On Interactive Knowledge with Bounded Communication

Ido Ben-Zvi[1] and Yoram Moses[2]

[1] The Technion, Israel idobz@cs.technion.ac.il
[2] The Technion, Israel moses@ee.technion.ac.il

Abstract. The effect of synchronous communication upon the dynamics of knowledge gain is investigated. In a previous paper [6] we have defined necessary conditions for knowledge gain in such systems. We now show that our definitions are tight. Utilizing these conditions, we show that when agents follow the full information protocol, a bound exists on the level of interactive knowledge that can be attained without common knowledge. This result is hedged by showing that in general no bound on the number of levels of E_G that will imply C_G exists.

1 Introduction

The growing body of research dealing with the dynamics of interactive epistemics has brought to light some of the intricate relations between time, knowledge and action [10, 19]. Yet the potential impact of another factor, namely communication, is often abstracted away.

In distributed computing however, the characteristics of the communication layer are vigorously studied. In particular, the distinction between synchronous and asynchronous communication is cardinal here. Knowledge oriented results pertaining to this distinction have mostly been concerned with asynchronous communication. Thus [12] have shown that common knowledge cannot be attained in the face of unreliable communication. Complementing this result, Chandy and Misra [8] have identified the necessary message exchange that must take place in order for agents in an asynchronous system to gain interactive knowledge to any finite degree. Yet in many if not most multi agent systems, some degree of synchrony can be assumed. This is the case in game theoretic analysis, and as has been shown by van Benthem et al. [18], this assumption is also implicit in Dynamic Epistemic Logic.

The main thrust of the current paper is to show that synchronous communication entails a rich underlying causal structure that imposes certain epistemic constraints, and that can be used to purposefully achieve others. Our analysis follows the approach of Chandy and Misra in characterizing interactive knowledge gain in terms of the underlying causal structure that relates agents in particular knowledge states. A different approach has recently been proposed By Apt et. al in [2].

Suppose that as a result of a nondeterministic occurance at time t, Alice comes to possess a juicy bit of information, namely that φ. Suppose that due to its nondeterministic nature, no one else could possibly come to know that φ indepedently. In an important paper [8], Chandy and Misra showed that in asynchronous systems, interactive knowledge is essentially determined by message sequences. Thus, if at time t' Bob knows that Alice knows φ, which we denote by $K_B K_A \varphi$, then there is a message sequence starting with Alice after time t, and reaching Bob by t'. As shown in [8], this property generalizes nicely: if we add Charlie into the story, then for $K_C K_B K_A \varphi$ to hold, there must be a message chain starting at Alice, passing through Bob, and reaching Charlie. We use the term *knowledge gain* to describe the process of attaining a high level of interactive knowledge. The crux of the argument here is that for Bob to gain knowledge about Alice, he must be somehow causally affected by her, and yet in asynchronous systems there is no way for Alice to causally affect Bob, epistemically or otherwise, except by means of a message chain.

Switching to a synchronous setting adds considerable complexity to any potential analysis of knowledge gain.[3] For example, if messages are known to arrive at most b time units after being sent, Alice could potentially let Bob know that $K_A \varphi$ by time $t + b$, by *not* sending him any message after time t (say if Bob and Alice have a previous agreement that Alice will send Bob messages at every time point until, if ever, she finds out that φ). Moreover, we could have $K_C K_B K_A \varphi$ at $t + b$, despite there being no message exchange between Bob and Charlie, or between Alice and Bob. In fact, this level of knowledge could be reached without a single message being sent (and received) in the duration $(t, t + b]$.

As these examples illustrate, the flow of information in synchronous systems is nontrivial. In a previous paper [6], we suggested a causal notion that we call *syncausality*, and showed that it is necessary for knowledge gain in synchronous settings. Interactive knowledge gain and common knowledge gain were shown to require richer structures that simple message chains, that we call *centipedes* and *centibrooms*. We briefly discuss these definitions, as well as our formal framework, in Section 2. Novel results start in Section 3, where we complete the characterization of knowledge gain, interactive knowledge gain and common knowledge gain by showing that the suggested causal structures cannot be strengthened. This is shown by defining the *full information protocol* (fip) for synchronous systems, where agents are fully cooperative in dispersing new information, and showing that under this protocol the causal structures are also sufficient for knowledge gain. We continue looking into fip in Section 4, where we show that there exists a bound on the depth of finite interactive knowledge that can be reached in a group G of agents, before common knowledge is attained. We provide a tight bound on this depth, as a function of $|G|$ and the elapsed time. Roughly speaking, it takes time to achieve deep knowledge *without* having common knowledge. We also show that, somewhat surprisingly, this bound depends not only on synchronous communication but also upon certain qualities of the protocol. These results interact with Parikh's [15] convincing motivation for the relevance of finite depth interactive knowledge without com-

[3] See Parikh and Ramanujam's [16] for some related analysis.

mon knowledge, both in real life scenarios as well as in game theoretic settings. Section 5 briefly discusses possible extensions of this paper.

This paper's main contributions are:

– The scope of causal analysis of the epistemic dynamics is extended. It is shown that the causal notions of syncausality, centipedes and centibrooms fully characterize knowledge gain, interactive knowledge gain and common knowledge gain in synchronous systems respectively.
– The causal definitions make explicit the role of time in determining attainable states of interactive knowledge.
– The usage of protocols, the commonly known aspects of the agents' behavior, in epistemic analysis is exemplified. Such common knowledge has a drastic effect on the evolution of knowledge in the system.
– The levels of interactive knowledge that can be attained without reaching common knowledge are studied.
  • When agents follow fip, there is a bound $M = (|G| - 1)(d - 1) + 2$, where d denotes the elapsed time since an event that has taken place, such that E_G^M implies common knowledge in G that the event has occured. This bound is shown to be tight: E_G^{M-1} can hold without implying C_G. This demonstrates the role played by the time elapsed in determining when levels of E_G give rise to common knowledge. Moreover, it shows that the time interacts with the size of G in this matter, in a nontrivial way.
  • No bound on the number of levels of E_G that will imply C_G exists in general. Namely, a protocol is presented in which E_G^k can be attained with arbitrarily large k in one round of communication, without C_G holding.
– Analysis of cases where both temporal and epistemic considerations exist is done by means of a timestamping technique reminiscent of the freeze quantifier of [1].

2 Definitions and Preliminary Results

In this section we present the background for the analysis performed in the paper. We start by describing the model and semantic definitions, Then, we present the material from our previous paper [6] discussing notions of causality and how they impose necessary conditions on knowledge gain.

2.1 The Model

This paper follows the *runs and systems* approach of Fagin et al. [10] to modeling multi-agent systems. Two essential building blocks are used to define the formal model. The *context* in which the agents are operating, 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 $\mathcal{R} = \mathcal{R}(P, \gamma)$, which is the set of all possible *runs*, or histories, of P in the context γ. The truth of facts of interest can change from one run to another, and within a run, from one time point to the next. We thus focus on facts that are true at *points* (r, t), specifying the run and time in question. In every point (r, t), the system is in a particular global state. The global state is identified with a tuple of local agent states, as well as a state for the *environment*, which accounts for, e.g., messages in transit.

In this paper we focus our attention a particular *synchronous context* γ^{s}. We use $\mathcal{R}^{\mathsf{s}}$ to denote a system $\mathcal{R}(P, \gamma^{\mathsf{s}})$ consisting of the set of all runs of protocol P in the synchronous context γ^{s}. A detailed definition of the models and the context γ^{s} are presented in Appendix A. Its main properties are highlighted below.

- We assume that agents can receive external inputs from the outside world. These 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.
- The set of agents is denoted by $\mathbb{P}$. The network consists of a weighted graph over $\mathbb{P}$, in which edges stand for communication channels, and the weights are natural numbers. We denote by b_{ij} the weight for the channels from i to j, and it represents the upper bound on transmission times for messages sent along this channel. A copy of the network, as well as the current global time, are part of every agent's local state at all times, and hence are common knowledge at all times.
- The environment agent is in charge of choosing these external inputs, and of determining message transmission times. The latter are also determined in a nondeterministic fashion, subject to the constraint that delivery satisfies the transmission bounds.
- Time is identified with the natural numbers, and each agent is assumed to take a step at each time t. 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.
- Events are message sends and receives, external inputs, and internal computations performed by the agent. All events in a run are distinct, and we denote a generic event by the letter e. Since an event occurs at a unique site, we assume disjoint sets $\{\mathcal{E}_i\}_{i\in\mathbb{P}}$, of events for each process. A *nondeterministic* (or *ND*) event is the arrival of an external input, or a message delivery that occurs strictly before the transmission bound b_{ij} for its channel is met.

2.2 Syntax and Semantics

We follow the framework of [10] very closely. We focus on a simple logical language in which the set Φ of primitive propositions consists of propositions of the form ϕ^e for all events e, as well as ones of the form $\mathtt{time} = t$ for times t. To obtain the logical language $\mathcal{L}$, we close Φ under propositional connectives and knowledge formulas. Thus, $\Phi \subset \mathcal{L}$, and if $\varphi \in \mathcal{L}$, $i \in \mathbb{P}$, and $G \subseteq \mathbb{P}$, then $\{K_i\varphi, E_G\varphi, C_G\varphi\} \subset \mathcal{L}$. The formula $K_i\varphi$ is read *agent i knows φ*, $E_G\varphi$ is read *everyone in G knows φ*, and $C_G\varphi$ is read *φ is common knowledge to G*. In addition, given the role that time plays in our analysis, we add a timestamping operator as well. Thus, if $\varphi \in \mathcal{L}$ and $t \in \mathbb{N}$, then $\varphi_{@t}$ is a formula.

The truth of a formula is defined with respect to a triple (R, r, t). We write $(R, r, t) \vDash \varphi$ to state that φ holds at time t in run r, with respect to system R. Unless stated otherwise, it is always assumed that $r \in R$ in a triple (R, r, t). Denoting by $r_i(t)$ agent i's local state at time t in r, we inductively define

$(R, r, t) \vDash \phi^e$ if event e occurs at time t in r;
$(R, r, t) \vDash \mathtt{time} = t'$ if $t' = t$;
$(R, r, t) \vDash \varphi_{@\hat{t}}$ if $(R, r, \hat{t}) \vDash \varphi$;
$(R, r, t) \vDash K_i\varphi$ if $(R, r', t') \vDash \varphi$ for every run r' satisfying $r_i(t) = r'_i(t')$;
$(R, r, t) \vDash E_G\varphi$ if $(R, r, t) \vDash K_i\varphi$ for every $i \in G$; and
$(R, r, t) \vDash C_G\varphi$ if $(R, r, t) \vDash (E_G)^k\varphi$ for every $k \geq 1$.

Propositional connectives are handled in the standard way, and their clauses are omitted above. By definition, $K_i\varphi$ is satisfied if φ holds at all points at which i has the same local state. Thus, given R, the local state determines what facts are true. Intuitively, a fact φ is common knowledge to G if everyone in G knows φ, everyone knows that everyone knows φ, and so on *ad infinitum*. In particular, if $(R, r, t) \vDash C_G\varphi$ then $(R, r, t) \vDash K_{i_h}K_{i_{h-1}} \cdots K_{i_1}\varphi$, for every string $K_{i_h}K_{i_{h-1}} \cdots K_{i_1}$ and $h > 0$.

We write $R \vDash \varphi$ and say that "φ is *valid* in R" if $(R, r, t) \vDash \varphi$ holds for all $r \in R$ and $t \geq 0$. A formula is valid, written $\vDash \varphi$, if it is valid in all systems R. Two simple but very useful properties of the timestamping operator @ are captured by the following:

TS1 $\quad \vDash \mathtt{time} = t \rightarrow (\varphi \leftrightarrow \varphi_{@t})$ $\quad$ and
TS2 $\quad \vDash (\varphi_{@t})_{@t'} \leftrightarrow \varphi_{@t}$

We now turn to summarize the preliminary results introduced in [6]. We start with a discussion of relevant notions of causality in synchronous contexts.

2.3 Causality

Notions of potential causality play a crucial role in the analysis of knowledge gain. Typically, causality is considered among events. For ease of exposition, we consider causal relations among agent-time pairs (i, t), which we call *nodes*. There is no loss of generality in doing so, because in our particular framework it is possible to identify an event with the node at which it occurs. Consider the following forms of precedence among nodes:

Local precedence: $(i, t) \overset{\ell oc}{\longmapsto} (i, t')$ if $t \leq t'$.

Message precedence: $(i, t) \overset{msg}{\longmapsto} (j, t')$ if i sends a message at (i, t) that is received by j at (j, t'). This notion is defined with respect to a particular run r.

Timeout precedence: $(i,t) \overset{t/o}{\longmapsto} (j, t+b_{ij})$ if i and j are connected by a channel with transmission bound b_{ij}. If agent j does not receive a message from i by time $t + b_{ij}$, it can "*timeout*" and determine that i did not send j a message at time t.

Different combinations of these relations can be used to describe different meaningful causal relations. Thus, for example, the transitive closure $\xrightarrow{L} = \{\overset{\ell oc}{\longmapsto}, \overset{msg}{\longmapsto}\}^*$ of local precedence and message precedence gives rise to *Lamport causality* [13], which plays a central role in the analysis of asynchronous systems. Under this notion, nodes at distinct sites can only be causally related by means of a message chain leading from one agent's node to the other's. Chandy and Misra used Lamport causality to characterize knowledge gain in asynchronous systems. Another causal relation that we make use of is $\dashrightarrow = \{\overset{\ell oc}{\longmapsto}, \overset{t/o}{\longmapsto}\}^*$, the transitive closure of local precedence and timeout precedence, which we call *silent causality*. Silent causality reflects the pace at which information can spread in the network without messages actually being sent. This pace is determined by the upper bounds on transmission time, so it is also the speed at which a message is delivered if agents relay it as fast as possible, but en route messages are maximally delayed. We think of this pace as the *speed of silence*. In the synchronous context γ^{s}, all three forms of precedence can give rise to a causal connection between nodes. As a result, we define *syncausality* among nodes in a run of γ^{s} to be $\rightsquigarrow = \{\overset{\ell oc}{\longmapsto}, \overset{msg}{\longmapsto}, \overset{t/o}{\longmapsto}\}^*$, the transitive closure of all three. Syncausality is thus a fusion of Lamport's causality and silent causality. As Theorem 2.1 below shows, syncausality is closely related to knowledge gain regarding events at a remote site.

Intuitively, we'd like to say that, at a given node for Alice, she can know about an event e that occurred at Bob's site only if the event syncausally precedes Alice's current node. This is too strong, however. For example, Alice may know at 10pm that Bob is in bed, because he always goes to sleep by 9:30. In this case, Alice's knowledge did not require any transfer of information; it was guaranteed based on Bob's protocol. In order to avoid such issues, we consider knowledge gain regarding *nondeterministic* events. Recall that these are either external inputs to the agents, or the arrival of messages that are received strictly before they must according on the bound b_{ij} for the channel they are sent on. Before such an event occurs, there is no guarantee that it will occur.

Theorem 2.1 (Basic Knowledge Gain [6]). *Let e be an ND event at (i_0, t) in run $r \in \mathcal{R}^{\mathsf{s}}$. If $(\mathcal{R}^{\mathsf{s}}, r, t') \models K_{i_1}\, \phi^{e}_{@t}$ then $(i_0, t) \rightsquigarrow (i_1, t')$.*

2.4 Centipedes and Centibrooms

Syncausality is necessary for the basic form of knowledge gain in the synchronous setting, just as Lamport causality (message chains) is necessary in the asynchronous case. However, while the same message chains also describe higher-level knowledge gain in the asynchronous setting, in synchronous systems a more complex casual structure is required, as stated by Theorem 2.3 below. This structure, that we call *centipede*, combines syncausality and silent causality restrictions. Centipedes are shown in Figure 1 and defined below. A guiding intuition behind the theorem is that the synchronous environment allows for a relaxation of the strict message sequence that characterizes asynchronous knowledge gain, by allowing nodes along the path from (i_0, t) to (i_k, t') to divert information to the intermediate agents $i_1, .., i_{k-1}$, if it can be guaranteed that the information will reach these agents by t'.

Definition 2.2 (Centipede)**.** Let $r \in \mathcal{R}^{\mathsf{s}}$, let $i_h \in \mathbb{P}$ for $0 \le h \le k$ and let $t \le t'$. A *centipede* for $\langle i_0, \ldots, i_k\rangle$ in the interval $(r, t..t')$ is a sequence $\theta_0 \rightsquigarrow \theta_1 \rightsquigarrow \cdots \rightsquigarrow \theta_k$ of nodes such that $\theta_0 = (i_0, t)$, $\theta_k = (i_k, t')$, and $\theta_h \dashrightarrow (i_h, t')$ holds for $h = 1, \ldots, k-1$.

$$\begin{tikzcd}
(i_0, t) \arrow[r, rightsquigarrow] & \theta_1 \arrow[r, rightsquigarrow] \arrow[d, dashed] & \theta_2 \arrow[r, rightsquigarrow] \arrow[d, dashed] & \cdots \arrow[r, rightsquigarrow] & \theta_{k-1} \arrow[r, rightsquigarrow] \arrow[d, dashed] & (i_k, t') \\
& (i_1, t') & (i_2, t') & & (i_{k-1}, t') &
\end{tikzcd}$$

Fig. 1. A Centipede for $\langle i_0, \ldots, i_k\rangle$ in $(r, t..t')$

Notice that each of the intermediate nodes $\theta_h = (j_h, t_h)$ with $h < k$ takes place at a time $t_h \le t'$. Moreover the nodes (i_1, t'),. . . ,(i_k, t') at the legs of the centipede all occur at time t'. Note that the legs of a centipede are not necessarily pairwise distinct, so that there could be $1 \le h < g \le k$ such that $(i_h, t') = (i_g, t')$.

Theorem 2.3 (Centipede Theorem [6])**.** *Let* $r \in \mathcal{R}^{\mathsf{s}} = \mathcal{R}(P, \gamma^{\mathsf{s}})$*, and let e be an ND event at* (i_0, t) *in r. If* $(\mathcal{R}^{\mathsf{s}}, r, t') \vDash K_{i_k} K_{i_{k-1}} \cdots K_{i_1} \phi^e_{@t}$*, then there is a centipede for* $\langle i_0, \ldots, i_k\rangle$ *in* $(r, t..t')$*.*

Turn now to common knowledge, which has been shown to be of particular import and use in many epistemic scenarios [3, 10]. The *centibroom* causal structure is depicted in Figure 2 and defined below. It roughly corresponds to a centipede in which all intermediate nodes coincide.

Definition 2.4 (Centibroom)**.** Let $t \le t'$ and $G \subseteq \mathbb{P}$. There is a *centibroom* ฿$\langle i_0, G\rangle$ in $(r, t..t')$ if there exists a node θ such that $(i_0, t) \rightsquigarrow \theta \dashrightarrow (i_h, t')$ holds for all $i_h \in G$.

Theorem 2.5 shows that the existence of a centibroom it is a necessary condition for common knowledge.

Theorem 2.5 (Common Knowledge Gain [6])**.** *Let* $r \in \mathcal{R}^{\mathsf{s}} = \mathcal{R}(P, \gamma^{\mathsf{s}})$*, and let e be an ND event at* (i_0, t) *in r. If* $(\mathcal{R}^{\mathsf{s}}, r, t') \vDash C_G\, \phi^e_{@t}$*, then there is a centibroom* ฿$\langle i_0, G\rangle$ *in* $(r, t..t')$*.*

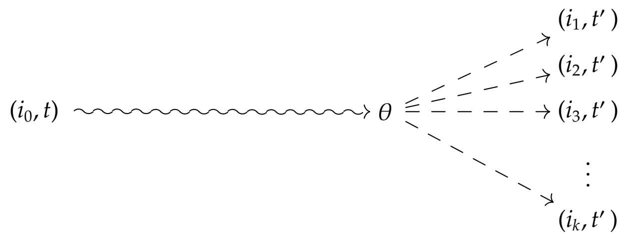

Fig. 2. A Centibroom, $\mathbb{B}\langle i_0, G\rangle$ in $(r, t..t')$

It is interesting to relate our results concerning common knowledge to those given by Apt et al. in [2]. In the formal system defined in [2], common knowledge in group G cannot be attained unless a message is broadcast to all agents in G simultaneously. In contrast, the setting discussed here allows for a wider range of ways to attain common knowledge. A series of point-to-point messages, or even a combination of messages and silence, can also serve. The difference in the results can be traced to the fact that in the framework we use, the protocol that the agents follow is an explicit parameter determining the system (i.e., the set of runs) in which knowledge is evaluated. The protocol, in turn, plays an important role in determining the information flow within an execution. In the frameowrk of [2], the class of protocols under consideration is somewhat restricted. Within this restricted class, they show that a series of point-to-point messages cannot give rise to common knowledge. Our results show that there are many reasonable protocols for which point-to-point messages *can* give rise to common knowledge.

3 Sufficiency of Causal Structures

Theorems 2.3 and 2.5 state that the centipede and centibroom structures are *necessary* for knowledge and, respectively, common knowledge gain regarding nondeterministic events, regardless of the protocol used by the agents. We now show that these results are tight. Namely, there is a protocol for which these structures are *sufficient*, as well as necessary for the corresponding form of epistemic gain.

Definition 3.1 (Full-information Protocol)**.** The *full information protocol* for synchronous systems, fip, is one in which every agent $i \in \mathbb{P}$ sends its local state on each of its outgoing channels at every time step.

In fip the agents convey all of their knowledge as fast as they can. Roughly speaking, knowledge is spread in the system as fast as the communication channels will allow. In fact, syncausality captures the spread of knowledge under the fip. Denoting $\mathcal{R}^{\mathsf{fip}} = \mathcal{R}(\mathsf{fip}, \gamma^{\mathsf{s}})$, we can show:

Lemma 3.2. *If* $(j, t) \overset{r}{\rightsquigarrow} (i, t')$ *and* $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_j\varphi$ *then* $(\mathcal{R}^{\mathsf{fip}}, r, t') \models K_i(\varphi_{@t})$.

Proofs not shown in this section can be found in Appendix B. Intuitively, if an event e takes place at (i,t), then all neighbors j of i will receive a message from i no later than time $t+b_{ij}$, after which $K_j\phi^e_{@t}$ will hold. Observe that there are no "silent messages" in fip, in the sense that agents constantly sendin explicit messages on all outgoing channels. Nevertheless, the $\dashrightarrow$ relation still relates nodes, and it does so even more strongly than $\rightsquigarrow$ does. Note first that $(j,t) \dashrightarrow (i,t')$ implies that $(j,t) \rightsquigarrow (i,t')$. Moreover, $\rightsquigarrow$ is dependent on the current run's nondeterministic behavior, since message precedence is dependent upon realized transmission times. The $\dashrightarrow$ relation however, is determined by the context γ^{s} alone. So that if $(i,t) \dashrightarrow (j,t')$ holds in a run r of the system, it will do so in *all* runs of the system. Thus, as we show in Lemma 3.3, agent j knows already at time t that its current knowledge will be available to i at $t+b(j,i)$. This serves to explain the crucial function played by silent causality in fip information flow, despite the fact that the agents are never really silent. Formally then, we have:

Lemma 3.3. *If* $(j,t) \dashrightarrow (i,t')$ *and* $(\mathcal{R}^{\mathsf{fip}},r,t) \models K_j\varphi$ *then* $(\mathcal{R}^{\mathsf{fip}},r,t) \models K_j(K_i(\varphi_{@t})_{@t'})$.

Lemmas 3.2 and 3.3 capture the essential epistemic aspects of the fip in the synchronous context γ^{s}. They facilitate proving that centipedes and centibrooms are sufficient for ensuring the appropriate levels of knowledge in this setting.

Theorem 3.4. *If* $(\mathcal{R}^{\mathsf{fip}},r,t) \models K_{i_0}\varphi$ *and there is a centipede for* $\langle i_0,\dots,i_k\rangle$ *in* $(r,t..t')$*, then* $(\mathcal{R}^{\mathsf{fip}},r,t') \models K_{i_k}K_{i_{k-1}}\cdots K_{i_1}K_{i_0}(\varphi_{@t})$.

Showing that the existence of a centibroom is sufficient for common knowledge gain in every $\mathcal{R}^{\mathsf{fip}}$ system makes use of the *Induction Rule for Common Knowledge*, which states that from $\mathcal{R}^{\mathsf{fip}} \models \alpha \to E_G(\alpha\wedge\beta)$ we can infer $\mathcal{R}^{\mathsf{fip}} \models \alpha \to C_G\beta$.

Theorem 3.5. *If* $(\mathcal{R}^{\mathsf{fip}},r,t) \models K_{i_0}\varphi$ *and there is a centibroom* $\text{\textbaht}\langle i_0,G\rangle$ *in* $(r,t..t')$*, then* $(\mathcal{R}^{\mathsf{fip}},r,t') \models C_G(\varphi_{@t})$.

Proof Assume that the conditions of the theorem hold, and let the single intermediate node in the centibroom $\text{\textbaht}\langle i_0,G\rangle$ be (j,t_j). In particular, $(j,t_j) \dashrightarrow (i,t')$ for every $i\in G$. From $(\mathcal{R}^{\mathsf{fip}},r,t) \models K_{i_0}\varphi$ and $(i_0,t) \overset{r}{\rightsquigarrow} (j,t_j)$ we have by Lemma 3.2 that $(\mathcal{R}^{\mathsf{fip}},r,t_j) \models K_j(\varphi_{@t})$. We now use the induction rule with α set to $(\mathtt{time}=t')\wedge(K_j(\varphi_{@t}))_{@t_j}$, and β being $\varphi_{@t}$. Since $\mathcal{R}^{\mathsf{fip}} \models \alpha\to\beta$ in this case, it suffices to show that $\mathcal{R}^{\mathsf{fip}} \models \alpha \to E_G\alpha$. Thus, let $r'\in\mathcal{R}^{\mathsf{fip}}$ and suppose that $(\mathcal{R}^{\mathsf{fip}},r',\hat{t}) \models (\mathtt{time}=t')\wedge(K_j(\varphi_{@t}))_{@t_j}$. In particular we have by TS1 and TS2 that $(\mathcal{R}^{\mathsf{fip}},r',t_j) \models K_j((K_j(\varphi_{@t}))_{@t_j})$. Fix $i\in G$. Since $(j,t_j) \dashrightarrow (i,t')$, we have by Lemma 3.3 that $(\mathcal{R}^{\mathsf{fip}},r',t') \models K_i((K_j(\varphi_{@t}))_{@t_j})$. Moreover, the fact that the time is part of the local state in γ^{s} implies that $(\mathcal{R}^{\mathsf{fip}},r',t') \models K_i(\mathtt{time}=t')$. It follows that $(\mathcal{R}^{\mathsf{fip}},r',t') \models K_i\alpha$, and since i was an arbitrarily chosen member of G then $(\mathcal{R}^{\mathsf{fip}},r',t') \models E_G\alpha$. It follows that $\mathcal{R}^{\mathsf{fip}} \models \alpha\to E_G\alpha$. Since $\beta=\varphi_{@t}$ we obtain by the Induction Rule that $\mathcal{R}^{\mathsf{fip}} \models \alpha \to C_G(\varphi_{@t})$. Finally, since $(\mathcal{R}^{\mathsf{fip}},r,t') \models \alpha$ we obtain that $(\mathcal{R}^{\mathsf{fip}},r,t') \models C_G(\varphi_{@t})$, as desired. $\Box_{\mathit{Theorem}\ 3.5}$

Note that both Theorems 3.4 and 3.5 provide us with a stronger result than is needed for showing sufficiency. We only claim the appropriateness of centipedes and centibrooms for gained knowledge about nondeterministic events, whereas the fip protocol allows us to show sufficiency for any fact φ such that $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_{i_0}\varphi$.

4 Common Knowledge as a Finite Conjunction

Common knowledge is typically perceived in terms of an infinite conjunction of E^k, for $k > 0$. There are definitions of common knowledge in terms of a fixed point (see, e.g., [14, 10, 5]). The centibroom structure and the necessity of centibrooms for common knowledge supports the fixed-point view.

Even though the fixed point definition implies the infinite conjunction, Fischer and Immerman [11] had shown that in *finite-state* systems, where the set of all global states in a system R is finite, there is a power k such that C_G is equivalent to E^k_G. The fip protocol, coupled with the perfect recall inherent in γ^{s}, produces a state space whose size is unbounded. Yet nevertheless, given the role of the centipede and centibroom structures in γ^{s}, we now show that there are cases in which common knowledge is a finite conjunction when running fip in γ^{s} as well.

Roughly speaking, it takes time to obtain deep knowledge *without* having common knowledge. Indeed, we obtain a sharp bound on the depth of E^k_G that can be obtained d time units after the occurrence of a nondeterministic event. Given a group of size $|G| = g$ and natural number $d > 0$, we denote by $M_{dg} = (d-1)(g-1) + 2$. We prove

Theorem 4.1. *Let $r \in \mathcal{R}^{\mathsf{fip}}$, $d > 0$, $|G| = g$, and assume that e is an ND event at (i_0, t) in r. If $\mathcal{R}^{\mathsf{fip}}, r, t+d \models E_G^{M_{dg}}\, \phi^e_{@t}$ then $\mathcal{R}^{\mathsf{fip}}, r, t+d \models C_G\, \phi^e_{@t}$.*

Theorem 4.1 follows directly by Theorem 3.5 from the following lemma:

Lemma 4.2. *Let $\mathcal{R}^{\mathsf{s}} = \mathcal{R}(P, \gamma^{\mathsf{s}})$, let $d > 0$, $|G| = g$, and assume that e is an ND event at (i_0, t) in r. If $(\mathcal{R}^{\mathsf{s}}, r, t+d) \models E_G^{M_{dg}} \phi^e_{@t}$ then there exists a centibroom $\rlap{/}\beta\langle i_0, G\rangle$ in $(r, t..t+d)$.*

Proof Assume that $(\mathcal{R}^{\mathsf{s}}, r, t+d) \models E_G^{M_{dg}} \phi^e_{@t}$. If (i_0, t) is a centibroom $\rlap{/}\beta\langle i_0, G\rangle$ in $(r, t..t+d)$ then we are done. Otherwise $|G| > 1$, and moreover there is some $j \in G$ such that $(i_0, t) \not\rightsquigarrow (j, t')$. For notational convenience, let us denote the agents of G by $\{j_0, \ldots, j_m\}$, where $(i_0, t) \not\rightsquigarrow (j_0, t')$. Denote $M = M_{dg} - 1$ and let $f(h) = j_{(h \bmod m+1)}$ for all $h \leq M$. Thus, f maps natural numbers into members of G, every sequence of $m+1$ adjacent numbers are mapped to the full set $\{j_0, \ldots, j_m\} = G$, and $f(0) = j_0$. We focus on a knowledge formula of the form

$$\Psi(e) = K_{f(M)}\, K_{f(M-1)} \cdots K_{f(1)} K_{f(0)}\ \phi^e_{@t}\ .$$

Observe that there are $M + 1 = (|G| - 1) \cdot (d - 1) + 2$ knowledge operators in $\Psi(e)$, all of which belong to agents in G. By assumption, $(\mathcal{R}^{\mathsf{s}}, r, t+d) \models$

$E_G^{M+1}\,\phi^e_{@t}$, and hence in particular $(\mathcal{R}^{\mathsf{s}}, r, t+d) \models \Psi(e)$. The Centipede Theorem implies that there exists a centipede for $\langle i_0, f(0), f(1), .., f(M)\rangle$ in $(r, t..t+d)$. Let $\langle (i_0, t), \Omega^0, \Omega^1, \ldots, \Omega^{M-1}, (f(M), t+d)\rangle$ be such a centipede. By definition of a centipede we have that $(i_0, t) \rightsquigarrow \Omega^0$ and $\Omega^0 \dashrightarrow (j_0, t+d)$. Since '$\dashrightarrow$' is transitive, the fact that $(i_0, t) \not\dashrightarrow (j_0, t+d)$ implies that $(i_0, t) \not\dashrightarrow \Omega^0$. Since '$\dashrightarrow$' is reflexive we have that $(i_0, t) \neq \Omega^0$. Recall by definition of f and the fact that $|G| > 1$ that $f(M) \neq f(M-1)$. Hence, by Lemma B.1 we have that $(f(M), t+d) \not\dashrightarrow (f(M-1), t+d)$. It follows that $\Omega^{M-1} \neq (f(M), t+d)$.

By Lemma B.1, if $\Omega^h = \Omega^{h'}$ then $\Omega^h = \Omega^{\hat{h}} = \Omega^{h'}$ for every $\hat{h}$ in the range $h \leq \hat{h} \leq h'$. Let $\Phi_1, \ldots, \Phi_D$ denote the maximal sub-sequence of distinct nodes in the sequence $\Omega^0, \ldots, \Omega^{M-1}$. Lemma B.1 implies that the times at which the nodes $(i_0, t), \Phi_1, \ldots, \Phi_D, (f(M), t+d)$ occur form a strictly increasing sequence, and so $D \leq d-1$. For all b in the range $1 \leq b \leq D$ define $s(b) = \{k : \Omega^k = \Phi_b\}$. Since $M = (|G|-1)\cdot(d-1)+1$ and $D \leq d-1$, we have by the pigeonhole principle that $|s(\hat{b})| \geq |G|$ for at least one index $\hat{b}$. Since the set $s(\hat{b})$ consists of at least $|G| = m+1$ consecutive natural numbers, we have that $\{f(k) : k \in s(\hat{b})\} = \{j_0, \ldots, j_m\} = G$. By definition of the centipede it follows that $\Omega_{\hat{b}} \dashrightarrow (j, t+d)$ for all $j \in s(\hat{b}) = G$, and so $\Omega_{\hat{b}}$ gives rise to a centibroom $\mathbb{B}\langle i_0, G\rangle$ in $(r, t..t+d)$, as required. ∎

Lemma 4.3. *For every $t \geq 0$, $d > 0$ and $g > 1$ there exists a run $r \in \mathcal{R}^{\mathsf{fip}}$, an ND event e at (i_0, t) in r and a set of agents $G \subseteq \mathbb{P}$ of size $|G| = g$, such that*

$$(\mathcal{R}^{\mathsf{fip}}, r, t+d) \models \quad E_G^{M_{dg}-1}\,\phi^e_{@t} \quad\wedge\quad \neg C_G\,\phi^e_{@t}.$$

Theorems 4.1 and 4.3 tightly bound the levels of E_G^k that can hold common knowledge necessarily arising. They draw an essential connection between this bound, the size of the set of agents in question, and the time that elapses since the ND event of interest occurs. It is natural to ask whether this property is restricted to fip, or perhaps may be true in general. We now show that it is not true for all protocols. In fact, there is a protocol that can attain arbitrary levels of interactive knowledge quickly, without giving rise to common knowledge.

Example 4.4. Let $\hat{\gamma}^{\mathsf{s}}$ be a context with $\mathbb{P} = \{i_0, 1, 2\}$, where the network is a star with center at i_0, and the communication bounds are $b_{i_0,1} = b_{i_0,2} = 1$. We assume that i_0 can receive nondeterministic external input x at time t, and that this input is of interest to agents 1 and 2. Further, i_0 receives external input on every time point, that allows it to choose a natural number $k \geq 0$ and an agent $h \in \{1, 2\}$. We consider a protocol $\hat{P}$ in which only i_0 can move, it moves only at time t, and acts as follows: it sends the message $\langle k+1, \phi^{e_x}_{@t}\rangle$ to the agent $\bar{h} \neq h$, and if $k > 0$ it also sends $\langle k, \phi^{e_x}_{@t}\rangle$ to the agent h. If $k = 0$ then no message is sent to agent h.

An inductive argument now shows

Lemma 4.5. *Let $r \in R = \mathcal{R}(\hat{P}, \hat{\gamma}^{\mathsf{s}})$, let $G = \{1, 2\}$ and assume that e_x takes place at (r, t). Then $(R, r, t+1) \models E_G^k\,\phi^{e_x}_{@t}$, yet $\quad R \models \neg C_G\,\phi^{e_x}_{@t}$.*

The example can be extended to sets $|G|$ of any size, by replacing treating all of the agents in $G \setminus \{h\}$ as $\bar{h}$ was in $\hat{P}$. We note that the epistemic structure obtained here is similar to that of the electronic mail game of Rubinstein [17]. One distinguishing feature is that in $\hat{P}$ the high degree of interactive knowledge is obtained in one step, with two messages, whereas k messages were required in the game. The same epistemic structure also arises in the analysis of the initial states of the *muddy children* puzzle [10], or of the Conway paradox [9].

5 Conclusions

The causal relations underlying synchronous systems can be further explored in a number of ways. For example, we would also like to explore the dynamics of ignorance in such a system. Moreover, synchrony is a working assumption in many multi agent conceptualizations, and hence we believe that the current analysis may be fruitfully applied. Noting that our defined causal relations are a prerequisite for any epistemic change, one such application may be to explore its ramifications in the context of belief revision [4]. Alternatively, utilizations for mechanism design [7] may also be of interest.

References

[1] Rajeev Alur and Thomas A. Henzinger. A really temporal logic. *J. ACM*, 41(1):181–204, 1994.

[2] Krzysztof R. Apt, Andreas Witzel, and Jonathan A. Zvesper. Common knowledge in interaction structures. In *TARK '09: Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 4–13, New York, NY, USA, 2009. ACM.

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

[4] Alexandru Baltag, L. S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. Technical report, Amsterdam, The Netherlands, The Netherlands, 1999.

[5] J. Barwise. Three views of common knowledge. In M. Y. Vardi, editor, *Proc. Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 365–379. Morgan Kaufmann, San Francisco, Calif., 1988.

[6] I. Ben-Zvi and Y. Moses. Causality, knowledge and common knowledge in synchronous systems. 2010.

[7] Dirk Bergemann and Stephen Morris. Robust mechanism design. *Econometrica*, 73:1771–1813, 2001.

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

[9] J. H. Conway, M. S. Paterson, and U. S. S. R. Moscow. A headache-causing problem. In J. K. Lenstra et al., editors, *Een pak met een korte broek: Papers presented to H. W. Lenstra on the occasion of the publication of his "Euclidische Getallenlichamen"*. Private publication, 1977.

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

[11] M. J. Fischer and N. Immerman. Foundations of knowledge for distributed systems. In J. Y. Halpern, editor, *Theoretical Aspects of Reasoning about Knowledge: Proc. 1986 Conference*, pages 171–186. Morgan Kaufmann, San Francisco, Calif., 1986.

[12] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990. A preliminary version appeared in *Proc. 3rd ACM Symposium on Principles of Distributed Computing*, 1984.

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

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

[15] R. Parikh. Levels of knowledge, games, and group action. *Research in Economics*, 57(3):267–281, September 2003.

[16] Rohit Parikh and Ramaswamy Ramanujam. A knowledge based semantics of messages. *Journal of Logic, Language and Information*, 12(4):453–467, 2003.

[17] Ariel Rubinstein. The electronic mail game: Strategic behavior under "almost common knowledge.". *American Economic Review*, 79(3):385–91, 1989.

[18] Johan van Benthem, Jelle Gerbrandy, and Eric Pacuit. Merging frameworks for interaction: Del and etl. In *TARK '07: Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge*, pages 72–81, New York, NY, USA, 2007. ACM.

[19] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer Publishing Company, Incorporated, 2007.

A Detailed Model

A.1 Contexts, Protocols and Systems

Formal semantics for knowledge in systems is given by Fagin, Halpern, Moses, and Vardi [10]. We shall simplify the exposition somewhat here, and review what we hope will be just enough of the details to support the analysis in this paper. Essentially all of the definitions in this section are taken from [10].

Informally, we view a multi agent system as consisting of a set $\mathbb{P} = \{1, \ldots, n\}$ of agents connected by a communication network. We assume that, at any given point in time, each agent in the system is in some *local state*. A *global state* is just a tuple $g = \langle \ell_e, \ell_1, \ldots, \ell_n \rangle$ consisting of local states of the agents, together with the state ℓ_e of the *environment*. The environment's state accounts for everything that is relevant to the system that is not contained in the state of the agents.

A *run* is a function from time to global states. Intuitively, a run is a complete description of what happens over time in one possible execution of the system. A *point* is a pair (r, t) consisting of a run r and a time t. If $r(m) = \langle \ell_e, \ell_1, \ldots, \ell_n \rangle$, then we use $r_i(m)$ to denote agent i's local state ℓ_i at the point (r, t), for $i = 1, \ldots, n$, and $r_e(m)$ to denote ℓ_e. For simplicity, time here is taken to range over the natural numbers rather than the reals (so that time is viewed as discrete, rather than dense or continuous). *Round* t in run r occurs between time $t - 1$ and t.

We identify a *protocol* for a agent i with a function from local states of i to nonempty sets of actions. (We often consider *deterministic* protocols, in which a local state is mapped to a singleton set of actions. Such a protocol essentially maps local states to actions.) A joint protocol is just a sequence of protocols $P = (P_1, \ldots, P_n)$, one for each agent.

We generally study knowledge in runs of a given protocol P in a particular setting of interest (in this paper the setting is synchronous, with global clocks and bounds on message transmission times). To do this, we separately describe the setting, or *context*, in which P is being executed. Formally, a context γ is a tuple $(\mathcal{G}_0, P_e, \tau)$, where $\mathcal{G}_0$ is a set of initial global states, P_e is a protocol for the environment, and τ is a *transition function*.[4] The environment is viewed as running a protocol (denoted by P_e) just like the agents; its protocol is used to capture nondeterministic aspects of the execution, such as the actual transmission times, external inputs into the system, etc. The transition function τ describes how the actions performed by the agents and the environment change the global state. Thus, if g is a global state and $\mathsf{a} = \langle \mathsf{a}_e, \mathsf{a}_1, \ldots, \mathsf{a}_n \rangle$ is a *joint action* (consisting of an action for the environment and one for each of the agents), then $\tau(\mathsf{a}, g) = g'$ specifies that g' is the state that results when a is performed in state g.

A run r is consistent with a protocol P if it could have been generated when running protocol P. Formally, run r is *consistent with joint protocol P in context γ* if

1. $r(0) \in \mathcal{G}_0$, so that it starts from a γ-legal initial global state, and
2. for all $t \geq 0$, the transition from global state $r(t)$ to $r(t+1)$ is the result of performing one of the joint actions specified by P and the environment protocol P_e (the latter is specified in γ) in the global state $r(m)$. That is, if $P = (P_1, \ldots, P_n)$, P_e is the environment's protocol in context γ, and $r(m) = \langle \ell_e, \ell_1, \ldots, \ell_n \rangle$, then there must be a joint action $\mathsf{a} = \langle \mathsf{a}_e, \mathsf{a}_1, \ldots, \mathsf{a}_n \rangle$ such that $\mathsf{a}_e \in P_e(\ell_e)$, $\mathsf{a}_i \in P_i(\ell_i)$ for $i = 1, \ldots, n$, and $r(m+1) = \tau(\mathsf{a}, r(m))$ (so that $r(m+1)$ is the result of applying the joint action a to $r(m)$).

We use $\mathcal{R}(P, \gamma)$ to denote the set of all runs of P in γ, and call it the *system representing P in context γ*.

A.2 The Synchronous Context $\gamma^{\mathsf{s}} = (\mathcal{G}_0^{\mathsf{s}}, P_e^{\mathsf{s}}, \tau^{\mathsf{s}})$

This paper considers knowledge in systems arising in a particular synchronous context. We now present some of the particular assumptions made, by describing the components of the context. In the following sections we will use $\mathcal{R}^{\mathsf{s}}$ to denote a system system representing P in context γ^{s}, where P is an arbitrary protocol over which we make no assumptions, and γ^{s} is the synchronous context described below.

The environment's state Recall that the environment's state keeps track of relevant aspects of the global state that are not represented in the local states of the agents. We assume that the environment's state has three components $\ell_e = (\mathsf{Net}, \overset{t/o}{\longmapsto}, Hist_e)$, where

[4] Depending on the application, a context can include additional components, to account for fairness assumptions, probabilistic assumptions, etc. Moreover, additional aspects of a context that are usually suppressed from the notation are nonempty sets Int and Ext of internal actions for the agents and external inputs, respectively.

1. Net is a labelled graph $(\mathbb{P}, E, w)$ describing the network topology and bounds on transmission times. Its nodes are agents, and a directed edge (i, j) captures the fact that there is a channel from i to j in the system. Moreover, the label $1 \le w(i, j) \in \mathbb{N}$ is an upper bound on the time that a message sent on (i, j) can be in transit. The contents of Net are not affected by τ^{s}, and so Net remains constant throughout the run.
2. The variable $\xmapsto{t/o}$ keeps track of global time. As we shall see its value starts at $\xmapsto{t/o}= 0$, and advances by 1 following each round. It is a notion of global time because we assume that agents have access to its value as described below. Finally,
3. $Hist_e$ records the sequence of joint actions performed so far. Since, as we will discuss below, there is no message loss in γ^{s}, the $Hist_e$ component uniquely determines the contents of all channels. Indeed, a message μ is *in transit* at a given global state g if the $Hist_e$ component in g records that μ has been sent, and does not record its delivery.

Agent local states We assume that local states have three components $\ell_i = (\mathsf{Net}_i, \xmapsto{t/o}_i, hist_i)$, where Net_i and $\xmapsto{t/o}_i$ are copies of the Net and $\xmapsto{t/o}$ values from the environment's state. The $hist_i$ component records the local events that take place at i. It consists of an initial local state for i, followed by a sequence recording all local events that have taken place so far. If we denote by $inp_i(k)$ the set of messages and external inputs delivered to i in round k, then $hist_i$ at time $t \ge 0$ consists of the sequence $\langle \mathsf{init}_i, \mathsf{a}_i(1), inp_i(1), \ldots, inp_i(m-1), \mathsf{a}_i(m-1), inp_i(m)\rangle$. In particular, in the global states of $\mathcal{G}_0^{\mathsf{s}}$, the $hist_i$ component contains only the initial state init_i.

The set $\mathcal{G}_0^{\mathsf{s}}$ of initial global states We assume that associated with γ^{s} there is a set Init_i of possible initial states for each agent $i \in \mathbb{P}$. We define $\mathcal{G}_0^{\mathsf{s}}$ to be the set of global states $g = (\ell_e, \ell_1, \ldots, \ell_n)$ satisfying: (1) $\xmapsto{t/o}_i= 0$ for all $i \in \mathbb{P}$ and $\xmapsto{t/o}= 0$; (2) the network components Net and Net_i are all identical; (3) for every $i \in \mathbb{P}$, $hist_i = \langle \mathsf{init}_i\rangle$,with $\mathsf{init}_i \in \mathrm{Init}_i$; and (4) $Hist_e$ is the empty sequence.

Actions and external inputs Associated with the context γ^{s} are sets Int of internal actions for the agents and sets Ext of external inputs, respectively. For ease of exposition we assume that $\perp \in \mathsf{Ext}$, where $\perp$ stands for the empty external input. Moreover, we generally assume that $\mathsf{Ext} \neq \{\perp\}$, so that there is at least one nontrivial possible external input. We assume that agents can perform send actions and internal actions. The local action $\mathsf{a}_i(k)$ that i contributes to the joint action in round $k+1$ consists of a finite sequence of distinct send and internal actions. (Recall that the local action is determined by the protocol, based on the local state.) We use external inputs to model spontaneous events. They are generated by the environment. In addition to external inputs, the environment is in charge of message delivery. Thus, the environment's action $\mathsf{a}_e(k)$ consists of a finite sequence of external inputs to be delivered to various individual agents, and a (possibly empty) set of messages that are to be delivered in the current round.

The environment's protocol P_e^{s} The environment in γ^{s} is in charge of delivering external inputs to agents, and determining message deliveries. We define $P_e^{\mathsf{s}}(g)$ to be the set of actions $\mathsf{a}_e = (\sigma_x, \sigma_d)$ such that

1. $\sigma_x \in (\mathbb{P} \times \mathsf{Ext})$ is a sequence assigning to each agent $i \in \mathbb{P}$ an external input (possibly the empty input $\perp$) it receives in the current round, and
2. σ_d is a sequence $\langle(1, M_1), \ldots, (n, M_n)\rangle$ where (i) for every $i \in Proc$ the set M_i consists of messages that are in transit in g, and (ii) M_i contains all messages in transit to i whose transmission time bounds, as specified in Net, will be violated (expire) if the message is not delivered in the current round.

Notice that P_e^{s} is genuinely nondeterministic. Exactly one of the actions in $P_e^{\mathsf{s}}(g)$ will be performed in global state g in any given instance. By definition of $\mathcal{R}(P, \gamma^{\mathsf{s}})$, however, if $r(k) = g$ then the system contains a run extending the prefix $r(0), \ldots, r(k)$ for every possible environment action in $P_e^{\mathsf{s}}(g)$. Another point to note is that our definition does not enforce (and hence does not assume) FIFO transmission; had we done so, channels would be considered to be queues, and the nondeterministic choices of messages to deliver would have to obey FIFO order. Finally, the fact that external inputs are delivered in a nondeterministic fashion implies they are not correlated in any way, and they do not depend on anything that happens before they are delivered. This is the sense in which external inputs can be viewed as independent, "spontaneous" events.

The transition function τ^{s} The transition function τ^{s} implements the joint actions in a rather straightforward manner. In every round: (i) the global clock variables $\xmapsto{t/o}$ and $\xmapsto{t/o}_i$ are advanced by one; (ii) a copy of the joint action is added to the environment's history log $Hist_e$; and (iv) a record for the current round is added to i's local history log $hist_i$, containing all message deliveries to agent i, all external inputs to i, and the sequence of snd actions and internal actions performed by i in the current round.

B Proofs

Lemma B.1 (Identity of centipede nodes [6])**.** *If* $(i, t) \rightsquigarrow (j, t')$ *then* $t \leq t'$*, with* $t = t'$ *holding only if* $i = j$*.*

Lemma 3.2. *If* $(j, t) \overset{r}{\rightsquigarrow} (i, t')$ *and* $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_j\varphi$ *then* $(\mathcal{R}^{\mathsf{fip}}, r, t') \models K_i(\varphi_{@t})$ *.*

Proof Let $\rightsquigarrow^{nt}$ be the relation defined by the non-transitive clauses in the definition of $\rightsquigarrow$. By definition of $\rightsquigarrow$, there exists a sequence of nodes $\langle(j_0, s_0), (j_1, s_1), .., (j_n, s_n)\rangle$ in r such that $(j_0, s_0) = (j, t)$, $(j_n, s_n) = (i, t')$, and $(j_l, s_l) \rightsquigarrow^{nt} (j_{l+1}, s_{l+1})$ for every $l < n$.

We prove the claim by induction on the length of the sequence.

n = 0 : Then $j = j_0 = j_n = i$ and $t = t'$. By assumption $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_j\varphi$. So by definition of K operator and since $j = i$ we obtain that for every r' such that $(r', t) \sim_i (r, t)$ $(\mathcal{R}^{\mathsf{fip}}, r', t) \models \varphi$. Using the definition of $_@$ we get that $(\mathcal{R}^{\mathsf{fip}}, r', t) \models \varphi_{@t}$, but since $t' = t$ this gives us $(\mathcal{R}^{\mathsf{fip}}, r', t') \models \varphi_{@t'}$. Now from our choice of r's we get $(\mathcal{R}^{\mathsf{fip}}, r, t') \models K_i(\varphi_{@t})$.

n > 0 : Assume that the inductive claim holds for $n-1$. Thus we have that $\mathcal{R}^{\mathsf{fip}}, r, s_{n-1} \models K_{j_{n-1}}(\varphi_{@t})$. Moreover, by definition of the sequence we have that $(j_{n-1}, s_{n-1}) \rightsquigarrow^{nt} (j_n, s_n)$. By definition of $\rightsquigarrow^{nt}$ there are three options to consider:

1. $j_{n-1} = j_n$ and $s_{n-1} \leq s_n$ - in this case by perfect recall we have that $\mathcal{R}^{\mathsf{fip}}, r, s_n \models K_{j_n}(\varphi_{@t})$.
2. Agent j_{n-1} sends a message in round s_{n-1}, which is received by j_n in round s_n - since all agents are running the fip protocol, message contents are the local state of sender. As we inductively assume $\mathcal{R}^{\mathsf{fip}}, r, s_{n-1} \models K_{j_{n-1}}(\varphi_{@t})$, we get that $\mathcal{R}^{\mathsf{fip}}, r, s_n \models K_{j_n}(K_{j_{n-1}}(\varphi_{@t})_{@s_{n-1}})$. Application of the Knowledge Axiom on every run r' such that $(r', s_n) \sim_{j_n} (r, s_n)$ gives us that $\mathcal{R}^{\mathsf{fip}}, r, s_n \models K_{j_n}(\varphi_{@t})$.
3. $(j_{n-1}, j_n) \in E$ and $s_n = s_{n-1} + b(j_{n-1}, j_n)$ - Since in fip every agent sends its local state to all neighbors on every round, and since we assume that $\mathcal{R}^{\mathsf{fip}}, r, s_{n-1} \models K_{j_{n-1}}(\varphi_{@t})$, then by s_n a message sent by j_{n-1} in round s_{n-1} is guaranteed to have arrived at j_n, and hence, as shown in case (2), we get $\mathcal{R}^{\mathsf{fip}}, r, s_n \models K_{j_n}(\varphi_{@t})$.

■

Lemma 3.3. *If* $(j, t) \dashrightarrow (i, t')$ *and* $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_j\varphi$ *then* $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_j(K_i(\varphi_{@t})_{@t'})$.

Proof Since $(j, t) \dashrightarrow (i, t')$, and since this property is determined by the context γ^{s}, we have that $(j, t) \dashrightarrow (i, t')$ in every run $r' \in \mathcal{R}^{\mathsf{fip}}$. Moreover, since $\dashrightarrow \subseteq \rightsquigarrow$, we have that $(j, t) \rightsquigarrow (i, t')$ in every such run. Applying Lemma 3.2 to every run r' such that $(r', t) \sim_j (r, t)$ we get that $(\mathcal{R}^{\mathsf{fip}}, r', t') \models K_i(\varphi_{@t})$. By definition of $_@$ this gives us $(\mathcal{R}^{\mathsf{fip}}, r', t') \models K_i(\varphi_{@t})_{@t'}$. By choice of runs r' we get that $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_j(K_i(\varphi_{@t})_{@t'})$ ■

Theorem 3.4. *If* $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_{i_0}\varphi$ *and there is a centipede for* $\langle i_0, \ldots, i_k\rangle$ *in* $(r, t..t')$, *then* $(\mathcal{R}^{\mathsf{fip}}, r, t') \models K_{i_k}K_{i_{k-1}} \cdots K_{i_1}K_{i_0}(\varphi_{@t})$.

Proof Let $\langle (i_0, t), (j_1, t_1), .., (j_{k-1}, t_{k-1}), (i_k, t')\rangle$ be a centipede for $\langle i_0, \ldots, i_k\rangle$ in $(r, t..t')$. We show by induction on h that $(\mathcal{R}^{\mathsf{fip}}, r, t_h) \models K_{j_h}((K_{i_h}K_{i_{h-1}} \ldots K_{i_0}(\varphi_{@t}))_{@t'})$. Recall that the global time t appears as a component of all local states. Thus, $(r, t) \sim_i (r', t')$ is possible only if $t' = t$.

$h = 0$: By assumption $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_{i_0}\varphi$. By the definition of $\models$ for K operator we obtain that $(\mathcal{R}^{\mathsf{fip}}, r', t') \models \varphi$ for every (r', t) such that $(r', t) \sim_{i_0} (r, t)$. The definition of $_@$ implies that $(\mathcal{R}^{\mathsf{fip}}, r', t) \models \varphi_{@t}$ for every such run r'. It follows that $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_{i_0}(\varphi_{@t})$. By positive introspection for K_{i_0} and the definition of $\models$ we have that $(\mathcal{R}^{\mathsf{fip}}, r', t) \models K_{i_0}(\varphi_{@t})$ for every (r', t) such that $(r', t) \sim_{i_0} (r, t)$.

Perfect recall and the definition of $\models$ for $_@$ imply that $(\mathcal{R}^{\mathsf{fip}}, r', t) \models (K_{i_0}(\varphi_{@t}))_{@t'}$ holds for all such r'. Finally, by the choice of r' we obtain that $(\mathcal{R}^{\mathsf{fip}}, r, t) \models K_{i_0}((K_{i_0}(\varphi_{@t}))_{@t'})$, as desired.

$h > 0$: Assume for $h-1$ and show for h. The inductive assumption gives us that $(\mathcal{R}^{\mathsf{fip}}, r, t_{h-1}) \models K_{j_{h-1}}((K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t}))_{@t'})$. By definition of centipede we have that $(j_{h-1}, t_{h-1}) \rightsquigarrow (j_h, t_h)$. Thus using Lemma 3.2 we obtain that $(\mathcal{R}^{\mathsf{fip}}, r, t_h) \models K_{j_h}((K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t}))_{@t'})_{@t_{h-1}}$. Using TS2 we reduce this back to $(\mathcal{R}^{\mathsf{fip}}, r, t_h) \models K_{j_h}((K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t}))_{@t'})$. Again, based on the centipede definition we have that $(j_h, t_h) \dashrightarrow (i_h, t')$. Using Lemma 3.3 it follows that

$$(\mathcal{R}^{\mathsf{fip}}, r, t_h) \models K_{j_h}((K_{i_h}((K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t}))_{@t'}))_{@t'}).$$

By definition of the K operator we get that $(\mathcal{R}^{\mathsf{fip}}, r', t_h) \models (K_{i_h}(K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t}))_{@t'})_{@t'}$ for every r' such that $(r', t_h) \sim_{j_h} (r, t_h)$. By definition of $_@$, we obtain that $(\mathcal{R}^{\mathsf{fip}}, r', t') \models K_{i_h}(K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t}))_{@t'})$. Applying the TS1 rule we obtain that $(\mathcal{R}^{\mathsf{fip}}, r', t') \models K_{i_h}K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t})$. We can thus re-introduce $_{@t'}$ and obtain that $(\mathcal{R}^{\mathsf{fip}}, r', t_h) \models K_{i_h}(K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t}))_{@t'}$. Finally, since this is true for every r' satisfying $(r', t_h) \sim_{i_h} (r, t_h)$ we obtain that $(\mathcal{R}^{\mathsf{fip}}, r, t_h) \models K_{j_h}(K_{i_h}K_{i_{h-1}} \dots K_{i_0}(\varphi_{@t})_{@t'})$, completing the inductive step.

In particular, we obtain $(\mathcal{R}^{\mathsf{fip}}, r, t') \models K_{i_k}(K_{i_k}K_{i_{k-1}} \dots K_{i_0}(\varphi_{@t})_{@t'})$ since $j_k = i_k$ and $t_k = t'$. We conclude the proof by using the Knowledge Axiom to derive $(\mathcal{R}^{\mathsf{fip}}, r, t') \models K_{i_k}K_{i_{k-1}} \dots K_{i_0}(\varphi_{@t})_{@t'}$ and from this, using TS1, $(\mathcal{R}^{\mathsf{fip}}, r, t') \models K_{i_k}K_{i_{k-1}} \dots K_{i_0}(\varphi_{@t})$ and we are done. $\square_{Theorem\ 3.4}$

Lemma 4.3. *For every $t \geq 0$, $d > 0$ and $g > 1$ there exists a run $r \in \mathcal{R}^{\mathsf{fip}}$, an ND event e at (i_0, t) in r and a set of agents $G \subseteq \mathbb{P}$ of size $|G| = g$, such that*

$$(\mathcal{R}^{\mathsf{fip}}, r, t+d) \models \quad E_G^{M_{dg}-1}\, \phi^e_{@t} \quad \wedge \quad \neg C_G\, \phi^e_{@t}.$$

Proof Fix d, g. Define $\gamma^{\mathsf{s}}_{d,g}$ to be a synchronic context with the following particular properties:

- Let $\mathbb{P} = \{j_0, j_1, .., j_{g-1}\} \bigcup \{i_0\} \bigcup \{h_{k,l}\}_{1 \leq k < d, 0 \leq l < m}$
- Let $G = \{j_0, .., j_{g-1}\}$. For every $l < g$, denote by G_{-l} the set $G \setminus \{j_l\}$.
- The network graph is complete, and the bounds set on the edges are as follows
  1. for every $k < d$ and $l \leq g$, $b(h_{k,l}, j) = 1$ for all $j \in G_{-l}$
  2. for every other $i, j \in \mathbb{P}$, $b(i, j) = d + 1$

For every $1 \leq k < d$, use H_k to denote the set $\{h_{k,l}\}_{0 \leq l < g}$. Note that for every $r \in \mathcal{R}^{\mathsf{s}}_{d,g}$, as the agents are running the fip, every agent sends every other agent a message on every time unit. Note also that there can be no $\nVdash\langle i_0, G\rangle$ node in $(r, 0..d)$, because for every agent i there exists at least one $j \in G$ such that $b(i, j) > d$.

Choose $r \in \mathcal{R}^{\mathsf{s}}_{d,g}$ such that $\mathcal{R}^{\mathsf{s}}_{d,g}, r, d \models \phi^e_{@0}$ where $e \in \mathcal{E}_i$ and such that all sent messages arrive at the speed of silence, except for the following ones:

1. For every $h \in H_1$, the message sent from i_0 to h at time 0 arrives at time 1.
2. For every $1 \leq k < d-1$, for every pair of agents $h_1 \in H_k$ and $h_2 \in H_{k+1}$, the message sent by h_1 to h_2 at time k arrives at $k+1$.
3. For every $h \in H_{d-1}$ and $j \in G$, the message sent from h to j at time $d-1$ arrives at time d

The existence of r is warranted by the exhaustiveness of the representing system $\mathcal{R}^{\mathrm{s}}_{d,g}$.

Use $f(k)$ to denote the value $(g-1)\cdot k$ for every $k > 0$. Fix a sequence $S = \langle i_0, i_1, .., i_{f(d-1)+1}\rangle$ such that $\{i_1, .., i_{f(d-1)+1}\} \subseteq G$. Observe that for every $1 \leq k < d$, the subsequence $s_k\langle i_{f(k-1)+1}, .., i_{f(k)}\rangle$ contains exactly $g-1$ elements, and so there must exist some $j(k) \in G$ such that $j(k) \neq i$ for every $i \in s_k$.

We now define a node sequence $\langle (i_0, t), \theta_1, .., \theta_{f(d-1)}, (i_{f(d-1)+1}, d)\rangle$ and show that it is a centipede for S in $(r, 0..d)$. For every $l = 1..f(d-1)$, let $k = \lceil \frac{l}{g-1} \rceil$, and define $\theta_l = (h_{k,j(k)}, k)$. Observe that $f(k-1) < l \leq f(k)$, and hence by choice of $j(k)$ that $b(h_{k,j(k)}, i_l) = 1$. Since $k \leq d-1$ we obtain that $\theta_l \dashrightarrow (i_l, d)$. Moreover, if $l < f(d-1)$ then $\theta_l \rightsquigarrow \theta_{l+1}$. For if $k > \frac{l}{g-1}$ then $\theta_l = \theta_{l+1}$ and the result stems from the reflexivity of $\rightsquigarrow$, while if $k = \frac{l}{g-1}$ then, noting that $\theta_l \in H_k$ and $\theta_{l+1} \in H_{k+1}$, we get the result from clause (2) above. Finally, we note that $(i_0, t) \rightsquigarrow \theta_1 = (h_{1,j(1)}, 1)$ since $h_{1,j(1)} \in H_1$ and using clause (1), and similarly that $\theta_{f(d-1)} = (h_{d-1,j(d-1)}, d-1) \rightsquigarrow (i_{f(d-1)+1}, d)$ since $h_{d-1,j(d-1)} \in H_{d-1}$ and from clause (3) above.

We have shown that there exists a centipede in $(r, 0..d)$ for every sequence $\langle i_0, i_1, .., i_{f(d-1)+1}\rangle$ such that $\{i_1, .., i_{f(d-1)+1}\} \subseteq G$. By Theorem 3.4 we get that $\mathcal{R}^{\mathrm{s}}_{d,g}, r, t' \vDash K_{i_{f(d-1)+1}} K_{i_{f(d-1)}} \cdots K_{i_1} \phi^{e}_{@0}$ for every such sequence. We thus obtain, considering that $f(d-1)+1 = (g-1)(d-1)+1 = M_{dg}-1$, that $\mathcal{R}^{\mathrm{s}}_{d,g}, r, t' \vDash E_G^{M_{dg}-1} \phi^{e}_{@0}$, by definition of E operator. ■

Lemma 4.5. *Let $r \in R = \mathcal{R}(\hat{P}, \hat{\gamma}^{\mathrm{s}})$, let $G = \{1, 2\}$ and assume that e_x takes place at (r, t) and that i_0 also chooses k nondeterministically. Then $(R, r, t+1) \vDash E^k_G \phi^{e_x}_{@t}$, yet $R \vDash \neg C_G \phi^{e_x}_{@t}$.*

Proof For every $h' \in \{1, 2\}$ and $k' \geq 0$, let $K_{h',k'}$ be the formula $K_{h'} K_{\bar{h}'} K_{h'} K_{\bar{h}'} \cdots \phi^{e_x}_{@t}$, where there are exactly k' K operators in the formula.

Assume that i_0 has chosen k, h. We now show by induction on $0 \leq d \leq k$ that if h receives the message $\langle d', \phi^{e_x}_{@t}\rangle$ for some $d' \geq d$, then $R, r, t+1 \vDash E^d_G \phi^{e_x}_{@t}$.

$d = 0$: By assumption we have that $R, r, t \vDash \phi^{e_x}_{@t}$. Since $\phi^{e_x}_{@t}$ is stable we also get $R, r, t+1 \vDash \phi^{e_x}_{@t}$. Since by definition, $E^0_G \phi^{e_x}_{@t}$ is an abbreviation for $\phi^{e_x}_{@t}$, we also have $R, r, t+1 \vDash E^d_G \phi^{e_x}_{@t}$.

$d > 0$: Assume for $d-1$ and show for d. Suppose that agent h receives the message $\langle d', \phi^{e_x}_{@t}\rangle$ for some $d' \geq d$. Let $r' \in R$ such that $(r', t+1) \sim_h (r, t+1)$. Then, as h receives the same message also in r', by the inductive assumption, and since $d' \geq d > d-1$, we have that $R, r', t+1 \vDash E^{d-1}_G \phi^{e_x}_{@t}$. Hence, by the semantics of the K operator we get that $R, r, t+1 \vDash K_h E^{d-1}_G \phi^{e_x}_{@t}$.

By protocol $\hat{P}$, $\bar{h}$ recives the message $\langle d'+1,\ \phi^{e_x}_{@t}\rangle$ in r. Let $r' \in R$ such that $(r', t+1) \sim_{\bar{h}} (r, t+1)$. The same protocol dictates that in every such run r', h receives the message $\langle d'',\ \phi^{e_x}_{@t}\rangle$, with $d'' = d'$ or $d'' = d'+2$. In either case, based on the inductive assumption and on $d'' \geq d' > d-1$, we get that $R, r', t+1 \vDash E^{d-1}_G\, \phi^{e_x}_{@t}$. Hence, again by the semantics of the K operator we get that $R, r, t+1 \vDash K_{\bar{h}} E^{d-1}_G\, \phi^{e_x}_{@t}$.

It follows, based on the definition of E operator, that $R, r, t+1 \vDash E_G E^{d-1}_G\, \phi^{e_x}_{@t}$, from which we obtain $R, r, t+1 \vDash E^d_G\, \phi^{e_x}_{@t}$ as required.

Since in r agent h receives the message $\langle k,\ \phi^{e_x}_{@t}\rangle$, it follows from the inductive proof that $R, r, t+1 \vDash E^k_G\, \phi^{e_x}_{@t}$.

We now show that $R \vDash \neg C_G\, \phi^{e_x}_{@t}$. Fix $r \in R$ and time $t' \geq 0$. Let $r_0 = r$ and assume wlog that h gets the message $\langle k,\ \phi^{e_x}_{@t}\rangle$ from i_0 in r. Based on the protocol $\hat{P}$ and on R being a representing system, there exists a run $r_1 \in R$, where h gets the same message from i_0, so that $(r_0, t') \sim_h (r_1, t')$ (regardless of whether or not $t' > t$), but where $\bar{h}$ receives the message $\langle k-1,\ \phi^{e_x}_{@t}\rangle$. Using the same justifications, there also exists a run $r_2 \in R$ such that $(r_1, t') \sim_{\bar{h}} (r_2, t')$, but where h gets the message $\langle k-2,\ \phi^{e_x}_{@t}\rangle$ from i_0. We continue to apply the same reasoning, eventually building a sequence of runs $r_0, r_1, ..., r_k$ such that $(r_0, t') \sim_h (r_1, t') \sim_{\bar{h}} (r_2, t') \sim_h \cdots (r_k, t')$ and such that in run r_k some agent $h' \in \{1, 2\}$ gets no message from i_0. Let r_{k+1} be a run such that $R, r_{k+1}, t \vDash \neg\, \phi^{e_x}_{@t}$. Noting that $(r_k, t') \sim_{h'} (r_{k+1}, t')$, we get that it is not the case that $R, r, t' \vDash K_h K_{\bar{h}} K_h K_{\bar{h}} \cdots K_{h'}\, \phi^{e_x}_{@t}$. We have shown that $\neg C_G\, \phi^{e_x}_{@t}$ holds for every choice of $r \in R$ and $t' \geq 0$. It follows then that $R \vDash \neg C_G\, \phi^{e_x}_{@t}$ ∎