OBSERVING, REPORTING, AND DECIDING IN NETWORKS OF SENTENCES

H. JEROME KEISLER AND JEFFREY M. KEISLER

Abstract. In prior work [7] we considered networks of agents who have knowledge bases in first order logic, and report facts to their neighbors that are in their common languages and are provable from their knowledge bases, in order to help a decider verify a single sentence. In report complete networks, the signatures of the agents and the links between agents are rich enough to verify any decider's sentence that can be proved from the combined knowledge base. This paper introduces a more general setting where new observations may be added to knowledge bases and the decider must choose a sentence from a set of alternatives. We consider the question of when it is possible to prepare in advance a finite plan to generate reports within the network. We obtain conditions under which such a plan exists and is guaranteed to produce the right choice under any new observations.

1. Introduction

This paper builds upon the paper [7]. In that paper, a signature network is a network of agents each labeled with a signature. Each agent has a knowledge base within its signature. Unless otherwise stated, all signatures and sentences are understood to be in first order logic. Agents with different but overlapping first order signatures may arise, for example, in an organization where the agents are experts in different areas of specialization, or in distributed networks where a large problem is broken up into small problems to be addressed by agents. A sentence in the language of an agent x is said to be report provable at x if x can verify D after the agents in the network report sentences to their neighbors in their common languages. A signature network is said to be report complete at x if whatever the agent knowledge bases are, every consequence of the union of these knowledge bases that is in the language of x is report provable at x. To obtain conditions for report completeness, the Craig interpolation theorem [4] is applied at each edge in the network.

The present paper considers situations that may arise in applications where there is some systematic relationship among the possible knowledge

Date: November 11, 2013.

bases and decision sentences. Specifically, we consider “observation networks” on a given signature network in which there are many potential observations that may be added to the underlying knowledge bases (i.e., learning), and a finite set of alternative sentences one of which must be proved (i.e., making a decision). We exploit the fact that report completeness holds for every knowledge base over a report complete signature network. We do this by introducing the notion of a “report plan” for an observation network — a finite scheme for finding a “report proof” for one of the alternative sentences once observations are added to the underlying knowledge bases. A report plan is decentralized; the agents only need to know their own observations and the reports they receive, not what happens elsewhere in the network. After the observations are made, a report plan is executed by a single pass through the network.

We give a brief preview in Section 2 and a summary of some facts from the literature in Section 3. In Section 4 we discuss single- and multiple-pass report proofs, and in Sections 5 and 6 we introduce the central notions of an observation network and a report plan. Our main result, Theorem 7.3, shows that an observation network will always have a report plan for a given finite set of alternatives provided that: (1) the observation network is sufficient for selecting from the set of alternatives, and (2) the underlying signature network contains a signature tree in the sense of the paper [7]. We also prove two results that are complementary to our main theorem. Theorem 5.6 is a “finiteness theorem” that shows that if an observation network is sufficient for selecting from a given infinite set of alternatives, then some finite part of the observation network is sufficient for selecting from some finite subset of the alternatives. Theorem 7.1 shows that a report plan guarantees that under every possible family of observations by the agents, some alternative has a single-pass report proof. In Proposition 8.5 we briefly indicate a way that report plans might be applied to obtain approximate values for unknown quantities depending on observations. At the end of the paper we give a list of open questions.

2. Preview

Before we formally define the main concepts in this paper, we give in Figure 1 a simple example that may help the reader to fix ideas. We will refer back to this example later on in this paper. In this example there are five agents arranged in a network as shown in the first picture. At the top is an agent d (the decider) who must decide between the two sentences $A, \neg A$ (indicated by a question mark). Agents y and z may report to w, agent z may also report to x, and agents w and x may report to d. Each agent is equipped with a set of sentences called its knowledge base. In addition, each agent has a set of potential observations that it might make. Before any observations are made, a report plan is prepared as in the picture. This plan gives a set of sentences that might be reported along each edge of the graph.

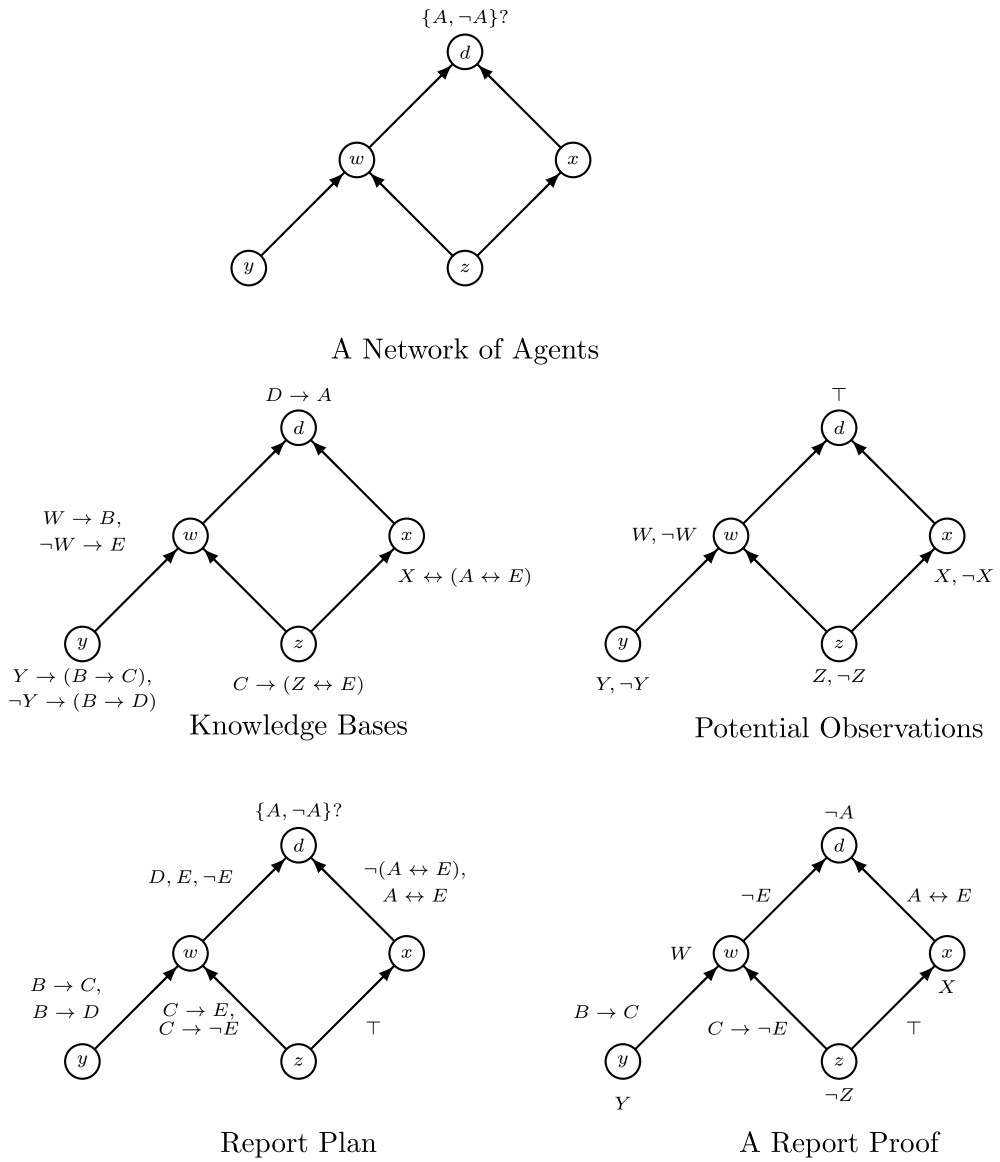

FIGURE 1. A Report Plan

Each agent will make one of its potential observations, independently of the other agents. Then the report plan is executed, building what we call a report proof. This is done with a single pass through the network, beginning at the bottom and ending at the top. Since four agents each have two potential observations, there are 16 possible cases where each agent makes an observation. One such case is shown in the last picture. Beginning at the bottom level, for each edge (r, s) in the graph, agent r proves a sentence from its knowledge base and observation and any sentences reported to it, and reports it to agent s. This sentence must be in the common language of the pair of agents involved. Finally, the decider d at the top is able to prove one of the alternatives, which happens to be $\neg A$ in this case.

We will now formally develop the ideas illustrated in the above example and obtain our results. A central feature is that agents report sentences to other agents in their common language. These sentences are Craig interpolants. Our arguments in this paper will apply results from the paper [7] that in turn depend on the Craig interpolation theorem.

3. Prerequisites

3.1. **Logic.** We assume familiarity with the notions of sentence, signature, and proof in first order logic with equality. For background, see [2] or [6]. The set of all first order sentences in a signature L is denoted by $[L]$. First order logic is formulated so that the true sentence $\top$ and false sentence $\bot$ belong to $[L]$ for every signature L. The notation $\mathcal{K} \vdash B$ means that the sentence B is provable from the set of sentences $\mathcal{K}$, and $\vdash B$ means that B is provable. A set of sentences $\mathcal{K}$ is **consistent** if it is not the case that $\mathcal{K} \vdash \bot$. Given a finite set $\mathcal{B}$ of sentences, $\bigvee \mathcal{B}$ is the disjunction of the sentences in $\mathcal{B}$, and $\bigwedge \mathcal{B}$ is the conjunction of the sentences in $\mathcal{B}$. The conjunction of the empty set of sentences is $\top$, and the disjunction of the empty set is $\bot$. We use $B \to D$ as an abbreviation for $\neg B \vee D$, and $B \leftrightarrow D$ for $(B \to D) \wedge (D \to B)$. Given a signature L and a set of sentences $\mathcal{K} \subseteq [L]$, a **complete** (or maximal consistent) extension of $\mathcal{K}$ (in L) is a set of sentences $\mathcal{M}$ such that $\mathcal{K} \subseteq \mathcal{M} \subseteq [L]$, $\mathcal{M}$ is consistent, and for each $B \in [L]$, either $B \in \mathcal{M}$ or $(\neg B) \in \mathcal{M}$. We will use the **Compactness Theorem** in the following form: *A set of sentences $\mathcal{K}$ is consistent if and only if every finite subset of $\mathcal{K}$ is consistent.*

3.2. **Graphs.** By a **(simple) directed graph** (V, E) we mean a non-empty finite set V of **vertices**, and a set $E \subseteq V \times V$ of **edges** (or arcs) (x, y) such that $x \neq y$. (We do not allow more than one edge from a vertex x to a vertex y, we do not allow edges from a vertex to itself, and we distinguish between the pair (x, y) and the pair (y, x).) When $(x, y) \in E$, we will say that x is a **child** of y and that y is a **parent** of x.

A **(directed) path** of length n from x to y is a sequence $(x_0, \ldots, x_n)$ of pairwise distinct vertices such that $x_0 = x, x_n = y$, and for each $i < n$, $(x_i, x_{i+1}) \in E$. (In particular, for each vertex x, (x) is a path of length 0 from x to itself.) A **source** is a vertex x such that there is no edge $(z, x) \in E$, and a **sink** is a vertex x such that there is no edge $(x, y) \in E$.

In a directed graph, by a **decider** we mean a vertex d such that for every other vertex x, there is at least one path from x to d. By a **pointed graph** we mean a directed graph (V, E) with at least one decider. Hereafter we will always assume that (V, E) is a pointed graph. We will always use the symbol d to denote a decider for (V, E). Note that we assume there is at least one decider, but allow the possibility that there is more than one decider.

A **directed cycle** of length n is a sequence $(x_0, \ldots, x_{n-1}, x_n)$ of vertices such that $(x_0, \ldots, x_{n-1})$ is a directed path, $x_n = x_0$, and $(x_{n-1}, x_n) \in E$ (hence $n \geq 2$).

By a **directed acyclic graph** we mean a pointed graph with no directed cycles. A directed acyclic graph has a unique decider, which is also the unique sink. In a directed acyclic graph, for every vertex x there is a path from a source to x.

By a **tree** we will mean a directed acyclic graph (V, E) such that for every vertex x there is at most one $(x, y) \in E$. In a tree, for every vertex x there is a unique path from x to the decider.

By a **connected symmetric graph** (or connected undirected graph) we mean a pointed graph (V, E) such that whenever $(x, y) \in E$ we also have $(y, x) \in E$. In a connected symmetric graph, every vertex is a decider.

3.3. **Signature and Knowledge Base Networks.** We now review some notions that we will need from the paper [7]. We attach signatures and knowledge bases to the vertices of pointed graphs. From now on we will call the vertices **agents**.

A **signature network** on (V, E) is an object

$$\mathbb{S} = (V, E, L(\cdot))$$

where (V, E) is a pointed graph with a labeling $L(\cdot)$ that assigns a first order signature $L(x)$ to each agent $x \in V$. We let $L(V) = \bigcup_{x \in V} L(x)$, and call $L(V)$ the **combined signature**. We say that a symbol S **occurs** at a vertex x if $S \in L(x)$.

We say that a signature network $\mathbb{S}$ **contains** a signature network $\mathbb{T} = (V, F, L(\cdot))$ if $\mathbb{S}$ and $\mathbb{T}$ have the same agents and signatures, and (V, F) can be obtained from (V, E) by removing edges.

Let $\mathbb{T} = (V, F, L(\cdot))$ be a signature network. We say that $\mathbb{T}$ is a **signature tree** if:

(1) (V, F) is a tree;
(2) for every pair of agents $x, y \in V$ and symbol S that occurs at both x and y, there is a vertex $z \in V$ such that S occurs at every vertex on a path from x to z and at every vertex on a path from y to z.

Note that Condition (2) allows the possibility that $z = x$ or $z = y$. Condition (2) occurs frequently in the literature, under the name "running intersection property" (e.g. see [3], [5]).

The example shown in Figure 1 in the Preview at the beginning of this paper is a knowledge base over a signature network $\mathbb{S}$ that contains the signature tree $\mathbb{T}$ shown in Figure 2. $\mathbb{T}$ is obtained from $\mathbb{S}$ by removing the edge (z, x).

A **knowledge base** is a set of first order sentences. Given a signature network $\mathbb{S} = (V, E, L(\cdot))$, a **knowledge base over** $\mathbb{S}$ (or knowledge base network over $\mathbb{S}$) is an object

$$\mathbb{K} = (V, E, L(\cdot), \mathcal{K}(\cdot))$$

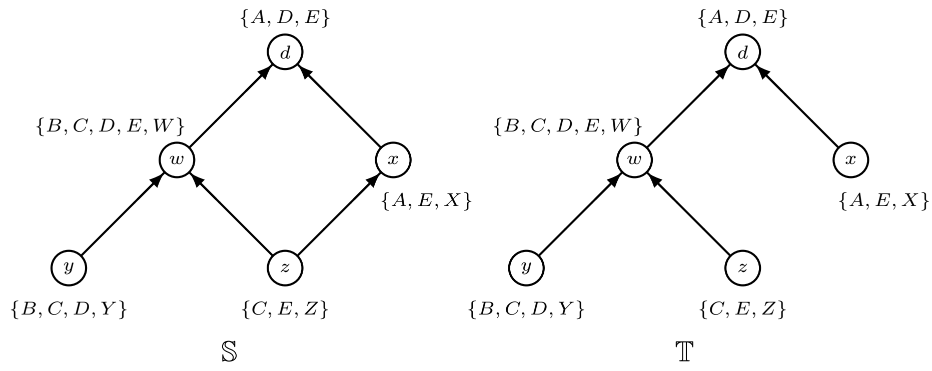

FIGURE 2. A Signature Network and Tree

where $\mathcal{K}(\cdot)$ is a labeling that assigns a knowledge base $\mathcal{K}(x) \subseteq [L(x)]$ to each agent $x \in V$. We write $\mathcal{K}(V) = \bigcup_{x \in V} \mathcal{K}(x)$, and we call the set $\mathcal{K}(V)$ the **combined knowledge base**.

Note that in a knowledge base $\mathbb{K}$ over $\mathbb{S}$, each symbol that occurs in a sentence in $\mathcal{K}(x)$ must belong to $L(x)$, but we allow the possibility that $L(x)$ also has additional symbols.

If a signature network $\mathbb{S} = (V, E, L(\cdot))$ contains a signature network $\mathbb{T} = (V, F, L(\cdot))$, then for any knowledge base $\mathbb{K}$ over $\mathbb{S}$ we get a knowledge base $\mathbb{K}_{\mathbb{T}}$ over $\mathbb{T}$ by replacing E by F and leaving everything else unchanged.

3.4. **Report Provability.** We summarize the concept of report provability and some related facts from [1] and [7].

Definition 3.1. *Let*

$$\mathbb{K} = (V, E, L(\cdot), \mathcal{K}(\cdot))$$

be a knowledge base over a signature network $\mathbb{S}$. *A sentence* C *is* 0-***reportable*** *in* $\mathbb{K}$ *along an edge* (x, y) *if*

$$C \in [L(x) \cap L(y)] \text{ and } \mathcal{K}(x) \vdash C.$$

C *is* $(n+1)$-***reportable*** *in* $\mathbb{K}$ *along an edge* (x, y) *if*

$$C \in [L(x) \cap L(y)] \text{ and } \mathcal{K}(x) \cup \mathcal{R} \vdash C,$$

where $\mathcal{R}$ *is a set of sentences each of which is* n*-reportable along some edge* (z, x).

The word "reportable" means "n*-reportable for some* n*", and "reportable to* x*" means "reportable along* (z, x) *for some* z*".*

Given a decider d *for* $\mathbb{S}$, *a sentence* $D \in [L(d)]$ *is* ***report provable*** *in* $\mathbb{K}$ *at* d *if* D *is provable from* $\mathcal{K}(d)$ *and a set* $\mathcal{R}$ *of sentences each of which is reportable to* d *in* $\mathbb{K}$.

This means that at each stage, for each edge (x, y), agent x can report to agent y a sentence C in their common language, where C is provable from the knowledge base $\mathcal{K}(x)$ and sentences reported to x at earlier stages.

Finally, D is provable from the knowledge base $\mathcal{K}(d)$ and sentences reported to d. Thus the sentence D is established using only proofs within the languages $[L(x)]$ of single agents x, and communications along edges (x, y) in the common language $[L(x) \cap L(y)]$.

The next fact shows that report provability implies provability.

Fact 3.2. *(Lemma 2.6 in* [7]*) Suppose d is a decider and a sentence $D \in [L(d)]$ is report provable at d in a knowledge base network $\mathbb{K}$. Then D is provable from the combined knowledge base, $\mathcal{K}(V) \vdash D$.*

Definition 3.3. *A knowledge base network $\mathbb{K}$ is **report complete** at a decider d if every sentence $D \in [L(d)]$ that is provable from the combined knowledge base $\mathcal{K}(V)$ is report provable in $\mathbb{K}$ at d. A signature network $\mathbb{S}$ is **report complete** at a decider d if every knowledge base $\mathbb{K}$ over $\mathbb{S}$ is report complete at d.*

The paper [7] deals with the following two questions. *Which signature networks are report complete at d? Which signature networks contain a signature tree at a decider d?*

Fact 3.4. *(Amir and McIlraith* [1]*) Every signature tree $\mathbb{S}$ is report complete at its unique decider.*

Corollary 3.5. *Every signature network $\mathbb{S}$ that contains a signature tree $\mathbb{T}$ with decider d is report complete at d.*

Proof. Let $\mathbb{K}$ be a knowledge base over $\mathbb{S}$. An easy induction shows that for each edge (x, y) in $\mathbb{T}$, every sentence that is n-reportable along (x, y) in $\mathbb{K}_{\mathbb{T}}$ is n-reportable along (x, y) in $\mathbb{K}$. By Fact 3.4, $\mathbb{K}_{\mathbb{T}}$ is report complete at d, so $\mathbb{K}$ is also report complete at d. □

Fact 3.6. *(Theorem 5.3 in* [7]*). Let $\mathbb{S}$ be a signature network on a directed acyclic graph and let d be its decider. $\mathbb{S}$ is report complete at d if and only if $\mathbb{S}$ contains a signature tree with decider d.*

Fact 3.7. *(Lemma 9.1 in* [7]*) Let $\mathbb{S}$ be a signature network on a connected symmetric graph. If for some agent x, $\mathbb{S}$ contains a signature tree with decider x, then for every agent y, $\mathbb{S}$ contains a signature tree with decider y.*

Fact 3.8. *(Theorem 9.7 in* [7]*). Let $\mathbb{S}$ be a signature network on a connected symmetric graph and let $d \in V$. $\mathbb{S}$ is report complete at d if and only if $\mathbb{S}$ contains a signature tree with decider d.*

As we mentioned in the Introduction, Theorem 7.3 will show that report plans at d always exist under the hypothesis that $\mathbb{S}$ is a signature network that contains a signature tree with decider d, like the example shown in the Preview. Facts 3.6—3.8 above show that this hypothesis is satisfied at d whenever $\mathbb{S}$ is over either a directed acyclic or a connected symmetric graph and is report complete at d. For this reason, we expect that the hypothesis will often be met in applications. Moreover, as noted in [7], the results hold not only for first order logic but for any logic that satisfies the compactness and Craig interpolation theorems and is closed under the Boolean connectives. In applications, signature networks that contain signature trees might also arise in various ways as a design goal. One possibility is suggested by an algorithm (called cut-cycles) in the paper [1] showing that for every signature network over a connected symmetric graph and every agent d, one can obtain a signature tree with decider d by removing edges and minimally adding symbols to signatures.

Our results in this paper deal with the situation where an agent x in a network receives reports from other agents and is faced with a decision or a family of related decisions. We have found it to be convenient to concentrate on the case that x is a decider, that is, for every agent y in the directed graph there is a path from y to x. But even if the agent x who is faced with a decision is not a decider in the directed graph, the results in this paper can still be applied in a slightly different manner. In that case, one can proceed as follows.

For the moment, we work with an arbitrary directed graph (V, E) which does not necessarily have a decider, and define the notions of a knowledge base network and report provability over (V, E) as before. For any signature network $\mathbb{S}$ over (V, E), knowledge base $\mathbb{K}$ over $\mathbb{S}$, and agent $x \in V$, one can always build a new directed graph (V_x, E_x), and a new signature network $\mathbb{S}_x$ over (V_x, E_x) and knowledge base $\mathbb{K}_x$ over $\mathbb{S}_x$ such that:

- x is a decider in (V_x, E_x);
- $\mathbb{S}$ contains a signature tree with decider x if and only if $\mathbb{S}_x$ contains a signature tree with decider x;
- a sentence is report provable at x in $\mathbb{K}$ if and only if it is report provable at x in $\mathbb{K}_x$.

All the results for the case where an agent x is a decider can be applied to the general situation where an arbitrary agent x is faced with a decision by passing to (V_x, E_x), $\mathbb{S}_x$, and $\mathbb{K}_x$.

The idea in building (V_x, E_x) is to ignore the agents z such that there is no path from z to x. Given a directed graph (V, E) and agent $x \in V$, let V_x be the set of all agents y such that (V, E) contains a path from y to x, and let $E_x = E \cap (V_x \times V_x)$. It is clear that x is a decider in (V_x, E_x). Now, for any signature network $\mathbb{S}$ over (V, E) and knowledge base $\mathbb{K}$ over $\mathbb{S}$, let $\mathbb{S}_x$ and $\mathbb{K}_x$ be the restrictions of $\mathbb{S}$ and $\mathbb{K}$ to (V_x, E_x). The properties listed above for $\mathbb{S}_x$ and $\mathbb{K}_x$ are easily verified by observing that there are no paths from agents in $V \setminus V_x$ to agents in V_x.

We now return to our original setting where (V, E) is always taken to be a pointed graph, and thus has at least one decider. While our main results will be about signature networks that contain signature trees, many of the examples in the following pages will have signature networks that do not contain signature trees, in order to illustrate the need for various assumptions in our results.

4. Single and Multiple Pass Report Proofs

The paper [7] introduced the notion of report provability, but refrained from introducing a notion of a report proof. In this section we introduce a kind of report proof. We will use this notion to explain what is meant by report provability by a single pass through the network, and give some examples where multiple passes are needed. Later on, in Section 6, we will develop the notion of a report plan, which is a plan that guarantees report provability by a single pass for a whole family of related knowledge base networks.

Informally, a report proof will consist of finitely many "reporting steps" denoted by (x, y, C) where an agent x reports a sentence C to a parent y, followed by one "deciding step" denoted by (x, x, D) where an agent x decides on a sentence D. In a reporting step (x, y, C), the reporting agent x first proves the sentence C from its knowledge base and the sentences that have been previously reported to x, and then reports C to y. In the deciding step, x proves D from its knowledge base and the sentences that have been reported to x. Here is the formal definition.

Definition 4.1. *By a **report proof** of a sentence D at a decider d in a knowledge base network $\mathbb{K}$ we mean a finite sequence*

$$\mathbb{R} = (x_1, y_1, C_1)(x_2, y_2, C_2) \cdots (x_k, y_k, C_k)$$

of edges labeled by sentences such that:

- *For each $i < k$ we have $(x_i, y_i) \in E$;*
- *$x_k = y_k = d$ and $C_k = D$;*
- *for each $i \leq k$, $C_i \in [L(x_i) \cap L(y_i)]$ and C_i is provable from $\mathcal{K}(x_i) \cup \mathcal{R}$ where $\mathcal{R}$ is the set of sentences previously reported to x_i in $\mathbb{R}$, i.e.,*

$$\mathcal{R} = \{C_j \, : \, j < i \quad and \quad y_j = x_i\}$$

*A report proof has a **single pass between** ℓ **and** m if the edges $(x_i, y_i), \ell \leq i \leq m$ are distinct, and for each agent $x \in V$, the i's such that $\ell \leq i \leq m$ and $x = x_i$ are consecutive.*

*A report proof has $\leq n$ **passes** if it can be broken up into at most n consecutive single passes.*

*$\mathbb{K}$ is n-**pass report complete** at d if every sentence in $[L(d)]$ that is provable from $\mathcal{K}(V)$ has a report proof with $\leq n$ passes in $\mathbb{K}$. A signature network $\mathbb{S}$ is n-**pass report complete** at d if every knowledge base network $\mathbb{K}$ over $\mathbb{S}$ is.*

Remark 4.2. *If there exists a report proof of D at d in $\mathbb{K}$, then D is report provable at d in $\mathbb{K}$.*

Proof. Let $\mathbb{R} = (x_1, y_1, C_1) \cdots (x_k, y_k, C_k)$ be a report proof of D at d in $\mathbb{K}$. Then $(x_k, y_k, C_k) = (d, d, D)$. An easy induction shows that for each $m < k$, C_m is $(m-1)$-reportable along (x_m, y_m) in $\mathbb{K}$. It follows that D is report provable at d in $\mathbb{K}$. □

Later on, we will show that the converse of Remark 4.2 is also true. First, we give some examples.

Example 4.3. Look again at Figure 1 in the Preview of this paper. In the picture labeled "A Report Proof" at the lower right, a knowledge base $\mathbb{K}^O$ is formed by adding each agent's observations to its original knowledge base:

$$\begin{aligned}
\mathbb{K}^O(d) &= \{D \to A\} \cup \{\top\}, \\
\mathbb{K}^O(w) &= \{W \to B, \neg W \to E\} \cup \{W\}, \\
\mathbb{K}^O(x) &= \{X \leftrightarrow (A \leftrightarrow E)\} \cup \{X\}, \\
\mathbb{K}^O(y) &= \{Y \to (B \to C), \neg Y \to (B \to D)\} \cup \{Y\}, \\
\mathbb{K}^O(z) &= \{C \to (Z \leftrightarrow E)\} \cup \{\neg Z\}.
\end{aligned}$$

The lower right picture shows the following single-pass report proof of $\neg A$ at d in $\mathbb{K}^O$, working from the bottom of the network to the top:

$$(y, w, B \to C)(z, w, C \to \neg E)(w, d, \neg E)(x, d, A \leftrightarrow E)(d, d, \neg A).$$

That is, y reports $B \to C$ to w, then z reports $C \to \neg E$ to w, then w reports $\neg E$ to d, then x reports $A \leftrightarrow E$ to d, and finally, d decides on $\neg A$.

We now give some examples of signature networks that are not single-pass report complete at a decider d. As we will see in Proposition 4.9 below, such a signature network cannot contain a signature tree with decider d. When $\mathbb{S}$ has a unique decider d, we sometimes omit the phrase "at d"; for instance, we may write "report complete" instead of "report complete at d".

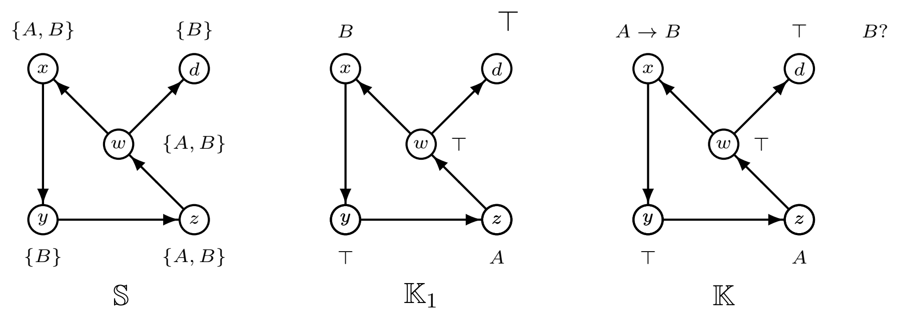

FIGURE 3. Example 4.4

Example 4.4. The signature network $\mathbb{S}$ shown in Figure 3 has a unique decider d and is two-pass report complete but not single-pass report complete. The knowledge base network $\mathbb{K}_1$ over $\mathbb{S}$ is single-pass report complete, but the knowledge base network $\mathbb{K}$ over $\mathbb{S}$ is not.

Proof. It is clear that $\mathbb{S}$ cannot contain a signature tree, because one would have to remove the edge (w, x), and then condition (2) in the definition of a signature tree would fail for the pair (x, z) and the symbol A. To prove that $\mathbb{S}$ is two-pass report complete at d, let $\mathbb{K}'$ be an arbitrary signature network over $\mathbb{S}$, let $D \in [L(d)] = \{\top, \bot, B, \neg B\}$, and suppose $\mathcal{K}'(V) \vdash D$. For each agent u, let K_u be the conjunction $\bigwedge \mathcal{K}'(u)$. Then

$$K_z \wedge K_w \wedge K_x \vdash K_y \to [K_d \to D].$$

Therefore D has the following two-pass the report proof, with the passes separated by a semicolon:

$$(z, w, K_z)(w, x, K_z \wedge K_w)(x, y, K_y \to [K_d \to D])(y, z, K_d \to D);$$

$$(z, w, K_d \to D)(w, d, K_d \to D)(d, d, D).$$

$\mathbb{K}_1$: We have $\mathcal{K}_1(V) = \{A, B\}$. The only non-trivial sentence $D \in [L(d)]$ such that $\{A, B\} \vdash D$ is B. B has the single-pass report proof

$$(x, y, B)(y, z, B)(z, w, B)(w, d, B)(d, d, B).$$

$\mathbb{K}$: We show that $\mathbb{K}$ is not single-pass report complete. We have

$$\mathcal{K}(V) = \{A, A \to B, \top\} \vdash B,$$

so there exists a report proof $\mathbb{R}$ of B at d in $\mathbb{K}$. Let us call a step (x_n, y_n, C_n) in $\mathbb{R}$ **significant** if the sentence C_n is not logically equivalent to $\top$. By working back from the end, one can see from the diagram that $\mathbb{R}$ must contain significant steps along the following edges in reverse order:

$$(w, d), (z, w), (y, z), (x, y), (w, x), (z, w).$$

Therefore $\mathbb{R}$ cannot be a single-pass report proof, and in fact, $\mathbb{R}$ must contain more than one significant step along the edge (z, w). □

Here is an example that requires more than two passes.

Example 4.5. The signature network $\mathbb{S}$ shown in Figure 4 has the unique decider d, and is three-pass report complete but not two-pass report complete. In the knowledge base $\mathbb{K}$, the sentence C is report provable in three passes but not in two passes. C has a three-pass report proof in $\mathbb{K}$ with only one step along each edge.

Using the same idea with a chain of n triangles, we can construct a signature network that is $(n+1)$-pass report complete but not n-pass report complete at its unique decider.

In our next example, the signature network is three-pass report complete, but is not two-pass report complete because a single agent must act at three different times during a report proof.

Example 4.6. In the knowledge base network $\mathbb{K}$ shown in Figure 5, the sentence C has the following three-pass report proof at d:

$$(w, x, A)(x, y, B)(y, w, B); (w, u, B)(u, v, C)(v, w, C); (w, d, C)(d, d, C),$$

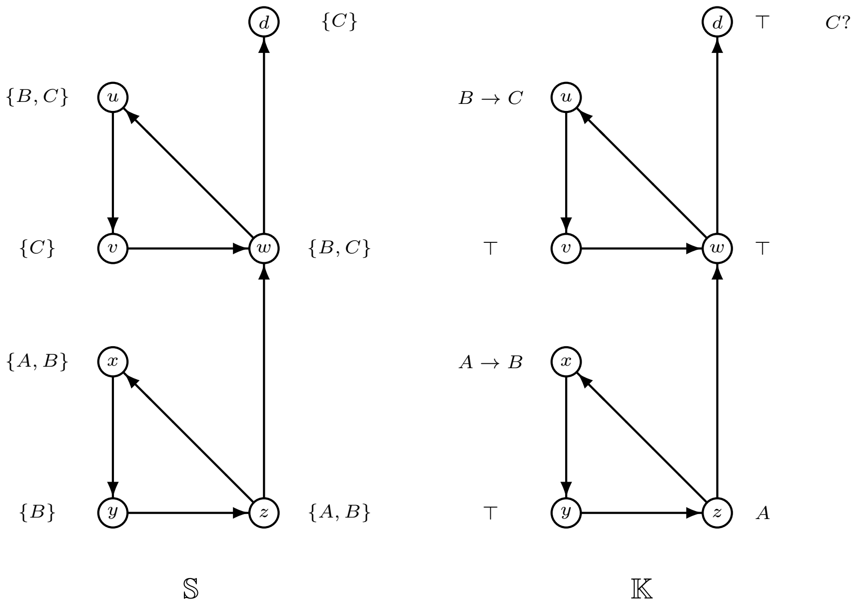

FIGURE 4. Example 4.5

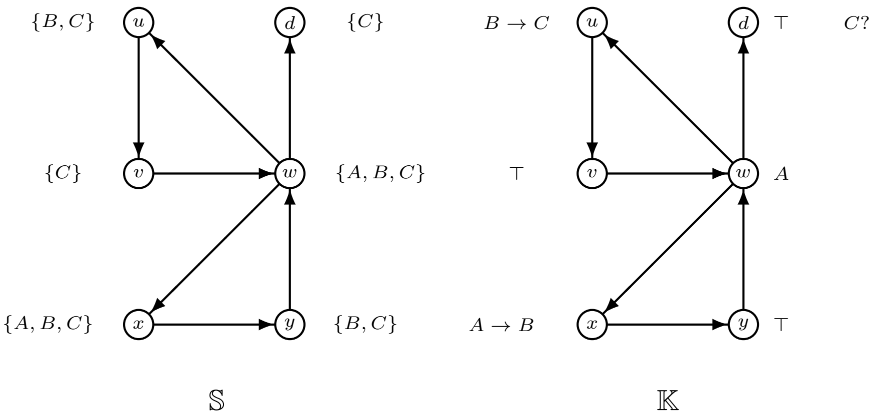

FIGURE 5. Example 4.6

but does not have a two-pass report proof at d because any report proof must contain a sequence of significant steps along the above edges.

The same idea with n triangles attached to w will give a signature network in which w must act more than n times in a report proof.

Here is an example of a knowledge base network that is report complete but there is no n such that it is n-pass report complete. However, the underlying signature network is not report complete. The analogous question for signature networks (Question 9.1) is open.

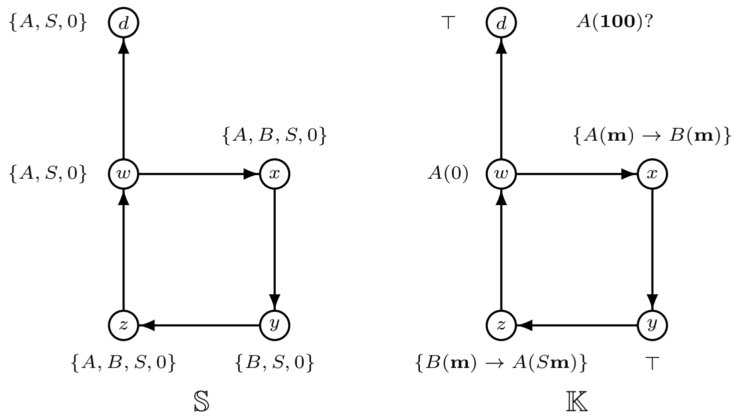

FIGURE 6. Example 4.7

Example 4.7. In Figure 6, $A(\cdot)$ and $B(\cdot)$ are unary predicate symbols, S is the successor function symbol, 0 is a constant symbol, and $\mathbf{m}$ is the term $S^m 0$. In $\mathbb{K}$, the knowledge bases at x and z are infinite sets of sentences, one for each $\mathbf{m}$. The signature network $\mathbb{S}$ is not report complete at the decider d, but the knowledge base network $\mathbb{K}$ is report complete at d. The sentence $A(\mathbf{100})$ is 101-pass report provable but not 100-pass report provable at d in $\mathbb{K}$. Using n instead of 100, we see that for any n, $\mathbb{K}$ is not n-pass report complete at d.

$\mathbb{S}$ is not report complete. To see this, let $\mathbb{K}'$ be the knowledge base over $\mathbb{S}$ with $\mathcal{K}'(x) = \{\forall t(A(t) \to B(t))\}$, $\mathcal{K}'(z) = \{\forall t(B(t) \to A(St))\}$, and with the other knowledge bases being $\{\top\}$. Then $\mathcal{K}'(V) \vdash \forall t(A(t) \to A(St))$ but $\forall t(A(t) \to A(St))$ is not report provable at d in $\mathbb{K}'$.

We now show that report provability lives up to its name.

Proposition 4.8. *In a knowledge base network $\mathbb{K}$, a sentence D is report provable at a decider d if an only if there exists a report proof of D at d.*

Proof. Remark 4.2 gives one direction. For the other direction, we first prove the following by induction on m:

In $\mathbb{K}$, if a sentence C is m-reportable along an edge (x, y), then for any decider d there is a report proof of $\top$ at d that contains the triple (x, y, C).

By definition, if C is 0-reportable along (x, y) then $(x, y, C)(d, d, \top)$ is a report proof. Assume the result holds for m and suppose C is $(m+1)$-reportable along (x, y). This means that $C \in [L(x) \cap L(y)]$ and $\mathcal{K}(x) \cup \mathcal{R} \vdash C$

where $\mathcal{R}$ is a set of sentences that are m-reportable to x. By the inductive hypothesis, for each $B \in \mathcal{R}$ there is a child z of x and a report proof $\mathbb{R}_B$ of $\top$ at d containing (z, x, B). By removing the last deciding step from each $\mathbb{R}_B$, stringing them together in any order, and adding $(x, y, C)(d, d, \top)$ at the end, we obtain a report proof containing (x, y, C). This completes the induction.

If D is report provable at d, then $D \in [L(d)]$ and D is provable from $\mathcal{K}(d)$ and sentences reportable to d. Arguing as in the preceding paragraph, we then obtain a report proof of D at d. □

The next result is an improvement of Corollary 3.5. It could be proved by an easy modification of the proof of Fact 3.4, but we will instead prove it as a consequence of Fact 3.4.

Proposition 4.9. *Every signature network $\mathbb{S}$ that contains a signature tree $\mathbb{T}$ with decider d is single-pass report complete at d.*

Proof. Let $\mathbb{K}$ be a knowledge base over $\mathbb{S}$, $A \in [L(d)]$, and $\mathcal{K}(V) \vdash A$. By Fact 3.4, $\mathbb{T}$ is report complete at d, so A is report provable at d in $\mathbb{K}_{\mathbb{T}}$. By Proposition 4.8, A has a report proof $\mathbb{R}$ at d in $\mathbb{K}_{\mathbb{T}}$. Since $\mathbb{T}$ is a tree, the set V of agents can be put in a list $(u_1, \ldots, u_n)$ such that $u_n = d$, and whenever there is a path from u_i to u_j in $\mathbb{T}$ we have $i \leq j$. The steps in $\mathbb{R}$ can be rearranged into a sequence $\mathbb{R}'$ so that all the steps with reporter u_1 come first, then all the steps with reporter u_2, and so on, ending with the deciding step (d, d, A). Since $i \leq j$ whenever there is a path from u_i to u_j in $\mathbb{T}$, $\mathbb{R}'$ is still a report proof in $\mathbb{K}_{\mathbb{T}}$ at d. Now combine the steps along each edge into a single step by taking the conjunction of the reported sentences. This gives us a single-pass report proof $\mathbb{R}''$ of A at d in $\mathbb{K}_{\mathbb{T}}$. $\mathbb{R}''$ is also a single-pass report proof of A at d in $\mathbb{K}$, as required. □

5. Observation Networks

An observation network is a signature network where each agent has both a knowledge base and a set of sentences called potential observations.

The story: An organization has a finite set of agents arranged in a network indexed by a graph, with a decider d. Each agent x has a signature $L(x)$ and a knowledge base $\mathcal{K}(x) \subseteq [L(x)]$. Each agent x also has a set of potential observations $\mathcal{O}(x) \subseteq [L(x)]$, and the decider d has a set $\mathcal{A} \subseteq [L(d)]$ of alternatives. In informal discussions in this paper, we will use the word "scenario" to mean a possible state of the world at the time that the agents make their observations. The organization is faced with a recurring situation where in every scenario, each agent x makes observations in $\mathcal{O}(x)$, and there is an alternative $A \in \mathcal{A}$ that is "correct" in the sense that it is provable from the combined knowledge base and the observations of the agents.

In Theorem 5.6 below we will use the Compactness Theorem to show that there must be a finite subset $\mathcal{A}_0 \subseteq \mathcal{A}$ and, for each agent $x \in V$, a finite set $\mathcal{O}_0(x) \subseteq \mathcal{O}(x)$, such that in every scenario, each agent x makes one observation in $\mathcal{O}_0(x)$, and some alternative $A \in \mathcal{A}_0$ is correct. We first formally state the definitions.

We say that a set $\mathcal{B}$ of sentences is **closed under finite conjunctions** if $\top \in \mathcal{B}$ and for each non-empty finite subset $\mathcal{B}_0 \subseteq \mathcal{B}$, the conjunction $\bigwedge \mathcal{B}_0$ is logically equivalent to some sentence in $\mathcal{B}$.

Remark 5.1.

- *If $\top \in \mathcal{B}$, and for each $A, B \in \mathcal{B}$ there exists $C \in \mathcal{B}$ that is logically equivalent to $A \wedge B$, then $\mathcal{B}$ is closed under finite conjunctions.*
- *If $\mathcal{B}$ is closed under finite conjunctions and $\mathcal{B}_0 \subseteq \mathcal{B}$ is finite, then there is a finite set $\mathcal{B}_1$ such that $\mathcal{B}_0 \subseteq \mathcal{B}_1 \subseteq \mathcal{B}$ and $\mathcal{B}_1$ is closed under finite conjunctions.*

Definition 5.2. *Given a knowledge base network*

$$\mathbb{K} = (V, E, L(\cdot), \mathcal{K}(\cdot)),$$

*an **observation network** (over $\mathbb{K}$) is an object*

$$\mathbb{O} = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}(\cdot))$$

consisting of $\mathbb{K}$ and a labeling $\mathcal{O}(\cdot)$ of (V, E) such that for each agent x in V, $\mathcal{O}(x) \subseteq [L(x)]$, and $\mathcal{O}(x)$ is closed under finite conjunctions.

The elements of $\mathcal{O}(x)$ are called **potential observations for** x. We write $L(V) = \bigcup_{x \in V} L(x)$ for the combined signature, $\mathcal{K}(V) = \bigcup_{x \in V} \mathcal{K}(x)$ for the combined knowledge base, and $\mathcal{O}(V) = \bigcup_{x \in V} \mathcal{O}(x)$ for the combined set of potential observations.

Definition 5.3. *A **finite observation network** is an observation network*

$$\mathbb{O}_0 = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}_0(\cdot))$$

such that the combined set of potential observations $\mathcal{O}_0(V)$ is finite.

Given an observation network

$$\mathbb{O} = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}(\cdot)),$$

*a **finite part** of $\mathbb{O}$ is a finite observation network*

$$\mathbb{O}_0 = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}_0(\cdot))$$

over the same knowledge base network such that $\mathcal{O}_0(x) \subseteq \mathcal{O}(x)$ for every agent $x \in V$.

Definition 5.4. *By a set of **alternatives** for an observation network $\mathbb{O}$*

$$\mathbb{O} = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}(\cdot))$$

*we will mean a non-empty set of sentences in $[L(V)]$. $\mathbb{O}$ is **sufficient** for a set of alternatives $\mathcal{A}$ if for every complete extension $\mathcal{M}$ of $\mathcal{K}(V)$, there exists an alternative $A^{\mathcal{M}} \in \mathcal{A}$ such that*

$$\mathcal{K}(V) \cup (\mathcal{M} \cap \mathcal{O}(V)) \vdash A^{\mathcal{M}}.$$

Definition 5.5. *By an **observation** in $\mathbb{O}$ we mean a function $O(\cdot)$ such that $O(x) \in \mathcal{O}(x)$ for each agent $x \in V$. We will write $O(V)$ for the sentence $\bigwedge_{x\in V} O(x)$. An observation $O(\cdot)$ in $\mathbb{O}$ is called **consistent** if $O(V)$ is consistent with $\mathcal{K}(V)$.*

For each observation $O(\cdot)$ in $\mathbb{O}$, let $\mathcal{K}^O(x) = \mathcal{K}(x) \cup \{O(x)\}$ and form the knowledge base network

$$\mathbb{K}^O = (V, E, L(\cdot), \mathcal{K}^O(\cdot))$$

by adding the sentence $O(x)$ to the knowledge base $\mathcal{K}(x)$ for each $x \in V$.

The following result is an application of the Compactness Theorem.

Theorem 5.6. *Let*

$$\mathbb{O} = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}(\cdot))$$

be an observation network, and let $\mathcal{A}$ be a set of alternatives for $\mathbb{O}$. Then the following are equivalent:

(i) *$\mathbb{O}$ is sufficient for $\mathcal{A}$;*
(ii) *there exists a finite set $\mathcal{A}_0 \subseteq \mathcal{A}$ and a finite part $\mathbb{O}_0$ of $\mathbb{O}$ such that $\mathbb{O}_0$ is sufficient for $\mathcal{A}_0$.*
(iii) *there are finitely many observations $O_1(\cdot), \ldots, O_n(\cdot)$ in $\mathbb{O}$ and finitely many alternatives $A_1, \ldots, A_n$ in $\mathcal{A}$ such that:*

(1) $$\mathcal{K}(V) \vdash O_1(V) \vee \cdots \vee O_n(V);$$

(2) $$\text{for each } k \leq n,\ \mathcal{K}(V) \vdash O_k(V) \to A_k.$$

Discussion: In this result, all that matters is the combined knowledge base $\mathcal{K}(V)$ and observation set $\mathcal{O}(V)$. The knowledge bases and observations of the individual agents will play a role later on in this paper. The formal counterpart of a possible scenario is a complete extension $\mathcal{M}$ of the combined knowledge base $\mathcal{K}(V)$. So Condition (i) says that in every possible scenario, some alternative $A \in \mathcal{A}$ can be proved from the combined knowledge bases and the observations of the agents. Condition (ii) says that there are predetermined finite sets of observations $\mathcal{O}_0(V)$ and of alternatives $\mathcal{A}_0$ such that in every possible scenario, some alternative in $\mathcal{A}_0$ can be proved from the combined knowledge bases and the observations of the agents in $\mathcal{O}_0(V)$.

Condition (iii) gives a characterization of sufficiency that does not mention complete extensions. (1) says that in every scenario, one of the observations $O_m(\cdot)$ will be made. (2) says that the alternative A_m can be proved from the combined knowledge base and the observations $O_m(x), x \in V$. Note that in (iii), one can remove any inconsistent observations $O_m(\cdot)$, so each of the observations $O_m(\cdot)$ can be taken to be consistent.

Proof of Theorem 5.6. (iii) $\Rightarrow$ (i): Assume (iii). Let $\mathcal{M}$ be a complete extension of $\mathcal{K}(V)$. By (1), $\mathcal{M} \vdash O_k(V)$ for some $k \leq n$, so by (2) we have $\mathcal{M} \vdash A_k$. Since $A_k \in \mathcal{A}$, this proves (i).

(i) $\Rightarrow$ (ii): Assume (i). By (i), for each complete extension $\mathcal{M}$ of $\mathcal{K}(V)$ there is an alternative $A^{\mathcal{M}} \in \mathcal{A}$ such that

$$\mathcal{K}(V) \cup (\mathcal{M} \cap \mathcal{O}(V)) \vdash A^{\mathcal{M}}.$$

By the Compactness Theorem, for each complete extension $\mathcal{M}$ of $\mathcal{K}(V)$, there is a finite set $\mathcal{O}^{\mathcal{M}} \subseteq (\mathcal{M} \cap \mathcal{O}(V))$ such that

$$\mathcal{K}(V) \cup \mathcal{O}^{\mathcal{M}} \vdash A^{\mathcal{M}}.$$

By the Compactness Theorem again, every consistent set of sentences in $L(V)$ has at least one complete extension. The set of sentences

$$\mathcal{K}(V) \cup \{\neg \bigwedge \mathcal{O}^{\mathcal{M}} : \mathcal{M} \text{ is a complete extension of } \mathcal{K}(V)\}$$

does not have a complete extension, and hence is not consistent. By the Compactness Theorem yet again, this set has a finite subset that is inconsistent. Therefore there are complete extensions $\mathcal{M}_1, \ldots, \mathcal{M}_n$ of $\mathcal{K}(V)$ such that

$$\mathcal{K}(V) \cup \{\neg O_1, \ldots, \neg O_n\}$$

is inconsistent, where $O_i = \bigwedge \mathcal{O}^{\mathcal{M}_i}$. Then

$$\mathcal{K}(V) \vdash O_1 \vee \cdots \vee O_n.$$

Let

$$\mathcal{A}_0 = \{A^{\mathcal{M}_1}, \cdots, A^{\mathcal{M}_n}\}.$$

Then $\mathcal{A}_0$ is a finite subset of $\mathcal{A}$. By Remark 5.1, for each agent $x \in V$ there is a finite set of sentences $\mathcal{O}_0(x)$ such that

$$\mathcal{O}(x) \supseteq \mathcal{O}_0(x) \supseteq \mathcal{O}(x) \cap (\mathcal{O}^{\mathcal{M}_1} \cup \cdots \cup \mathcal{O}^{\mathcal{M}_n}),$$

and $\mathcal{O}_0(x)$ is closed under finite conjunctions. Then

$$\mathbb{O}_0 = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}_0(\cdot))$$

is a finite part of $\mathbb{O}$. For each complete extension $\mathcal{N}$ of $\mathcal{K}(V)$ we have $\mathcal{N} \vdash O_k$ for some $k \leq n$, and therefore

$$(\mathcal{K}(V) \cup \mathcal{O}^{\mathcal{M}_k}) \subseteq (\mathcal{N} \cap \mathcal{O}_0(V)).$$

We have $\mathcal{K}(V) \vdash O_k \to A^{\mathcal{M}_i}$, so $\mathcal{K}(V) \cup \mathcal{O}^{\mathcal{M}_k} \vdash A^{\mathcal{M}_k}$. It follows that

$$(\mathcal{N} \cap \mathcal{O}_0(V)) \vdash A^{\mathcal{M}_k}$$

and $A^{\mathcal{M}_k} \in \mathcal{A}_0$. Then $\mathbb{O}_0$ is sufficient for $\mathcal{A}_0$, and (ii) is proved.

(ii) $\Rightarrow$ (iii): Assume (ii). Since $\mathcal{O}_0(V)$ is finite, there are finitely many subsets $\mathcal{P}_1, \ldots, \mathcal{P}_n$ of $\mathcal{O}_0(V)$ such that:

(a) For every complete extension $\mathcal{M}$ of $\mathcal{K}(V)$, $\mathcal{M} \cap \mathcal{O}_0(V) \in \{\mathcal{P}_1, \ldots, \mathcal{P}_n\}$;
(b) for each $k \leq n$, there is a complete extension $\mathcal{M}_k$ of $\mathcal{K}(V)$ such that $\mathcal{M}_k \cap \mathcal{O}_0(V) = \mathcal{P}_k$.

For each $k \leq n$ and $x \in V$, let $O_k(x)$ be a sentence in $\mathcal{O}(x)$ that is logically equivalent to $\bigwedge(\mathcal{P}_k \cap \mathcal{O}_0(x))$. Then $O_k(V)$ is logically equivalent to $\bigwedge \mathcal{P}_k$. Condition (1) follows easily from (a). By (b), each $O_k(\cdot)$ is a consistent observation in $\mathbb{O}$. By (ii), $\mathbb{O}_0$ is sufficient for $\mathcal{A}_0$, so by (b), for each $k \leq n$ there exists $A_k \in \mathcal{A}_0$ such that

$$\mathcal{K}(V) \cup (\mathcal{M}_k \cap \mathcal{O}_0(V)) = \mathcal{K}(V) \cup \mathcal{P}_k \vdash A_k,$$

and (2) follows. This proves (iii). □

6. Report Plans

In the rest of this paper, we will focus on what happens when information is passed between neighboring agents in an observation network. In this section we will formally define the notion of a report plan. The intuitive idea of a report plan and a simple example were given in the Preview to this paper.

Suppose $\mathbb{S}$ is a signature network. Then for every observation network

$$\mathbb{O} = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}(\cdot))$$

over $\mathbb{S}$ and every observation $O(\cdot)$ in $\mathbb{O}$, $\mathbb{K}^O$ is a knowledge base network over $\mathbb{S}$. In this way, each observation network over $\mathbb{S}$ gives rise to a whole family of knowledge base networks over $\mathbb{S}$. So if $\mathbb{S}$ is report complete, then for every knowledge base network $\mathbb{K}^O$ in this family, every sentence that is provable from the combined knowledge base is report provable.

By a set of **alternatives at** d we mean a non-empty set of sentences in $[L(d)]$. The following is an immediate consequence of Theorem 5.6

Corollary 6.1. *Let*

$$\mathbb{O} = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}(\cdot))$$

be an observation network over a signature network $\mathbb{S}$ that is report complete at d. Let $\mathcal{A} \subseteq [L(d)]$ be a set of alternatives at a decider d. Suppose that $\mathbb{O}$ is sufficient for $\mathcal{A}$. Then there are finitely many observations $O_1(\cdot), \ldots, O_n(\cdot)$ in $\mathbb{O}$ and finitely many alternatives $A_1, \ldots, A_n$ in $\mathcal{A}$ such that:

(a) $\mathcal{K}(V) \vdash (O_1(V) \vee \cdots \vee O_n(V))$;
(b) *for each $k \leq n$, A_k is report provable in $\mathbb{K}^{O_k}$ at d.*

This corollary shows that there is a predetermined finite set of observations $O_k(\cdot)$ and alternatives A_k, $k \leq m$, such that in every scenario,

- one of the observations $O_k(\cdot)$ will be made;
- the alternative A_k will be correct in the sense that it is provable from $\mathcal{K}(V) \cup O_k(V)$;
- reports between neighboring agents will propagate through the network, and then the decider will be able to prove the alternative A_k from its knowledge base and the sentences that are reported to it.

A set of sentences $\mathcal{B}$ is called **Boolean closed** if $\mathcal{B}$ is closed under finite conjunctions, and for any $B \in \mathcal{B}$ the negation $\neg B$ is logically equivalent to a sentence in $\mathcal{B}$. An observation network $\mathbb{O}$ is called **Boolean closed** if $\mathcal{O}(x)$ is Boolean closed for every agent $x \in V$.

In the case that the observation network is Boolean closed and its signature network contains a signature tree with decider d, one can improve upon Corollary 6.1. We will show that in that case there must be a finite plan, called a report plan, that will always enable the decider to arrive at a correct alternative. For each agent x and edge (x, y), a report plan will provide x with a finite rule that gives a sentence C to report along (x, y), where C depends only on the observation by x and the reports received by x. For any potential observations by the agents, the report plan will produce a single-pass report proof of one of the alternatives.

A finite set of sentences $\{C_1, \dots, C_n\}$ is a **partition** if each sentence alone is logically consistent, the sentences are mutually exclusive, and their disjunction is logically valid, that is,

$$C_i \nvdash \bot \text{ for each } i, \quad \vdash \neg(C_i \wedge C_j) \text{ whenever } i < j, \text{ and } \vdash C_1 \vee \cdots \vee C_n.$$

Definition 6.2. *Let $\mathbb{O}$ be an observation network and $\mathcal{A}$ be a set of alternatives at a decider d. A **report plan** for $(\mathbb{O}, \mathcal{A}, d)$ is an object*

$$\mathbb{P} = (\mathcal{O}_0(\cdot), \mathcal{C}_0(\cdot), \mathcal{A}_0)$$

such that:

(1) *for each agent $x \in V$, $\mathcal{O}_0(x) \subseteq \mathcal{O}(x)$ and $\mathcal{O}_0(x)$ is a partition;*
(2) *for each edge $(x, y) \in E$, $\mathcal{C}_0(x, y)$ is a finite set of sentences in the common language $[L(x) \cap L(y)]$,*
(3) *$\mathcal{A}_0$ is a finite subset of $\mathcal{A}$;*
(4) *by a **potential report** (to x) we mean a set $\mathcal{R}(x)$ of sentences consisting of one element of $\mathcal{C}_0(z, x)$ for each edge (z, x) ($\mathcal{R}(x)$ is empty if x is a source);*
(5) *for each edge $(x, y) \in E$, observation $O(x) \in \mathcal{O}_0(x)$, and potential report $\mathcal{R}(x)$, there is a sentence $C \in \mathcal{C}_0(x, y)$ such that $\mathcal{K}(x) \cup \{O(x)\} \cup \mathcal{R}(x) \vdash C$;*
(6) *for each observation $O(d) \in \mathcal{O}_0(d)$ and potential report $\mathcal{R}(d)$, there is a sentence $A \in \mathcal{A}_0$ such that $\mathcal{K}(d) \cup \{O(d)\} \cup \mathcal{R}(d) \vdash A$;*
(7) *$\mathbb{P}$ **avoids cycles**, that is, every directed cycle contains an edge e such that $\mathcal{C}_0(e) = \{\top\}$.*

*By an **observation in** $\mathbb{P}$ we mean an observation $O(\cdot)$ in $\mathbb{O}$ such that $O(x) \in \mathcal{O}_0(x)$ for each agent x. By a **report proof in** $\mathbb{P}$ from an observation $O(\cdot)$ in $\mathbb{P}$ we mean a single-pass report proof $\mathbb{R}$ in $\mathbb{K}^O$ such that $C \in \mathcal{C}_0(x, y)$ for each triple $(x, y, C) \in \mathbb{R}$.*

Remark 6.3. *Let $\mathbb{P}$ be a report plan for $(\mathbb{O}, \mathcal{A}, d)$. Then for each complete extension $\mathcal{M}$ of $\mathcal{K}(V)$ there is a unique observation $O^{\mathbb{P},\mathcal{M}}(\cdot)$ in $\mathbb{P}$ such that*

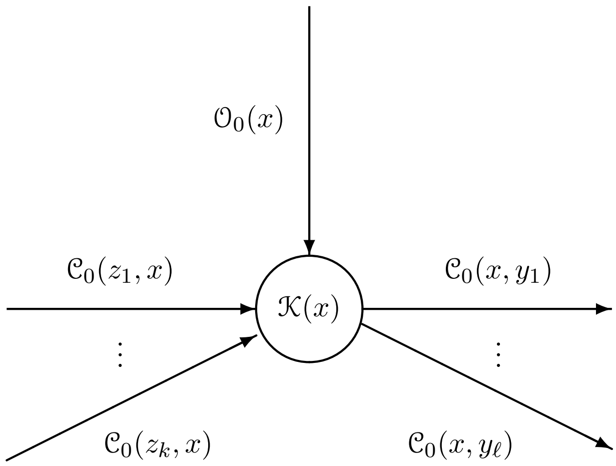

FIGURE 7. Checklist for agent x.

$\mathcal{M} \vdash O^{\mathbb{P},\mathcal{M}}(V)$. *This follows from the fact that each set $\mathcal{O}_0(x)$ in a report plan is a partition.*

The story: Intuitively, a report plan provides each agent x with a checklist consisting of a finite partition $\mathcal{O}_0(x)$ of potential observations, and finite sets of sentences $\mathcal{C}_0(z, x)$ and $\mathcal{C}_0(x, y)$ for each child z of x and each parent y of x (see Figure **??**). From the point of view of agent x, $\mathcal{C}_0(z, x)$ is a finite set of "question sentences" to watch for, and $\mathcal{C}_0(x, y)$ is a finite set of "answer sentences" to try to prove. (Thus for each edge (x, y), $\mathcal{C}_0(x, y)$ is both a set of questions asked by y and a set of possible answers by x). The decider d is also provided with a finite set of alternatives $\mathcal{A}_0$. Meaningful reports can only be passed along edges (x, y) such that $\mathcal{C}_0(x, y) \neq \{\top\}$.

Remark 6.3 says that in each scenario $\mathcal{M}$, every agent x makes exactly one observation $O^{\mathbb{P},\mathcal{M}}(x) \in \mathcal{O}_0(x)$. Each agent x asks each of its children z the finite set of questions $\mathcal{C}_0(z, x)$ and receives an answer in this set. These answers together form one of the potential reports $\mathcal{R}(x)$ to x, which we will call the report received by x. Then for each parent y of x, agent x proves one of the sentences in $\mathcal{C}_0(x, y)$ from its knowledge base $\mathcal{K}(x)$, its observation $O^{\mathbb{P},\mathcal{M}}(x)$, and the report $\mathcal{R}(x)$ it receives, and reports this sentence as its answer to y. Finally, the decider d receives a report $\mathcal{R}(d)$ from its children and proves an alternative in $\mathcal{A}_0$ from its knowledge base $\mathcal{K}(d)$, its observation $O^{\mathcal{M}}(d)$, and $\mathcal{R}(d)$.

Using the fact that the report plan avoids cycles, we will see that this process will always provide a single-pass report proof of one of the alternatives.

In this process, each agent x selects an answer sentence $C(x,y) \in \mathcal{C}_0(x,y)$ for each edge (x,y), and the decider d selects an alternative $A \in \mathcal{A}$. To do its part, each agent x only needs its own knowledge base $\mathcal{K}(x)$ and observation $O^{\mathbb{P},\mathcal{M}}(x)$, and the sets of sentences $\mathcal{C}_0(z,x)$ for each child z and $\mathcal{C}_0(x,y)$ for each parent y. The agents do not need to know the observations or question sentences of the other agents. This process is a particular realization of the report plan where each agent makes an observation, receives a single potential report from its children, and proves one of finitely many possible answers to report to each of its parents.

What does a report plan look like? We have already seen an example of a report plan $\mathbb{P}$ in the Preview at the beginning of this paper (Figure 1, revisited in Example 4.3). Let us take a closer look.

Example 6.4. The observation network shown in Figure 1 is Boolean closed (provided that we put $\mathcal{O}(d) = \{\top, \bot\}$ instead of $\mathcal{O}(d) = \{\top\}$). The lower left diagram shows a report plan $\mathbb{P}$. The lower right diagram shows the report proof

$$\mathbb{R} = (y, w, B \to C)(z, w, C \to \neg E)(w, d, \neg E)(x, d, A \leftrightarrow E)(d, d, \neg A)$$

in $\mathbb{P}$ of $\neg A$ at d from the observation O, where

$$O(y) = Y, \quad O(z) = \neg Z, \quad O(w) = W, \quad O(x) = X, \quad O(d) = \top.$$

Let us examine this report plan from the viewpoint of agent w. Agent w has signature, knowledge base, and potential observations

$$L(w) = \{W, B, C, D, E\}, \quad \mathcal{K}(w) = \{W \to B, \neg W \to E\}, \quad \mathcal{O}(w) = \{W, \neg W\}.$$

In the report plan $\mathbb{P}$, w's checklist consists of the sets of potential observations $\mathcal{O}(w)$, question sentences

$$\mathcal{C}_0(y,w) = \{B \to C, B \to D\}, \quad \mathcal{C}_0(z,w) = \{C \to E, C \to \neg E\},$$

and possible answers

$$\mathcal{C}_0(w,d) = \{D, E, \neg E\}.$$

To verify that $\mathbb{P}$ actually is a report plan, the non-trivial things to check are Conditions (5) and (6) in Definition 6.2. For instance, Condition (5) for the edge (w,d) says that for each of the two potential observations for w and four potential reports to w, some sentence in $\mathcal{C}_0(w,d)$ is deducible from w's knowledge base and observation and the potential report. So there are eight things to verify for the edge (w,d). Here are two of them.

$$\{W \to B, \neg W \to E\} \cup \{W\} \cup \{B \to C, C \to E\} \vdash E,$$

$$\{W \to B, \neg W \to E\} \cup \{W\} \cup \{B \to C, C \to \neg E\} \vdash \neg E.$$

Why do we require that a report plan avoids cycles? A key requirement for a report plan is that for every observation $O(\cdot)$ in $\mathbb{P}$, some alternative $A \in \mathcal{A}_0$ has a report proof in $\mathbb{P}$ (Theorem 7.1). The following example gives an object $\mathbb{P}$ that does not do this, but satisfies all the requirements for a report plan except avoiding cycles.

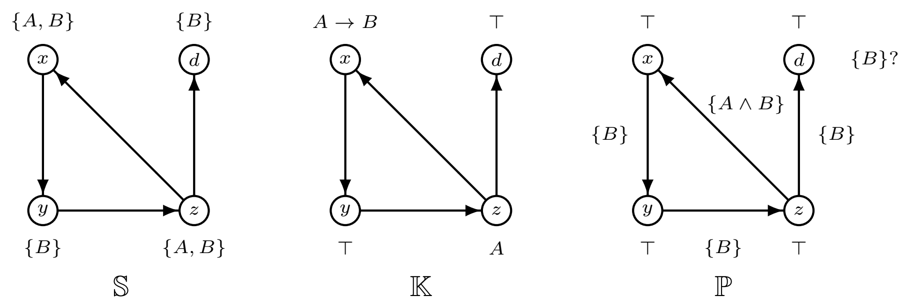

FIGURE 8. Example 6.5

Example 6.5. In Figure 8, the agent d is the unique decider in the signature network $\mathbb{S}$. Using an argument like the proof of Example 4.4, one can see that $\mathbb{S}$ is two-pass report complete but not single-pass report complete at d. Now let $\mathbb{O}$ be the Boolean closed observation network over $\mathbb{K}$ where $\mathcal{O}(x) = \{\top, \bot\}$ for each agent x. Let $\mathbb{P} = (\mathcal{O}_0(\cdot), \mathcal{C}_0(\cdot), \mathcal{A}_0)$ be the object shown at the right in Figure 8. $\mathbb{P}$ assigns the observation $\mathcal{O}_0(x) = \{\top\}$ to every agent x, has the single alternative $\mathcal{A}_0 = \{B\}$, and has $\mathcal{C}_0(z, x) = \{A \wedge B\}$ and $\mathcal{C}_0(\cdot) = \{B\}$ elsewhere. The only observation in $\mathbb{P}$ is the function $O(x) = \top$, and the combined knowledge base for $\mathbb{K}^O$ is $\mathcal{K}^O(V) = \{A\}$. $\mathbb{P}$ satisfies all the requirements for being a report plan for $(\mathbb{O}, \mathcal{A}_0, d)$ except for avoiding cycles. The sentence B is two-pass report provable at d in $\mathbb{K}^O$, but B does not have a report proof in $\mathbb{P}$, because none of the agents in the triangle can report anything before it receives a report from another agent.

The remark below says that being a report plan is preserved under adding new sentences to the knowledge bases, new potential observations, or new alternatives.

Remark 6.6. *Suppose $\mathbb{P}$ is a report plan for $(\mathbb{O}, \mathcal{A}, d)$. Let*

$$\mathbb{O}' = (V, E, L(\cdot), \mathcal{K}'(\cdot), \mathcal{O}'(\cdot))$$

be an observation network over the same $\mathbb{S}$ such that $\mathcal{K}(x) \subseteq \mathcal{K}'(x)$ and $\mathcal{O}(x) \subseteq \mathcal{O}'(x)$ for each agent $x \in V$. Also suppose that $\mathcal{A} \subseteq \mathcal{A}' \subseteq [L(d)]$. Then $\mathbb{P}$ is still a report plan for $(\mathbb{O}', \mathcal{A}', d)$.

7. Main Results

In this section we prove two theorems about report plans. Theorem 7.1 will show that a report plan guarantees that for every observation, some alternative is single-pass report provable from the observation and the knowledge bases. The proof of Theorem 7.1 will describe how a report plan $\mathbb{P}$ can be executed after each each agent x makes an observation. Theorem 7.3 will show that if the signature network $\mathbb{S}$ contains a signature tree with decider d, then every Boolean closed observation network over $\mathbb{S}$ that is sufficient for the set of alternatives has a report plan at d.

Theorem 7.1. *Let $\mathbb{O}$ be an observation network with decider d, and let $\mathbb{P} = (\mathcal{O}_0(\cdot), \mathcal{C}_0(\cdot), \mathcal{A}_0)$ be a report plan for $(\mathbb{O}, \mathcal{A}, d)$. Then for every observation $O(\cdot)$ in $\mathbb{P}$ there exists an alternative $A \in \mathcal{A}_0$ and a report proof in $\mathbb{P}$ of A at d from $O(\cdot)$.*

Proof. Let F be the set of all edges $(x, y) \in E$ such that $\mathcal{C}_0(x, y) \neq \{\top\}$. Let U be the set consisting of d and all agents $x \in V$ such that (V, F) contains a path from x to d. Let $G = F \cap (U \times U)$. Note that if $y \in U$ and $(x, y) \in F$ then $x \in U$ and $(x, y) \in G$. Define the **height** $h(x)$ of an agent $x \in U$ to be the length of the longest path in (U, G) that ends in x, with the convention that $h(x) = 0$ if G contains no edge of the form (z, x). Since $\mathbb{R}$ avoids cycles, (U, G) contains no directed cycles. Hence each path in (U, G) has length at most $|U|$. Therefore each agent in U has finite height. (U, G) contains a path from every agent $x \in U$ to d, so d has the greatest height. The set $G \cup \{(d, d)\}$ can be arranged in a list $(u_1, v_1)(u_2, v_2), \ldots, (u_n, v_n)$ such that:

- $u_n = v_n = d$;
- if $i < j \leq n$ then $h(u_i) \leq h(u_j)$;
- for each $x \in U$, the i's such that $1 \leq i \leq n$ and $u_i = x$ are consecutive.

Since $\mathbb{P}$ is a report plan, we can choose sentences $C_i \in \mathcal{C}_0(u_i, v_i)$ such that for each $i < n$, if $\mathcal{R}_i = \{C_j : j < i, v_j = u_i\}$ is a potential report to u_i then $\mathcal{K}(u_i) \cup \{O(u_i)\} \cup \mathcal{R}_i \vdash C_i$. It then follows by induction that for each $i < n$, $\mathcal{R}_i$ is a potential report to u_i. Moreover, $\mathcal{R}_n = \{C_j : j < n, v_j = d\}$ is a potential report to d, and there is a sentence $A \in \mathcal{A}_0$ such that $\mathcal{K}(d) \cup \{O(d)\} \cup \mathcal{R}_n \vdash A$. This shows that

$$\mathbb{R} = (u_1, v_1, C_1)(u_2, v_2, C_2) \cdots (u_{n-1} v_{n-1} C_{n-1})(d, d, A)$$

is a single-pass report proof in $\mathbb{K}^O$ of A at d from $O(\cdot)$. Since $C_i \in \mathcal{C}_0(u_i, v_i)$ for each i and $A \in \mathcal{A}_0$, $\mathbb{R}$ is a report proof in $\mathbb{P}$. □

Definition 7.2. *We say that an observation network $\mathbb{O}$ is **plan complete** at a decider d if for every set of sentences $\mathcal{A} \subseteq [L(d)]$ for which $\mathbb{O}$ is sufficient, there exists a report plan for $(\mathbb{O}, \mathcal{A}, d)$. A signature network $\mathbb{S}$ is **plan complete** at d if every Boolean closed observation network $\mathbb{O}$ over $\mathbb{S}$ is plan complete at d.*

Theorem 7.3. *If a signature network $\mathbb{S}$ contains a signature tree $\mathbb{T} = (V, F, L(\cdot))$ with decider d, then $\mathbb{S}$ is plan complete at decider d.*

The proof will show a bit more: For each Boolean closed observation network $\mathbb{O}$ over $\mathbb{S}$ and $\mathcal{A} \subseteq [L(d)]$ there is report plan $\mathbb{P}$ for $(\mathbb{O}, \mathcal{A}, d)$ such that $\mathcal{C}_0(x, y) = \{\top\}$ whenever $(x, y) \in E \setminus F$. Intuitively, this means that each agent reports only to its unique parent in the tree (V, F).

Proof of Theorem 7.3. Let $\mathcal{A} \subseteq [L(d)]$. Note that if we are able to get a report plan for $(\mathbb{O}', \mathcal{A}, d)$ where $\mathbb{O}'$ is the observation network

$$\mathbb{O}' = (V, F, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}(\cdot))$$

formed by replacing E by F and leaving everything else unchanged, then we at once get a report plan for $(\mathbb{O}, \mathcal{A}, d)$ by putting $\mathcal{C}_0(x, y) = \{\top\}$ for each edge $(x, y) \in E \setminus F$. We may therefore assume without loss of generality that $\mathbb{S}$ is already a signature tree, so that $E = F$.

Since (V, E) is a tree, each agent $x \neq d$ has a unique parent $y = p(x)$, and d has no parents. Moreover, (V, E) has no directed cycles. By Fact 3.4, $\mathbb{S}$ is report complete at d. Then by Corollary 6.1, there are finitely many observations $O_1(\cdot), \ldots, O_n(\cdot)$ in $\mathbb{O}$ and finitely many alternatives $A_1, \ldots, A_n$ in $\mathcal{A}$ such that conditions 6.1 (a) and (b) hold. Let $\mathcal{A}_0 = \{A_1, \ldots, A_n\}$.

For each agent $x \in V$, let $\mathcal{O}_0(x)$ be the set of all logically consistent sentences of the form $P_1(x) \wedge \cdots \wedge P_n(x)$ where for each $k \leq n$, $P_k(x) \in \{O_k(x), \neg O_k(x)\}$. It is clear that each $\mathcal{O}_0(x)$ is a finite partition. Since $\mathbb{O}$ is Boolean closed, we have $\mathcal{O}_0(x) \subseteq \mathcal{O}(x)$ for each x.

We are going to build sets of question sentences $\mathcal{C}_0(x, y)$ for each $(x, y) \in E$ such that $\mathbb{P} = (\mathcal{O}_0(\cdot), \mathcal{C}_0(\cdot), \mathcal{A}_0)$ is a report plan for $(\mathbb{O}, \mathcal{A}, d)$. For now, we let $\mathbb{P}^-$ be the pair $(\mathcal{O}_0(\cdot), \mathcal{A}_0)$, and say that $O(\cdot)$ is an observation in $\mathbb{P}^-$ if $O(x) \in \mathcal{O}_0(x)$ for each $x \in V$. Consider an observation $O(\cdot)$ in $\mathbb{P}^-$. Recall that we write $O(V)$ for $\bigwedge_{x \in V} O(x)$. Since each set $\mathcal{O}_0(x)$ is a partition, for each $k \leq n$, $O(V)$ is either equal to $O_k(V)$ or is inconsistent with $O_k(V)$. By 6.1 (a), we have

$$\mathcal{K}(V) \vdash (O_1(V) \vee \cdots \vee O_n(V)).$$

So if $O(V)$ is not one of $O_1(V), \ldots, O_n(V)$, then $O(V)$ is not consistent with $\mathcal{K}(V)$, and hence every sentence is provable from $\mathcal{K}(V)$ and $O(V)$. So in every case there exists $k \leq n$ such that

$$\mathcal{K}(V) \vdash O(V) \to O_k(V),$$

and hence

$$\mathcal{K}^O(V) \vdash O_k(V).$$

By 6.1 (b), A_k is report provable in $\mathbb{K}^{O_k}$ at d, so A_k is provable from $\mathcal{K}^O(V)$. We let $A^O = A_k$ and note that $A^O \in \mathcal{A}_0$.

By Proposition 4.9, for each observation $O(\cdot)$ in $\mathbb{P}^-$, $\mathbb{K}^O$ is single-pass report complete at d, so we may choose a single-pass report proof of A^O at d in $\mathbb{K}^O$. For each edge $(x, y) \in E$, let $B^O(x, y)$ be the unique sentence reported along (x, y) in this proof if there is one, and otherwise let $B^O(x, y) = \top$. Then for each $(x, y) \in E$ we have:

- $B^O(x, y) \in [L(x) \cap L(y)]$;
- $\mathcal{K}(x) \cup \{O(x)\} \cup \{B^O(z, x) : (z, x) \in E\} \vdash B^O(x, y)$;
- $\mathcal{K}(d) \cup \{O(d)\} \cup \{B^O(z, d) : (z, d) \in E\} \vdash A^O$.

For each edge $(x, y) \in E$, let

$$\mathcal{D}_0(x, y) = \{B^O(x, y) : O(\cdot) \text{ is an observation in } \mathbb{P}^-\}.$$

We write $x \leq y$ if there is a path from x to y in (V, E), and write $x < y$ if $x \leq y$ and $x \neq y$. As before, the **height** of y is the length of the longest path ending in y. Note that if $x < y$ then the height of x is less than the height of y.

Now, for each observation $O(\cdot)$ in $\mathbb{P}^-$, define the sentences $D^O(x,y)$, $(x,y) \in E$, inductively on the height of x in the tree (V,E) as follows: $D^O(x,y)$ is the conjunction of all sentences in $\mathcal{D}_0(x,y)$ that are provable from

$$\mathcal{K}(x) \cup \{O(x)\} \cup \{D^O(z,x) : (z,x) \in E\}.$$

Finally, for each edge $(x,y) \in E$ we define

$$\mathcal{C}_0(x,y) = \{D^O(x,y) : O(\cdot) \text{ is an observation in } \mathbb{P}^-\}.$$

Since there are finitely many observations in $\mathbb{P}^-$, $\mathcal{D}_0(x,y)$ and $\mathcal{C}_0(x,y)$ are finite subsets of $[L(x) \cap L(y)]$ for each edge $(x,y) \in E$. We let

$$\mathbb{P} = (\mathcal{O}_0(\cdot), \mathcal{C}_0(\cdot), \mathcal{A}_0),$$

and will show that $\mathbb{P}$ is a report plan for $(\mathbb{O}, \mathcal{A}, d)$. Note that the observations in $\mathbb{P}$ are just the observations in $\mathbb{P}^-$.

Consider an arbitrary agent $x \in V$ and let $B \in \mathcal{O}_0(x)$. An argument by induction on the height of z in (V,E) shows that for each edge $(z,y) \in E$, the sentences $D^O(z,y)$ depend only on the observations $O(u)$ such that $u \le z$. Let $\{z_1, \ldots, z_m\}$ be the set of all children of x in (V,E), and let $\mathcal{R}(x)$ be a potential report to x in $\mathbb{P}$. Then $\mathcal{R}(x)$ has the form

$$\mathcal{R}(x) = \{D_1, \ldots, D_m\}$$

where each D_j is the conjunction of one or more sentences in $\mathcal{C}_0(z_j, x)$. Then for each $j \le m$ there is an observation $O_j(\cdot)$ in $\mathbb{P}$ such that the sentence $D^{O_j}(z_j, x)$ occurs in the conjunction D_j.

Since (V,E) is a tree, for any distinct $i,j \le m$, the sets $\{u : u \le z_i\}$ and $\{u : u \le z_j\}$ are disjoint. Therefore there is an observation $O(\cdot)$ in $\mathbb{P}$ such that $O(u) = O_j(u)$ for each $j \le m$ and agent $u \le z_j$, and $O(x) = B$. Then for each $j \le m$ we have $D^O(z_j, x) = D^{O_j}(z_j, x)$, so $\mathcal{R}(x) \vdash D^O(z,x)$ whenever $(z,x) \in E$.

Whenever $(x,y) \in E$, $D^O(x,y)$ is a conjunction of sentences that are provable from $\mathcal{K}(x) \cup \{O(x)\} \cup \mathcal{R}(x)$, so

$$\mathcal{K}(x) \cup \{O(x)\} \cup \mathcal{R}(x) \vdash D^O(x,y).$$

By the definition of $\mathcal{C}_0(x,y)$ we have $D^O(x,y) \in \mathcal{C}_0(x,y)$.

In the case $x = d$, we have $\mathcal{R}(d) \vdash D^O(z,d)$ whenever $(z,d) \in E$. By induction on height, for each observation $O(\cdot)$ in $\mathbb{P}$ and edge $(z,y) \in E$, we have $D^O(z,y) \vdash B^O(z,y)$. Then $\mathcal{R}(d) \vdash B^O(z,d)$ whenever $(z,d) \in E$. We recall that

$$\mathcal{K}(d) \cup \{O(d)\} \cup \{B^O(z,d) : (z,d) \in E\} \vdash A^O.$$

Therefore

$$\mathcal{K}(d) \cup \{O(d)\} \cup \mathcal{R}(d) \vdash A^O.$$

This completes the proof that $\mathbb{P}$ is a report plan for $(\mathbb{O}, \mathcal{A}, d)$, and hence that $\mathbb{O}$ is plan complete at d. □

8. Special Cases

8.1. **Minimal Observation Networks.** What do the theorems in Section 7 tell us if we just have a knowledge base network instead of an observation network? In other words, what do they tell us when the agents do not learn (or only observe the true sentence $\top$)? To answer this, we note that every knowledge base network

$$\mathbb{K} = (V, E, L(\cdot), \mathcal{K}(\cdot))$$

can be made into a finite Boolean closed observation network

$$\mathbb{O} = (V, E, L(\cdot), \mathcal{K}(\cdot), \mathcal{O}(\cdot))$$

by putting $\mathcal{O}(x) = \{\top, \bot\}$ for each agent x. We call $\mathbb{O}$ the **minimal observation network** over $\mathbb{K}$. In the minimal observation network over $\mathbb{K}$, in any scenario every agent will observe the true sentence $\top$, because $\mathcal{O}(V) = \{\top, \bot\}$, and for every complete extension $\mathcal{M}$ of $\mathcal{K}(V)$, we have $\mathcal{M} \cap \mathcal{O}(V) = \{\top\}$. Therefore, the minimal observation network $\mathbb{O}$ over $\mathbb{K}$ is sufficient for a set of alternatives $\mathcal{A}$ if and only if $\mathcal{K}(V) \vdash A$ for some $A \in \mathcal{A}$.

Proposition 8.1. *Suppose that $\mathbb{O}$ is a minimal observation network, and $\mathbb{P} = (\mathcal{O}_0(\cdot), \mathcal{C}_0(\cdot), \mathcal{A}_0)$ is a report plan for $(\mathbb{O}, \mathcal{A}, d)$. Then $\mathcal{O}_0(x) = \{\top\}$ for each agent $x \in V$, and there exists an alternative $A \in \mathcal{A}_0$ and elements $C(x, y) \in \mathcal{C}_0(x, y)$, $(x, y) \in E$, such that the triple*

$$\mathbb{P}' = (\mathcal{O}_0(\cdot), \mathcal{C}'_0(\cdot), \mathcal{A}'_0), \quad \mathcal{C}'_0(x, y) = \{C(x, y)\}, \quad \mathcal{A}'_0 = \{A\}$$

is a report plan for $(\mathbb{O}, \mathcal{A}, d)$.

Proof. We first note that the only partition contained in $\mathcal{O}(V)$ is $\{\top\}$, so $\mathcal{O}_0(x) = \{\top\}$ for each $x \in V$. Then the only observation in $\mathbb{P}$ is the function $O(x) = \top$. By Theorem 7.1, there is a single-pass report proof $\mathbb{R}$ in $\mathbb{P}$ of a sentence $A \in \mathcal{A}_0$ at d from $O(\cdot)$. Each edge (x, y) occurs at most once in $\mathbb{R}$. If $\mathbb{R}$ contains the triple (x, y, B), put $C(x, y) = B$, and otherwise put $C(x, y) = \top$. Then $C(x, y) \in \mathcal{C}_0(x, y)$ for each edge, and $\mathbb{P}'$ satisfies all the requirements for a report plan for $(\mathbb{O}, \mathcal{A}, d)$. □

We now use a minimal observation network to show that plan completeness implies single-pass report completeness for signature networks.

Proposition 8.2. *If a signature network $\mathbb{S}$ is plan complete at d, then $\mathbb{S}$ is single-pass report complete at d.*

Proof. Suppose $\mathbb{S}$ is plan complete at d, and let $\mathbb{K}$ be a knowledge base network $\mathbb{K}$ over $\mathbb{S}$. Assume that $D \in [L(d)]$ and D is provable from $\mathcal{K}(V)$. Let $\mathbb{O}$ be the minimal observation network over $\mathbb{K}$. Then $\mathbb{O}$ is plan complete at d. Since $\mathcal{K}(V) \vdash D$, D holds in every complete extension $\mathcal{M}$ of $\mathcal{K}(V)$, and hence $\mathbb{O}$ is sufficient for D at d. By the plan completeness of $\mathbb{S}$ at d, there exists a report plan for $(\mathbb{O}, \{D\}, d)$. Then by Theorem 7.1, D is single-pass report provable in $\mathbb{K}$ at d, so $\mathbb{S}$ is single-pass report complete at d. □

8.2. **Fragments of First Order Logic.** As pointed out in the Prerequisites, all our results hold for propositional logic as well as first order logic. In particular, Theorem 7.3 holds for propositional logic.

Corollary 8.3. *Assume the hypotheses of Theorem 7.3. If each sentence in the combined knowledge base $\mathcal{K}(V)$, the combined set of potential observations $\mathcal{O}(V)$, and the set of alternatives $\mathcal{A}$ is a sentence in propositional logic, then there is a report plan $\mathbb{P}$ composed entirely of sentences in propositional logic.*

Similarly, all the results of [7] and hence Theorem 7.3 hold for the fragment of first order logic without quantifiers. A first order sentence is said to be **quantifier-free** if it has no quantifiers.

Corollary 8.4. *Assume the hypotheses of Theorem 7.3. If each sentence in the combined knowledge base $\mathcal{K}(V)$, the combined set of potential observations $\mathcal{O}(V)$, and the set of alternatives $\mathcal{A}$ is quantifier-free, then there is a report plan $\mathbb{P}$ composed entirely of quantifier-free sentences.*

8.3. **Approximate Values.** In many situations, there is a need to determine an approximate value of one or more unknown quantities. We briefly indicate how report plans might be used to approximate one unknown quantity. The idea can easily be generalized to the case of finitely many unknown quantities.

Suppose $\mathbb{O}$ is an observation network in which the signature $L(d)$ of the decider has at least the symbols $+, -, \leq$, a constant symbol for each rational number, and one extra constant symbol c for an "unknown quantity". Suppose the knowledge base $\mathcal{K}(d)$ has at least the axioms for an ordered abelian group and all true equations and inequalities involving rational numbers. We allow the possibility that $L(d)$ also has other symbols and $\mathcal{K}(d)$ has additional sentences, and make no restrictions about the other agents. We let $\mathbb{Q}$ denote the set of rational numbers. For each positive rational number r, let $\mathcal{A}(r)$ be the set of sentences

$$\{q \leq c \wedge c \leq q + r \,:\, q \in \mathbb{Q}\}.$$

Each sentence in $\mathcal{A}(r)$ says that c belongs to a closed interval of length r with rational endpoints.

Theorem 7.3 tells us that if $\mathbb{O}$ is Boolean closed, contains a signature tree with decider d, and is sufficient for $\mathcal{A}(r)$, then there exists a report plan $\mathbb{P}$ for $(\mathbb{O}, \mathcal{A}(r), d)$.

If $\mathbb{P} = (\mathcal{O}_0(\cdot), \mathcal{C}_0(\cdot), \mathcal{A}_0)$ is a report plan for $(\mathbb{O}, \mathcal{A}(r), d)$, and $O(x) \in \mathcal{O}_0(x)$ for each agent x, then Theorem 7.1 tells us that some sentence $A \in \mathcal{A}_0$ is report provable in the knowledge base network $\mathbb{K}^O$. This sentence A belongs to $\mathcal{A}(r)$, so it approximates the unknown value c within r.

Now let $r_0, r_1, \ldots$ be a sequence of positive rational numbers that converges to 0, and consider a sequence of report plans $\mathbb{P}^n$ for $(\mathbb{O}, \mathcal{A}(r_n), d)$. In each scenario, the sequence of report plans will produce a sequence of better and better approximations of the unknown quantity c. The following corollary to Theorem 7.1 shows that in every scenario, the sequence of report plans produces a unique real value for the unknown constant c.

Proposition 8.5. *Suppose that $r_0, r_1, \ldots$ is a sequence of positive rational numbers that converges to 0, and that for each n, $\mathbb{P}^n = (\mathcal{O}_0^n(\cdot), \mathcal{C}_0^n(\cdot), \mathcal{A}_0^n)$ is a report plan for $(\mathbb{O}, \mathcal{A}(r_n), d)$. Let $\mathcal{M}$ be a complete extension of $\mathcal{K}(V)$ and let $O^n(\cdot) = O^{\mathbb{P}^n,\mathcal{M}}(\cdot)$ be the unique observation given in Remark 6.3. Then the real number*

$$s = \sup\{q \in \mathbb{Q} : \mathcal{M} \vdash q \leq c\}$$

exists and for each n, s belongs to an interval $[q_n, q_n + r_n]$ where the sentence

$$q_n \leq c \wedge c \leq q_n + r_n$$

belongs to $\mathcal{A}_0^n$ and is report provable in $\mathbb{K}^{O^n}$.

Proof. By Theorem 7.1, for each n there is a rational q_n such that the sentence $A_n = (q_n \leq c \wedge c \leq q_n + r_n)$ belongs to $\mathcal{A}_0^n$ and is report provable in $\mathbb{K}^{O^n}$. By Fact 3.2, A_n is provable from the combined knowledge base $\mathbb{K}^{O^n}(V)$. But $\mathcal{K}(V) \subseteq \mathcal{M}$ and $\mathcal{M} \vdash O^n(x)$ for each n and x. Therefore $\mathcal{M} \vdash A_n$ for each n.

Since $\mathcal{M}$ is consistent and contains the axioms for ordered abelian groups with constants for each rational, the set

$$\{q \in \mathbb{Q} : \mathcal{M} \vdash q \leq c\}$$

contains q_1 and is bounded above by $q_1 + r_1 + 1$. Therefore the supremum of this set exists and is a real number s. Moreover, s belongs to the interval $[q_n, q_n + r_n]$ for each n, as required. □

9. Some Open Questions

In this section we pose some questions that were left open in this paper. Example 4.7 gives a report complete knowledge base network $\mathbb{K}$ that is not n-pass report complete for any n. The analogous question for signature networks is open.

Question 9.1. *Suppose $\mathbb{S}$ is a report complete signature network. Is there a finite n such that $\mathbb{S}$ is n-pass report complete?*

Our next open question is a possible converse to Theorem 7.3.

Question 9.2. *Let $\mathbb{S}$ be a signature network that is plan complete at a decider d. Must $\mathbb{S}$ contain a signature tree with decider d?*

The following corollary gives two cases where the above question has an affirmative answer.

Corollary 9.3. *Let $\mathbb{S}$ be a signature network that is either over a directed acyclic graph or a connected symmetric graph. If $\mathbb{S}$ is plan complete at d then $\mathbb{S}$ contains a signature tree with decider d.*

Proof. By Proposition 8.2 and Facts 3.6 and 3.8. □

Here is another open question.

Question 9.4. *Let $\mathbb{S}$ be a signature network and d a decider. Suppose that every minimal observation network $\mathbb{O}$ over $\mathbb{S}$ is plan complete. Must every Boolean closed observation network over $\mathbb{S}$ also be plan complete?*

Remark 9.5. *The following are equivalent:*

(a) *The answers to Questions 9.2 and 9.4 are both "yes";*
(b) *If every minimal observation network over signature network $\mathbb{S}$ is plan complete at d, then $\mathbb{S}$ contains a signature tree with decider d.*

Our last open question is a possible converse to Proposition 8.2.

Question 9.6. *If a signature network is single-pass report complete at a decider d, must it be plan complete at d?*

10. Conclusion

This paper and the paper [7] provide a "report theory" for analyzing situations of the following kind: Agents make inferences based on the information they have, and report them other agents in order to make decisions. Such reporting is natural when decentralized information in a network is to be incorporated by an agent who is faced with a decision. We provide conditions under which reporting can lead to correct decisions. These conditions involve signatures, knowledge bases, and potential observations of the agents, the way the agents are connected in a network, and the alternatives to be decided. Future work might involve evaluating and designing knowledge bases and networks tailored to specific applications.

References

[1] Amir, E. and S. McIlraith. Partition-Based Logical Reasoning for First-Order and Propositional Theories. Artificial Intelligence Journal, 162 49-88 (2005).
[2] Chang, C. C. and H. J. Keisler. Model Theory, Dover Publications (2012).
[3] Cowell, R, A. P. Dawid, S. Lauritzen, and D. Spiegelhalter. Probabilistic Networks and Expert Systems, Springer (2007).
[4] Craig, W. Linear Reasoning. Journal of Symbolic Logic 22 250-268. Three Uses of the Herbrand-Gentzen Theorem in Relating Model Theory and Proof Theory. Ibid. 269-285 (1957).
[5] Dechter, R. Constraint Processing. Morgan Kauffman (2003).
[6] Enderton, H. A Mathematical Introduction to Logic. Harcourt Academic Press (2001).
[7] Keisler, H. J. and J. M. Keisler. Craig Interpolation for Networks of Sentences. Annals of Pure and Applied Logic 163, 1322–1344 (2012).

Department of Mathematics, University of Wisconsin-Madison, Madison, WI 53706
E-mail address: keisler@math.wisc.edu
URL: www.math.wisc.edu/∼keisler

Department of Management Science and Information Systems, University of Massachusetts-Boston, Boston MA 02125
E-mail address: jeff.keisler@umb.edu
URL: www.umb.edu/academics/cm/faculty_staff/faculty/jeffrey_m._keisler