Reasoning with Protocols Under Imperfect Information: Extended Abstract

Eric Pacuit

Tilburg Institute for Logic and Philosophy of Science
Department of Philosophy
Tilbrug University
e.j.pacuit@uvt.nl

Sunil Easaw Simon

Institute for Logic, Language and Information
University of Amsterdam
s.e.simon@uva.nl

Intuitively, a *protocol*, or *plan*, specifies actions (observations, communications) that are available (or permitted) at any given moment. For example, in a conversation, it is typically not polite to "blurt everything out at the beginning", as we must speak in small chunks. Other natural conversational rules include "do not repeat yourself", "let others speak in turn", and "be honest". Imposing such rules *restricts* the legitimate sequences of possible statements. In general, a protocol identifies a subtree from the "grand stage" of all possible sequences of events that could take place in an interactive situation. A number of authors have forcefully argued that the underlying protocol is an important component of any analysis of a (social) interactive situation and should be explicitly represented in a formal model. Indeed, much of the work over the past 20 years using epistemic logic to reason about distributed algorithms has provided interesting case studies highlighting the interplay between protocol analysis and epistemic reasoning.

We represent a decision situation (ignoring preferences and other motivational attitudes) as a labeled transition system which we call an **arena**:

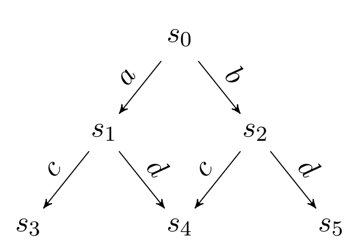

A **protocol** is a tree with labels from the (finite) set of possible actions. We are interested in what properties the agent can *guarantee* will be true by *adopting a given protocol*. The idea is that adopting a protocol at a given state restricts the paths that the agent will follow from that state. For example, suppose that the agent is in state s_0 and consider the protocol "either choose c or choose d". The protocol only gives partial information about what actions to follow at a given state (eg., neither c nor d is an option at s_0, so the protocol does not offer any advice about what to do at s_0). This protocol can be described by the PDL expression $(a \cup b); c \cup (a \cup b); d$. Note that every path in the above arena is consistent with this protocol, so we can say that this protocol is *enabled* at s_0. However, as pointed out by Johan van Benthem is his new book on dynamic logic, this way of thinking about the protocol misses a crucial point: the agent must commit to do either c or d *independent* of which action is chosen at state s_0. In other words, the agent must choose (at s_0) between the following two restrictions on future choices:

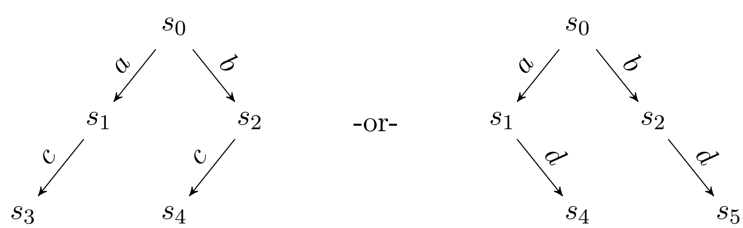

This distinction is not important if we are only interested in the states that can result by following this protocol: in this case, $\{s_3, s_4\} \cup \{s_4, s_5\}$. However, it is crucial when defining what it means for a protocol to be *enabled* at a state. This is particularly important in situations of *imperfect information*. For example, consider the following situation where the agent cannot distinguish between nodes s_1 and s_2 and the protocol pictured to the right (do a followed by c or do b followed by d):

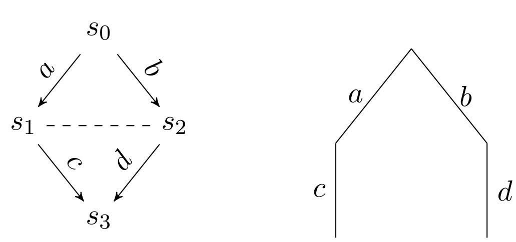

This protocol is clearly enabled in the situation without the uncertainty relation between s_1 and s_2. However, in the above situation at s_0, the agent cannot agree to "*knowingly*" follow the protocol since the agent is uncertain about the actions that are available at states[1] s_1 and s_2.

1 Key Definitions

We assume the reader is familiar with standard definitions of trees and arenas (i.e., labeled transition systems or Kripke models). A **protocol** is a finite labelled trees. Let Σ be a finite set whose elements are called **actions**. A Σ-labelled (finite) tree T

[1] Alternatively, we can say that the agent forgets at state s_1 (and s_2) the choice made at state s_0.

is a tuple $(S, \{\Rightarrow_a\}_{a\in\Sigma}, s_0)$ where S is a (finite) set of nodes, $s_0 \in S$ is the root and for each $a \in \Sigma$, $\Rightarrow_a \subseteq S \times S$ is the edge relation satisfying the usual properties. For a node $s \in S$, let $\mathcal{A}(s) = \{a \in \Sigma \mid \exists s' \in S \text{ where } s \Rightarrow_a s'\}$ denote the set of *actions available at s*.

We formally model an interactive (or decision-theoretic) situation in a standard way as a labelled transition system which we call **arenas**: Let W be a nonempty finite set, whose elements are called **positions** or **states**, and Σ a finite set of basic actions. An **arena** is a structure $\mathcal{G} = (W, \{\rightarrow_a\}_{a\in\Sigma})$ where for each $a \in \Sigma$, $\rightarrow_a \subseteq W \times W$. Following standard notation, we write $w \rightarrow_a v$ if $(w, v) \in \rightarrow_a$. The above notation for available actions and paths are readily applied to finite arenas.

A protocol or plan *restricts* the available choices for the agent(s). Intuitively, if an agent agrees to follow a finite protocol, then she agrees to restrict her choices to all and only those actions compatible with the protocol. Of course, not all protocols *can* be followed in any situation. This lead us to the key notion of a protocol being **enabled** at a state u in an arena. If there is no uncertainty in the arena, then the formal definition of a protocol being enabled is completely straightforward: a protocol T is enabled at u in $\mathcal{G}$ if T can be embedded in the unwinding of u.

Intuitively, if a protocol T is enabled at a state u in an arena $\mathcal{G}$, then it is (physically, objectively) *possible* for the agent to *agree* to follow T. Of course, this does not necessarily mean that the agent *knows* (or *believes*) she can follow T, *wants* to follow T or it is in the agent's interests for follow T. Our main goal in this paper is to explore a different sense in which a protocol is "possible" taking into account the agent's *point-of-view*. Our first task is to extend the definition of an arena with an explicit representation of the agent's "point-of-view" at each position in the arena. As is standard in the epistemic logic literature, we use a relation on the set of states in an arena to represent this uncertainty of the agent about her position in the arena.

Definition 1.1 An **arena with imperfect information** is a structure $\mathcal{G}^I = (W, \{\rightarrow_a\}_{a\in\Sigma}, \rightsquigarrow)$ where $(W, \{\rightarrow_a\}_{a\in\Sigma})$ is a finite arena and $\rightsquigarrow \subseteq W \times W$.

For each position u in an arena, let $\mathcal{I}(u) = \{w \mid u \rightsquigarrow w\}$ be the agent's "*point-of-view*". The above models do not impose any structural properties on the action and $\rightsquigarrow$ relations. However, a number of properties have been discussed in the literature are natural in many situations which we list below without discussion:

- **No Miracles**: for all $a \in \Sigma$ and all $w, v, w', v' \in W$, if $w \rightsquigarrow v$, $w \rightarrow_a w'$, and $v \rightarrow_a v'$, then $w' \rightsquigarrow v'$.
- **Success**: If $w \rightsquigarrow v$ then $\mathcal{A}(v) \subseteq \mathcal{A}(w)$
- **Awareness**: If $w \rightsquigarrow v$ then $\mathcal{A}(w) \subseteq \mathcal{A}(v)$
- **Certainty of available actions**: If $w \rightsquigarrow v$ and $w \rightsquigarrow v'$ then $\mathcal{A}(v) = \mathcal{A}(v')$

A protocol being enabled is an *objective* notion from the point-of-view of the modeler that does not take into account the agent's point-of-view. What we need is a *subjective* version of being enabled. One idea is to mimic the restriction operation, but ensure *at each step* to take into account all and only the positions that the agent has access to via the $\rightsquigarrow$ relation. Intuitively, a protocol T is *subjectively enabled* at a position u in an arena with imperfect information if

(i) the agent is *certain that* T is enabled (for all $v \in \mathcal{I}(u)$, T is enabled at v), and

(ii) the agent will be certain that she is in fact following the protocol at *every stage* of the protocol.

This second point is important as there is a difference between "knowing that a protocol is enabled" and "being able to *knowingly follow* a protocol". This difference is crucial when analyzing long-term plans. Thus, our definition must take into account the **forest** $\{T_v \mid v \in \mathcal{I}(u)\}$ *for every position u not ruled out by the protocol* (T_v is the unwinding at v in the arena). The formal definition of a protocol being **subjectively enabled** is provided in the full paper. For now, we discuss simple example which illustrates the main idea. Notice that without additional structural assumptions on $\rightsquigarrow$ a protocol being subjectively enabled does *not* imply that the protocol is enabled. For example, consider the arena below and the protocol discussed in the introduction: "either do a followed by c or do b followed by d". This protocol is subjectively enabled but not enabled at state s_0.

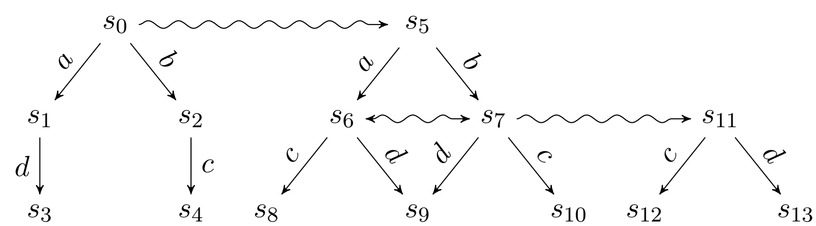

(*Note that the protocol is still subjectively enabled if we impose the no miracle property which would add a number $\rightsquigarrow$ edges.*)

However, in situations of *perfect information*, we can prove that subjectively enabled is equivalent to enabled.

2 Epistemic Protocol Logic

An arena with imperfect information describes what *can happen* in an interactive situation both objectively (from the modeler's point-of-view) and subjectively (from the agent's point-of-view via the $\rightsquigarrow$ relations). That is, they describe both what is physically possible for the agent to do and what she thinks she can do in an interactive situation.

Committing to a basic protocol T *restricts* the choices available to the agent, but there is a trade-off: it also *increases* the ability of the agent to *guarantee* that certain propositions are true. Formally, each basic protocol (which is a finite tree) is associated with a set of states X (the *frontier* of T in an arena). These are the states that the agent can "force" the situation to end up in by making choices consistent with the protocol. There are a number of ways to make precise what it means for an agent to "guarantee" that some proposition is true because she adopts the protocols T. One options is to see what is true no matter what the agent does as long as it is consistent T. A second option recognizes that T still represents choices for the agent which will be settled in the course of the interaction. In this case, we are interested in what the agent can force by doing something consistent with

T. The situation is even more interesting when the agent committs to a complex protocol. If the protocol involves the operators $\cup$ or Kleene star then the agent first must choose which set of states she wants to have the abilitiy to force. For example, consider the protocol $T_1 \cup T_2$, in order to commit to this protocol the agent must choose which of the two basic protocols to follow. More generally, given a complex protocol π, the agent must first decide both *how* to go about adopting π then make her choices"in the moment" consistent with this plan.

This discussion suggests that our basic modality will be interpreted as a sequence of *two* quantifiers (each corresponding to the different "types" of decisions the agent makes when committing to a protocol). This is familiar from other modal logics of ability (eg., STIT) and game logics. Of the four possible combinations of quantifiers, we list take the following two as primitive (corresponding to $\exists\forall$ and $\exists\exists$ respectively):

- $\langle\pi\rangle^{\forall}\alpha$: By adopting the protocol π, α is guaranteed to be true.
- $\langle\pi\rangle^{\exists}\alpha$: By adopting the protocol π, the agent can do something consistent with the protocol that will make α true.

As usual, the remaining two possible combinations of quantifiers are dual to these. We take "adopting a protocol" to mean that the agent decides how to follow the protocol (so an existential quantifier over the different sets of states the agent can force). The second quantifier is over the different ways that the agent actually implements the protocol. These notions are objective since they do not take into account the fact that the agent may be imperfectly informed about her current position in the arena. This suggests the following "epistemized" versions of the above operators:

- $\langle\pi\rangle^{\Box}\alpha$: By *agreeing* to adopt the protocol π, the agent is certain that α is guaranteed to be true.
- $\langle\pi\rangle^{\Diamond}\alpha$: By *agreeing* to adopt the protocol π, the agent is can "knowingly" do something consistent with the protocol that will make α true.

3 Results

This paper focuses on the interplay between epistemic reasoning and protocol analysis. The full version of the paper discusses a number of key issues surrounding the problem of how to model agents "knowing a protocol, or plan".

- The main technical contribution of this paper is a sound and (weakly) complete axiom system for reasoning about agents' abilities in situations of imperfect information. We also compare with related logics (CPDL, Game Logic and CTL).
- This extended abstract focused on the single agent situation. While this situation already raises a number of conceptual and technical issues, we also discuss the many agent case defining (and axiomatizing) **joint protocols**. This many-agent extension may also shed new light on combining strategic and epistemic reasoning (cf. discussions in the AT(E)L literature).
- There is no reason to focus only on *regular operations*. Indeed, there may be other natural operations in our context, such as "merging" or "revising" protocols.