mathsf{past}(r, \theta_2)$, and where
- all messages sent and delivered outside $\mathsf{past}(r, \theta_2)$ have minimal transmission times.

That such a run exists is given by $\mathcal{R}$ being a representing system and by the non-dependence of spontaneous events on the past of the run in which they occur. By Lemma 1 we get that $r \sim_{\theta_2} r'$, and hence that $(\mathcal{R}, r') \vDash \theta_0 \not\mapsto \theta_1$. Moreover, as $\mathsf{past}(r', \theta_2) = \mathsf{past}(r, \theta_2)$ we also have that

$$\texttt{Front}_{\theta_2}(r', \theta_0, \theta_1) = \texttt{Front}_{\theta_2}(r, \theta_0, \theta_1) = F.$$

We now have two choices:

F **is not a** $\theta_0\theta_1$**-cut:** Then there exists $\Psi = \langle \psi_0 = \theta_0, \psi_1, \cdots, \psi_n = \theta_1 \rangle \in \mathsf{Legal}_{\theta_0\theta_1}$ such that $\Psi \cap \mathsf{past}(r', \theta_2) = \emptyset$. By definition of r' we get that $\psi_0 \twoheadrightarrow \psi_1 \cdots \twoheadrightarrow \psi_n$ in r', and hence that $(\mathcal{R}, r') \vDash \theta_0 \mapsto \theta_1$, contradiction.

F **is not** θ_0**-clean:** In this case there exists $\Psi = \langle \psi_0 = \theta_0, \psi_1, \cdots, \psi_n = \theta_1 \rangle \in \mathsf{Legal}_{\theta_0\theta_1}$ and some $k < n$ such that $\psi_0 \twoheadrightarrow \psi_k$ and $\langle \psi_{k+1} \cdots \psi_n \rangle \cap \mathsf{past}(r', \theta_2) = \emptyset$. Again by definition of r' we get that $\psi_{k+1} \twoheadrightarrow \cdots \twoheadrightarrow \psi_n$ in r'. We obtain that $\psi_0 \twoheadrightarrow \psi_n$ and hence that $(\mathcal{R}, r') \vDash \theta_0 \mapsto \theta_1$, again contradicting the assumption.

$\Leftarrow$: Suppose that $(\mathcal{R}, r) \nvDash K_{\theta_2}(\theta_0 \not\mapsto \theta_1)$. Then there exists a run r' such that $r \sim_{\theta_2} r'$, where $(\mathcal{R}, r') \vDash \theta_0 \mapsto \theta_1$. It follows that there exists a message chain $\Psi \in \mathsf{Legal}_{\theta_0\theta_1}$ that is realized in r'.

Since F is a $\theta_0\theta_1$-cut, there must exist some $\phi \in F \cap \Psi$. Since $\phi \in \Psi$ we get that $\theta_0 \in \mathsf{past}(r', \phi)$. Since F is a $\theta_0\theta_1$ causal front with respect to θ_2, we have that $\phi \in \mathsf{past}(r, \theta_2)$, and as $r' \sim_{\theta_2} r$ and the agents are following fip we obtain using Lemma 2 that $\mathsf{past}(r', \phi) = \mathsf{past}(r, \phi)$. This gives us that $\theta_0 \twoheadrightarrow \phi$ in r too, contradicting the assumption that F is θ_0-clean in r.

□

Theorem 3 characterizes knowledge of non-causality under fip in a system with lower bounds on transmission times. Based on the knowledge gain theorem, we can translate this into conditions on when an agent will know that another agent is ignorant of the occurrence of an event of interest. Consider an event e_0 that can occur only at i_0 (we call it an i_0 event). We are interested in when $K_{\theta_2}\neg K_{\theta_1}\mathsf{occurred}_{t_1}(e_0)$ holds. Clearly, if i_2 knows that e_0 did not take place, then she would know that i_1 does not know that e_0 took place. Theorem 3 provides a condition enabling knowledge at θ_2 that $\theta_0 \not\mapsto \theta_1$. Suppose that $\theta_0 = \langle i_0, t_0 \rangle$. Since $\twoheadrightarrow$ is transitive, however, $\theta_0 \not\twoheadrightarrow \theta_1$ implies that $\theta' \not\twoheadrightarrow \theta_1$ for all $\theta' = \langle i_0, t' \rangle$ with $t' > t_0$. So, by the Knowledge Gain theorem, θ_1 could not have knowledge that e_0 happened at any time after t_0 too! Combining these observations, we are able to obtain a tight characterization of knowledge about ignorance regarding the occurrence of a spontaneous event:

THEOREM 4. ***[Knowledge of Ignorance Theorem]*** *Let $r \in \mathcal{R}(\mathsf{fip}, \gamma^{min})$, fix a node θ_2, and let e_0 be an i_0-event. Let t' be the latest time for which $(\mathcal{R}, r) \vDash K_{\theta_2}\neg\mathsf{occurred}_{t'}(e_0)$ holds, and denote $\theta_0 = \langle i_0, t'+1 \rangle$. Then $(\mathcal{R}, r) \vDash K_{\theta_2}\neg K_{\theta_1}\mathsf{occurred}_{t_1}(e_0)$ iff $\texttt{Front}_{\theta_2}(r, \theta_0, \theta_1)$ is a θ_0-clean, $\theta_0\theta_1$-cut.*

PROOF.

$\Rightarrow$: We will prove the counter-position. Suppose that F is not a $\theta_0\theta_1$-cut or it is not is θ_0-clean. Thus, in particular, $\mathsf{Legal}_{\theta_0\theta_1} \neq \emptyset$ and hence $t_0 \leq t_1$. As t_0 is the latest time for which $(\mathcal{R}, r) \vDash K_{\theta_2}\neg\mathsf{occurred}_{t_0-1}(e_0)$ holds, there must exist a run $r' \sim_{\theta_2} r$ where e_0 occurs at $\langle i_0, t_0 \rangle$. As $P = \mathsf{fip}$ we get, using Lemma 2, that $\mathsf{past}(r', \theta_2) = \mathsf{past}(r, \theta_2)$ and hence that

$$\texttt{Front}_{\theta_2}(r', \theta_0, \theta_1) = \texttt{Front}_{\theta_2}(r, \theta_0, \theta_1) = F.$$

Theorem 3 now shows that $(\mathcal{R}, r') \nvDash K_{\theta_2}(\theta_0 \not\mapsto \theta_1)$. So there must exist a run $r'' \in \mathcal{R}$ such that $r'' \sim_{\theta_2} r'$, where $(\mathcal{R}, r'') \vDash \theta_0 \mapsto \theta_1$. Since the agents are following fip we get, using Theorem 2, that $(\mathcal{R}, r') \vDash K_{\theta_1}\mathsf{occurred}_{t_0}(e_0)$. Since $t_0 \leq t_1$, we get $(\mathcal{R}, r') \vDash K_{\theta_1}\mathsf{occurred}_{t_1}(e_0)$. Finally, since $r'' \sim_{\theta_2} r' \sim_{\theta_2} r$, we get that $(\mathcal{R}, r) \nvDash K_{\theta_2}\neg K_{\theta_1}\mathsf{occurred}_{t_1}(e_0)$, contradicting our assumptions.

$\Leftarrow$: Choose an arbitrary $r' \in \mathcal{R}$ such that $r \sim_{\theta_2} r'$. We consider three options for the occurrence of event e_0:

- e_0 does not occur in run r': in this case we have, in particular, that $(\mathcal{R}, r') \vDash \neg K_{\theta_1}\mathsf{occurred}_{t_1}(e_0)$.
- e_0 occurs before time t_0: in this case we obtain a contradiction to the theorem's assumption that $(\mathcal{R}, r) \vDash K_{\theta_2}\neg\mathsf{occurred}_{t_0-1}(e_0)$.
- e_0 occurs at some time $t' \geq t_0$: from $r \sim_{\theta_2} r'$ and causal trace we get that
$$\texttt{Front}_{\theta_2}(r', \theta_0, \theta_1) = \texttt{Front}_{\theta_2}(r, \theta_0, \theta_1) = F.$$
Theorem 3 is now used to show that $(\mathcal{R}, r') \vDash K_{\theta_2}(\theta_0 \not\mapsto \theta_1)$, and thus that $(\mathcal{R}, r') \vDash \theta_0 \not\mapsto \theta_1$. By definition of $\twoheadrightarrow$ we also get that $\langle i_0, t' \rangle \not\twoheadrightarrow \theta_1$. Using Theorem 2 we conclude that $(\mathcal{R}, r') \vDash \neg K_{\theta_1}\mathsf{occurred}_{t_1}(e_0)$.

We showed that $(\mathcal{R}, r') \vDash \neg K_{\theta_1}\mathsf{occurred}_{t_1}(e_0)$ for all $r \sim_{\theta_2} r'$. By definition of $\vDash$ we thus obtain that $(\mathcal{R}, r) \vDash K_{\theta_2}\neg K_{\theta_1}\mathsf{occurred}_{t_1}(e_0)$, as required.

□

5. CONCLUSIONS

The interaction between knowledge and communication is natural, and while it has been widely recognized [14, 12, 16, 7, 5], it has not been explored at the level of detail that it deserves. While information flow and causality in the totally asynchronous model are well understood following Lamport's work [10] and Chandy and Misra's follow-on [4], understanding the interaction between clocks and the dynamics of epistemics is a very new topic. We have recently extended Chandy and Misra's knowledge gain results (as well as Lamport's notions of causality) in a nontrivial way to synchronous systems with upper bounds on transmission times [2, 3]. The current paper extends further and studies the impact of lower bounds.

While lower bounds on transmission times are typically of limited impact in timing-based algorithms such as clock-synchronization algorithms, our thesis in the current paper is that lower bounds play a crucial role in determining agents' ignorance and their knowledge about ignorance. This, in turn, can be of value in player's considerations in non-cooperative settings. In the current paper, we characterized when knowledge about ignorance is obtained in runs of the full-information protocol, in the presence of lower bounds. A natural question involves characterizing knowledge of ignorance for general protocols, or in strategic settings in which a player has uncertainty concerning other players' strategies [9, 17]. Our results have a natural implication applicable in such settings: If Alice knows that even under the full-information protocol Bob cannot know about a spontaneous event occurring at Charlie's, then she can conclude the same even under lesser communication.

While our analysis holds for event occurrences, this line of reasoning cannot be used to deduce Bob's ignorance about more general facts. For example, Bob may know that Charlie is ignorant of events occurring at Bob's, precisely *because* there had been no messages sent that disclose this. Thus, we may consider several possible lines for future research. First, under fip, what are the conditions that guarantee knowledge of *nested* ignorance: when does Alice know that Bob is ignorant of Charlie's ignorance of events occurring at Dafney's? Alternatively, what conditions guarantee knowledge of (non-nested) ignorance when we relax the protocol requirements? Assuming we keep to non-encrypted truthful communication, we can still relax fip in two ways: either send fewer messages, or send less information in each message.

One class of relaxed protocols, for which Theorems 3 and 4 still hold, are protocols that are *causally traced*. Intuitively, these are protocols in which whenever an agent-time pair θ_0 does send a message, it also makes sure to include the skeletal form of its past, so that the receiver will be able to tell whether $\theta_2 \twoheadrightarrow \theta_1 \twoheadrightarrow \theta_0$ for any arbitrary θ_1, θ_2. Under such a protocol however, it is not necessary to also send the contents of all such previous communications, as is done in fip.

Several other simplistic relaxations come to mind, such as a protocol where agents send messages every n turns. These can be easily accommodated within the current framework, however in a rather *ad hoc* way. A deeper and more general inquiry into relations between knowledge and communication under specific (arbitrary) protocols is yet to be undertaken.

The system γ^{min} considered throughout most of this paper assumes that communication bounds and the current time are common knowledge. In distributed computing it is natural to require a relaxation of this assumption, in the form of a system where each agent keeps its own subjective clock. In non-cooperative settings, it may be more natural to keep the global clock but to assume partial knowledge of the communication bounds, or even of the protocol being used.

In this paper, we draw an analogy between the causal cones that are formed by information in synchronous systems with bounds, and the notion of causal light-cones in physics. The invariance of the speed of light causes the causal cone of a given point in 4-dimensional Einstein-Minkowski spacetime to be fixed *a priori* and unchanging as time proceeds. In contrast, in the digital space of communication networks, upper bounds induce $\Box\mathsf{aff}(\theta)$, a region of points that are definitely affected by a spontaneous event occurring at θ, while lower bounds define $\Box\mathsf{unaff}(\theta)$, a region of points guaranteed to be causally *unaffected* by θ. These regions grow with time, converging at the end of time to partition the space of all points into those actually affected by θ and those unaffected by it. We used this view to motivate our analysis of knowledge of ignorance. We believe that further study of the causal cones and their evolution over time will provide insights into the fundamental properties of synchronous environments.

6. REFERENCES

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

[2] Ido Ben-Zvi and Yoram Moses. Beyond Lamport's happened-before: On the role of time bounds in synchronous systems. In *Proc. 24th International Symposium on Distributed Computing (DISC)*, pages 421–436, 2010.

[3] Ido Ben-Zvi and Yoram Moses. On interactive knowledge with bounded communication. In *Conference on Logic and the Foundations of Game and Decision Theory (LOFT)*, 2010.

[4] K. M. Chandy and J. Misra. How processes learn. *Distributed Computing*, 1(1):40–52, 1986.

[5] Fred I. Dretske. *Knowledge and the flow of information.* Blackwell, Oxford, 1981.

[6] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge.* MIT Press, Cambridge, Mass., 2003.

[7] Joseph Farrell and Matthew Rabin. Cheap talk. *Journal of Economic Perspectives*, 10(3):514–531, 1996.

[8] R. P. Feynman, R. B. Leighton, and M. Sands. *The Feynman Lectures on Physics including Feynman's Tips on Physics: The Definitive and Extended Edition.* Addison Wesley, 2nd edition, 2005.

[9] Aviad Heifetz Martin and Martin Meier. Interactive unawareness. *Journal of Economic Theory*, 130:78–94, 2006.

[10] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.

[11] L. Lamport. On interprocess communication, Part I: basic formalism. *Distributed Computing*, 1(2):77–85, 1986.

[12] D. Lewis. *Convention, A Philosophical Study*. Harvard University Press, Cambridge, Mass., 1969.

[13] F. Morton. *The Rothschilds: Portrait of a Dynasty*. Kodansha International, 1998.

[14] Y. Moses. Knowledge and communication (a tutorial). In *Theoretical Aspects of Reasoning about Knowledge: Proc. Fourth Conference*, pages 1–14. Morgan Kaufmann, San Francisco, Calif., 1992.

[15] G. L. Naber. *The Geometry of Minkowski Spacetime: An Introduction to the Mathematics of the Special Theory of Relativity*. Dover Publications, 2003.

[16] A Michael Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 1973.

[17] Y. Wang. *Epistemic Modelling And Protocol Dynamics*. PhD thesis, ILLC, 2010.