Observation algebras:
Heyting algebra over coherence spaces

Paul Brunet

Université Paris-Est Créteil
paul.brunet-zamansky.fr
paul@brunet-zamansky.fr

I - Introduction

When designing formal semantics for program verification, one needs a language (logic) to describe states: such a language forms the static side of an otherwise dynamic framework. In many cases, boolean algebras (BA), i.e. classical propositional logic, are used for this purpose. Examples include the KAT family of models [6], as well as (concurrent) Kleene algebras with boolean observations [5, 4]. However, as discussed in [7], this is not the most useful model in the presence of concurrent behaviour. Indeed, in *op. cit.* it was argued that a more general class of logics, namely pseudo-complemented distributive lattices or PCDL, were more appropriate. A tailor-made instance of this logic was presented then, and used to model weak memory properties. In this paper, we study a class of distributive lattices that may be used as alternative models of partial observations. This model generalises the PCDL model of [7] in two ways: we add an implication operator, and abstract over the atomic predicates of the language. We intend this work to be used to define purpose-built models to tackle various verification or modelling tasks in concurrent settings.

To get intuitions about our model, consider a situation where a system is being monitored by an observer. We are given a (possibly infinite) set of atomic propositions, called observations in the following, that describe static properties of the system. We also know which pairs of observations are coherent, i.e. may hold simultaneously, and which pairs are incompatible, meaning they are mutually exclusive. The observer may report a set of observations that they witnessed. This set need not be exhaustive: the fact that an atomic proposition does not appear in a report means that the observer failed to evaluate its veracity, not that they saw it to be false. To establish the falsity of an observation o, the observer has to witness some other observation, incoherent with o. This gives the model an intuitionist flavour (existence of witness), as well as an epistemic one: knowledge may be partial, and is obtained by positive observations. The experiment may be repeated, for instance observing each possible state of the system during an execution, or tying the observer to a program point and checking the system at this point for various runs with different inputs. This way we obtain a set of reports, one for each experiment, each describing partially some possible state of the system.

Formally, this situation is modeled by an undirected graph, whose nodes are the atomic propositions, and whose adjacency relation is interpreted as coherence. A single report, containing a set of pairwise coherent nodes, is simply a clique of the graph. Non-exhaustivity is represented by an ordering between cliques. We say that c_2 is more general than c_1 if c_1 contains c_2: every observation made in the report c_2 is also present in c_1, but the latter also deals with properties left unspecified by c_2. A sentence in our model is then interpreted as a set of cliques X such that if c belongs to X, then any clique less general than c also belongs to X. This way, we capture all partial states that are compatible with the report we received.

This abstract is structured as follows. First we describe the syntax and semantics of Observation Algebra, and provide some sound (but incomplete) axiomatisation. We then study classes of graphs for which a complete set of axioms is available. Finally, we discuss how this formalism can be used to model memory states before a brief overview of future work.

All the proofs in this paper have been formalized in Rocq, and are available online [2]. Additionally, an extended version of this abstract can be found on arxiv and hal [1].

II - Syntax and semantics of observation algebra

We fix a so-called coherence graph [3], i.e. a pair $G = \langle O(G), \frown_G \rangle$ consisting of a set $O(G)$ equipped with a symmetric and reflexive relation $\frown_G$. In other words, a possibly infinite undirected graph. $a, b, \ldots$ will range over $O(G)$. We will be interested in cliques of this graphs ($x, y, \ldots$ will range over cliques), whose set is called the coherence space generated by G, and written $\mathscr{C}(G) := \{x \subseteq O(G) \mid \forall a, b \in x,\ a \frown_G b\}$.

Terms $s, t \in \mathcal{T}_{\text{obs}}(G)$ are built out of the connectives $\vee$, $\wedge$ and $\to$, the constants $\top$ and $\bot$, as well as predicates a for each $a \in O(G)$. A clique x is said to satisfy a term t when $x \vDash t$, where $\vDash$ is defined inductively:

$$\begin{array}{ll} x \vDash \top & x \vDash s \text{ or } x \vDash t \Leftrightarrow x \vDash s \vee t \\ a \in x \Leftrightarrow x \vDash a & x \vDash s \text{ and } x \vDash t \Leftrightarrow x \vDash s \wedge t \\ \multicolumn{2}{l}{(\forall y \in \mathscr{C}(G),\, y \vDash s \text{ and } x \cup y \in \mathscr{C}(G) \Rightarrow x \cup y \vDash t) \Leftrightarrow x \vDash s \to t.} \end{array}$$

This yields an interpretation function $[\![-]\!] : \mathcal{T}_{\text{obs}}(G) \to \mathcal{P}(\mathscr{C}(G))$, defined as $[\![t]\!] := \{x \in \mathscr{C}(G) \mid x \vDash t\}$. Notice that the interpretation of any term is upwards-closed for containment: if $x \subseteq y$ and $x \vDash t$, then $y \vDash t$.

The question we address in the paper is finding a set of axioms that define a syntactic equivalence relation $\equiv$ on terms that is sound and complete, i.e. such that $s \equiv t \Leftrightarrow [\![s]\!] = [\![t]\!]$.

The axioms of bounded distributive lattices (BDL), which we omit here for space reasons, form a suitable starting point: they are obviously all sound with respect to our interpretation. In fact, if we augment them with axiom (A1) below, we obtain a relation $\equiv_1$ which is sound and complete when we remove $\to$ from our set of term formers (i.e. for the signature of BDL).

$$\forall a \smile b, a \wedge b \equiv \bot \tag{A1}$$

If we want to add the implication, we need to add the axioms of Heyting algebras, which we do not recall for space reasons. We supplement them with axioms that deal with the effect of putting a clique as the left argument of an implication. We write $\mathscr{C}_f(G)$ for the set of finite cliques of G, and for a finite set $x = \{x_1, \ldots, x_n\}$, we write $\bigwedge x$ for the term $x_1 \wedge \cdots \wedge x_n$.

$$\forall x \in \mathscr{C}_f, \qquad \left(\bigwedge x\right) \to (s \vee t) \equiv \left(\bigwedge x \to s\right) \vee \left(\bigwedge x \to t\right) \tag{B1}$$

$$\forall x \in \mathscr{C}_f,\ \forall a \notin x, \qquad \left(\bigwedge x\right) \to a \equiv \left(\bigwedge x \to \bot\right) \vee a. \tag{B2}$$

All these axioms are sound, but fail to be complete. In fact we haven't been able to find one generic axiomatisation for all graphs. Instead, we focus on specific graphs or classes of graphs to derive complete axiomatisations.

III - Tractable graphs

FAN graphs The first class we consider consists of all graphs that have the finite anti-neighbourhood (FAN) property: for any vertex $a \in O(G)$, the set $\{b \in O(G) \mid a \smile_G b\}$ should be finite. This class obviously includes all finite graphs. For such graphs, we can formulate the following axiom scheme:

$$\forall x \in \mathscr{C}_f, \qquad \left(\bigwedge x\right) \to \bot \equiv \bigvee \{a \in O(G) \mid \exists b \in x : a \smile_G b\} \tag{C1}$$

This yields an sound and complete axiomatisation, when we take the axioms of Heyting algebra, as well as (B1), (B2), and (C1) to define the relation $\equiv_{fan}$.

Infinite anticlique To go beyond FAN graphs we consider an extreme opposite case: an infinite graph with no edge, so that the anti-neighbourhood of any vertex is the rest of the graph. A first important observation is that cliques in this graphs are either empty or singleton, and that for any vertex a the term $a \vee (a \to \bot)$ is interpreted as the set of all singleton cliques. As such, this set does not depend on a, which motivates axiom (D1). In Heyting algebras, double negation is different from the identity. However here it does hold partially as stated in (D2).

$$\forall a, b \in O(G), \qquad a \vee (a \to \bot) \equiv b \vee (b \to \bot) \tag{D1}$$

$$\forall x \in \mathcal{P}_f\left(O(G)\right), \qquad \left(\left(\bigvee x\right) \to \bot\right) \to \bot \equiv \bigvee x \tag{D2}$$

In fact, augmenting the axioms of Heyting algebras with (B1), (B2), (D1), and (D2) yields a sound and complete relation $\equiv_\omega$. The proof relies on the fact that terms are always interpreted as either finite or cofinite sets of cliques.

Products Let $(G_i)_{i \in \mathcal{I}}$ be a collection of pairwise-disjoint graphs indexed by a (possibly infinite) set of dimensions $\mathcal{I}$. The graph $\mathcal{P} = \bigotimes_{i \in \mathcal{I}} G_i$ is defined as the union of the graphs G_i (all vertices and edges are imported), with extra edges added to connect to one another the vertices of the different graphs. In other words:

$$a \mathrel{\frown\!\!\!\!\smile}_{\mathcal{P}} b \Leftrightarrow \left(\exists i, a \mathrel{\frown\!\!\!\!\smile}_{G_i} b\right) \text{ or } \left(\exists i \neq j : a \in O(G_i) \text{ and } b \in O\left(G_j\right)\right)$$

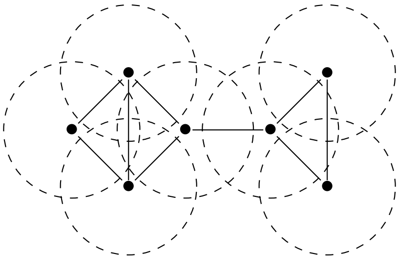

Figure 1: Example of coherence graph based on a 2-dimensional distance

The question here is the following: if we are given axiomatisations (or indeed any sound and complete syntactic relations) for each of the G_is, can we build one of the composite graph?

The first observation we make is that for any $i \in \mathcal{I}$, the terms over G_i can be seen as being terms over $\mathcal{P}$, and that this shift in perspective preserves semantic equivalence, so that the following axiom is sound:

$$\forall s, t \in \mathcal{T}_{\text{obs}}(G_i),\ s \equiv_i t \Rightarrow s \equiv_{\mathcal{P}} t \tag{E1}$$

The second axiom is more subtle, and requires the notion of term vectors, crucially used in the proofs. A term vector v is a partial function mapping a finite set of dimensions $i \in \mathcal{I}$ to component terms $v_i \in \mathcal{T}_{\text{obs}}(G_i)$. They can be mapped to terms in two ways: either take the conjunction of all components ($\prod v$), or their disjunction ($\coprod v$). The three binary operations $\vee$, $\wedge$, and $\rightarrow$ can be defined coordinate-wise on term vectors so we can state the following:

$$(\prod u) \rightarrow (\coprod v) \equiv_{\mathcal{P}} \coprod (u \rightarrow v) \tag{E2}$$

With these two new axioms as well as the Heyting algebra axioms and (B1)-(B2), we obtain a sound and complete equivalence relation $\equiv_{\mathcal{P}}$.

IV - A short example

We can represent a infinite memory with two types of cells, containing either a boolean or a natural number value. Indeed, take a copy of the binary (FAN) graph $\mathcal{B}_x = \langle \{0_x, 1_x\}, \emptyset \rangle$ for each boolean variable x, and a copy $\mathbb{N}_y$ of the infinite anticlique over natural numbers for each numeric variable y. Their product $\bigotimes_x \mathcal{B}_x \otimes \bigotimes_y \mathbb{N}_y$ is a graph whose vertices correspond to predicates $var == val$ with var a variable and val a value of the appropriate type.

In this model, the semantics of a predicate $\neg(x == v) = (x == v) \rightarrow \bot$ is the set of cliques that contain an observation incoherent with $x == v$, i.e. some $x == v'$ with $v' \neq v$.

Predicates $x == y$ can be added for boolean variables, as a macro for the formula $(x == 1 \wedge y == 1) \vee (x == 0 \wedge y == 0)$ but not for numeric ones as it would entail an infinite disjunction. On the other hand a predicate $x \leqslant n$ for $n \in \mathbb{N}$ can be encoded as the formula $x == 0 \vee \cdots \vee x == n$. Therefore a predicate like $n < x$ may be added as the negation of the previous one. Some data structures may also be added to the language as syntactic sugar. For example, arrays of fixed size are easily encoded as sets of variables. In Figure 1 another finite graph is drawn that can be interpreted as the value space of a 2-dimensional variable. Observations (i.e. atomic propositions) check whether the value belongs to some disc. Two observations are coherent if their discs overlap.

V - Future work

One direction for future research is to capture more refined memory models, for instance with more datatypes such as lists. This means populating the class of Observation algebras with more coherence graphs. One would then need to check whether these are instances of our constructions (products of FAN graphs and anticliques). In the negative case, hopefully graphs of use would still allow for some reasonnable axiomatisation.

Another obvious direction to explore is the study of concurrent processes using this model inside Concurrent Kleene Algebra with Hypotheses [4]. On paper this should work exactly like partially observable CKA [7], except with a more expressive specification logic for memory states. An investigation into the sort of properties that such a system can capture might yield some interesting results.

References

[1] Paul Brunet.. Observation algebras: Heyting algebra over coherence spaces. Mar. 2025. doi: 10.48550/arXiv.2503.07130.

[2] Paul Brunet.. Repository of Rocq proofs: github.com/monstrencage/obs-alg-proofs/. 2022.

[3] Jean-Yves Girard. "Linear logic". In:. Theoretical Computer Science 50.1 (Jan. 1987). Publisher: Elsevier. doi: 10.1016/0304-3975(87)90045-4.

[4] Tobias Kappé et al. "Concurrent Kleene Algebra with Observations: From Hypotheses to Completeness". In: . FoSSaCS. Lecture Notes in Computer Science. 2020. doi: 10.1007/978-3-030-45231-5_20.

[5] Tobias Kappé et al. "Kleene Algebra with Observations". In:. CONCUR. Vol. 140. LIPIcs. 2019. doi: 10.4230/LIPIcs.CONCUR.2019.41.

[6] Dexter Kozen and Frederick Smith. "Kleene algebra with tests: Completeness and decidability". In:. CSL. Lecture Notes in Computer Science. 1997.

[7] Jana Wagemaker et al. "Partially Observable Concurrent Kleene Algebra". In:. CONCUR. LIPIcs. 2020. doi: 10.4230/LIPIcs.CONCUR.2020.20.