

Satisfiability of Arbitrary Public Announcement Logic with Common Knowledge is Σ_1^1 -hard

Rustam Galimullin

University of Bergen
Bergen, Norway

rustam.galimullin@uib.no

Louwe B. Kuijer

University of Liverpool
Liverpool, UK

lbkuijer@liverpool.ac.uk

Arbitrary Public Announcement Logic with Common Knowledge (APALC) is an extension of Public Announcement Logic with common knowledge modality and quantifiers over announcements. We show that the satisfiability problem of APALC on $S5$ -models, as well as that of two other related logics with quantification and common knowledge, is Σ_1^1 -hard. This implies that neither the validities nor the satisfiable formulas of APALC are recursively enumerable. Which, in turn, implies that APALC is not finitely axiomatisable.

1 Introduction

Quantified Public Announcement Logics. *Epistemic logic* (EL) [21] is one of the better-known formalisms for reasoning about knowledge of agents in multi-agent systems. It extends the language of propositional logic with constructs $\Box_a \varphi$ meaning that ‘agent a knows φ ’. Formulas of EL are interpreted on epistemic models (or, equivalently, $S5$ -models) that comprise a set of states, equivalence relations for each agent between states, and a valuation function that specifies in which states propositional variables are true. However, EL provides only a static description of distribution of knowledge in a system. Extensions of the logic that allow one to reason about how information of individual agents and groups thereof changes as a result of some epistemic event are generally collectively known as *dynamic epistemic logics* (DELs) [10].

The prime example of a DEL and arguably the most well-studied logic in the family is *public announcement logic* (PAL) [25]. A public announcement is an event of all agents publicly and simultaneously receiving the same piece of information. The language of PAL extends that of EL with formulas $[\psi]\varphi$ that are read as ‘after public announcement of ψ , φ is true’.

Quantification over various epistemic actions, and in particular over public announcements, has been explored in the last 15 or so years [9]. Adding quantification over public announcements allows one to shift the emphasis from the effects of a particular announcement to the question of (non-)existence of an announcement leading to a desired epistemic goal. In this paper, we focus on the three, perhaps most well-known, *quantified PALs* (QPALs). The first of the three is *arbitrary PAL* (APAL) [6] that extends the language of PAL with constructs $[!]\varphi$ meaning ‘after *any* public announcement, φ is true’. A formula with the dual existential quantifier $\langle ! \rangle \varphi$ is read as ‘*there is* a public announcement, after which φ is true’.

Observe that quantifiers of APAL do not specify whether an announcement can be made by any of the agents, or groups thereof, modelled in a system. Hence, a more ‘agent-centric’ quantified PAL was proposed. *Group announcement logic* (GAL) [1] extends the language of PAL with formulas $[G]\varphi$ meaning ‘after *any* announcement by agents from group G , φ is true’. A formula with the dual of the universal GAL quantifier is $\langle G \rangle \varphi$ that is read ‘*there is* an announcement by agents from group G that makes φ true’.

Once we start reasoning about what groups of agents can achieve by making public announcements, it is only too natural to consider their abilities in a game-theoretic setting. In particular, we may let agents outside of the group make their own announcements in an attempt to preclude the group from reaching their epistemic goals. A QPAL with such a competitive flavour to it is called *coalition announcement logic* (CAL) [2, 15]. The logic extends PAL with modalities $\llbracket G \rrbracket \varphi$ that are read as ‘whatever agents from coalition G announce, *there is* a counter-announcement by the anti-coalition that makes φ true’. The diamond version $\langle\langle G \rangle\rangle \varphi$ is then means that ‘*there is* an announcement by coalition G , such that *whatever* the anti-coalition announces at the same time, they cannot avoid φ ’. Observe, that compared to APAL and GAL, modalities of CAL contain double quantification: $\forall\exists$ and $\exists\forall$ correspondingly. As the name of the logic suggests, modalities of CAL were inspired by coalition logic [24], and they capture game-theoretic notions of α - and β -effectivity [5].

Some Logical Properties of QPALs. One of the most pressing open problems in the area is the existence of finitary axiomatisations of QPALs. Both finitary and infinitary axiom systems for APAL were proposed in [6], but later the finitary version was shown to be unsound [19]. The infinitary axiomatisation is, however, sound and complete [7]. As the axiomatisation of GAL [1] is quite similar to that of APAL, its finitary version is also not sound [13, Footnote 4], and its infinitary version can be shown to be sound and complete by a modification of the proof from [7]. To the best of our knowledge, there are no known sound and complete proof systems, finitary or infinitary, for CAL¹.

The satisfiability problem for QPALs is known to be undecidable [3]. The result is achieved by a reduction from the classic tiling problem that consists in answering the question whether a given finite set of tiles can tile the $\mathbb{N} \times \mathbb{N}$ plane. Since this problem is co-RE-complete [8, 17], or, equivalently, Π_1^0 -complete, the reduction amounts to the fact that the satisfiability problem for QPALs is co-RE-hard (or Π_1^0 -hard). Note that this result does not rule out the existence of finitary axiomatisations of QPALs. A prime example of a logic with a co-RE-complete satisfiability problem and a finitary axiomatisation is first-order logic.

Overview of the paper and our result. In this paper we consider extensions of QPALs with *common knowledge* [12], which is a classic variant of group knowledge in multi-agent systems. Its intuitive meaning is that ‘ φ is common knowledge among agents in group G if everyone in G knows φ , everyone in G knows that everyone in G knows φ and so on ad infinitum’. Semantically, common knowledge among agents from G corresponds to the reflexive transitive closure of equivalence relations of all agents from group G . We call extensions of APAL, GAL, and CAL with common knowledge APALC [4], GALC, and CALC, correspondingly, or QPALCs if we refer to all of them at the same time.

The result we prove in this paper is that the satisfiability problems for QPALCs are Σ_1^1 -hard. We do this by showing that the *recurring tiling problem*, which is known to be Σ_1^1 -complete [18], can be reduced to satisfiability of QPALC formulas. Because the satisfiability problems are Σ_1^1 -hard, it follows that, in particular, the set of valid QPALC formulas is not recursively enumerable. That, in turn, implies that QPALCs have no finitary axiomatisations. The non-existence of a finitary axiomatisation of a somewhat related arbitrary arrow update logic [11] with common knowledge was shown in [20] by the reduction from the non-halting problem. Moreover, the recurring tiling problem was used in [22] to demonstrate that the satisfiability problem of PAL with iterated announcements and common knowledge is Σ_1^1 -complete.

The use of common knowledge is instrumental in our paper, since it allows us to have a ‘tighter’ grid than the ones from [3] and [14]. We deem our result important in at least two ways. First, the non-existence of finitary axiomatisations of QPALCs is interesting in its own right as it demonstrates

¹A complete infinitary axiomatisation with CAL modalities and additional operators was given in [16]

that presence of common knowledge in QPALCs is a sufficient condition for Σ_1^1 -hardness. Second, having both our construction (with common knowledge) and the constructions from [3] and [14] side by side, allows one to flesh out crucial differences between Σ_1^1 -hardness and Σ_1^0 -hardness arguments, and, hopefully, move closer to tackling the open problem of (non-)existence of finitary axiomatisations of QPALs.

Outline of the paper. The rest of the paper is organised as follows. In Section 2 we cover the background on QPALCs. After that, in Section 3, we prove the main claim of this paper, and, finally, we conclude in Section 4.

2 Quantified Public Announcement Logics with Common Knowledge

Let A be a finite set of agents, and P be a countable set of propositional variables.

Definition 2.1. The languages of arbitrary public announcement logic with common knowledge APALC, group announcement logic with common knowledge GALC, and coalition announcement logic with common knowledge CALC are inductively defined as

$$\begin{aligned} \text{APALC} &\ni \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid [\varphi]\varphi \mid \blacksquare_G\varphi \mid [!]\varphi \\ \text{GALC} &\ni \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid [\varphi]\varphi \mid \blacksquare_G\varphi \mid [G]\varphi \\ \text{CALC} &\ni \varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid [\varphi]\varphi \mid \blacksquare_G\varphi \mid \llbracket G \rrbracket\varphi \end{aligned}$$

where $p \in P$, $a \in A$, and $G \subseteq A$. Duals are defined as $\Diamond_a\varphi := \neg\Box_a\neg\varphi$, $\langle \psi \rangle\varphi := \neg[\psi]\neg\varphi$, $\blacklozenge_G\varphi := \neg\blacksquare_G\neg\varphi$, $\langle ! \rangle\varphi := \neg[!]\neg\varphi$, $\langle G \rangle\varphi := \neg[G]\neg\varphi$ and $\llbracket G \rrbracket\varphi := \neg\llbracket G \rrbracket\neg\varphi$.

The fragment of APALC without $[!]\varphi$ is called *public announcement logic with common knowledge* PALC; the latter without $[\varphi]\varphi$ is *epistemic logic with common knowledge* ELC; PALC and ELC minus $\blacksquare_G\varphi$ are, correspondingly, *public announcement logic* PAL and *epistemic logic* EL. Finally, fragments of APALC, GALC and CALC without $\blacksquare_G\varphi$ are called *arbitrary public announcement logic* APAL, *group announcement logic* GAL and *coalition announcement logic* CAL respectively.

Definition 2.2. A model M is a tuple (S, \sim, V) , where S is a non-empty set of states, $\sim: A \rightarrow 2^{S \times S}$ gives an equivalence relation for each agent, and $V: P \rightarrow 2^S$ is the valuation function. By \sim_G we mean reflexive transitive closure of $\bigcup_{a \in G} \sim_a$. We will denote model M with a distinguished state s as M_s .

We would like to stress that agent relations in our models are *equivalence relations* (and hence our models are *S5* models). The results of this paper do not generalise to arbitrary agent relations in any obvious way.

It is assumed that for group announcements, agents know the formulas they announce. In the following, we write $\text{PALC}^G = \{\bigwedge_{i \in G} \Box_i \psi_i \mid \text{for all } i \in G, \psi_i \in \text{PALC}\}$ to denote the set of all possible announcements by agents from group G . We will use ψ_G to denote arbitrary elements of PALC^G .

Definition 2.3. Let $M_s = (S, R, V)$ be a model, $p \in P$, $G \subseteq A$, and $\varphi, \psi \in \text{APALC} \cup \text{GALC} \cup \text{CALC}$.

$$\begin{aligned}
M_s \models p & \quad \text{iff} \quad s \in V(p) \\
M_s \models \neg\varphi & \quad \text{iff} \quad M_s \not\models \varphi \\
M_s \models \varphi \wedge \psi & \quad \text{iff} \quad M_s \models \varphi \text{ and } M_s \models \psi \\
M_s \models \Box_a \varphi & \quad \text{iff} \quad \forall t \in S : s \sim_a t \text{ implies } M_t \models \varphi \\
M_s \models \blacksquare_G \varphi & \quad \text{iff} \quad \forall t \in S : s \sim_G t \text{ implies } M_t \models \varphi \\
M_s \models [\psi]\varphi & \quad \text{iff} \quad M_s \models \psi \text{ implies } M_s^\psi \models \varphi \\
M_s \models [!]\varphi & \quad \text{iff} \quad \forall \psi \in \text{PALC} : M_s \models [\psi]\varphi \\
M_s \models [G]\varphi & \quad \text{iff} \quad \forall \psi_G \in \text{PALC}^G : M_s \models [\psi_G]\varphi \\
M_s \models \llbracket G \rrbracket \varphi & \quad \text{iff} \quad \forall \psi_G \in \text{PALC}^G, \exists \chi_{A \setminus G} \in \text{PALC}^{A \setminus G} : M_s \models \psi_G \text{ implies } M_s \models \langle \psi_G \wedge \chi_{A \setminus G} \rangle \varphi
\end{aligned}$$

where $M_s^\psi = (S^\psi, R^\psi, V^\psi)$ with $S^\psi = \{s \in S \mid M_s \models \psi\}$, $R^\psi(a)$ is the restriction of $R(a)$ to S^ψ for all $a \in A$, and $V^\psi(p) = V(p) \cap S^\psi$ for all $p \in P$.

Observe, that it follows from the definition of the semantics that in the case of the grand coalition A , $M_s \models [A]\varphi$ if and only if $M_s \models \llbracket A \rrbracket \varphi$. For the case of the empty group \emptyset , we assume that the conjunction of an empty set of formulas is a tautology.

Remark 1. For APAL, GAL, and CAL, we assume that quantification ranges over a quantifier-free fragment of the language, i.e. over PAL, which is equally expressive as EL [25]. This is, however, not as straightforward once we consider ELC and PALC. The latter is strictly more expressive than ELC [10, Theorem 8.48], and ELC, in its turn, is strictly more expressive than EL, and thus it matters, expressivity-wise, which quantifier-free fragment of a QPALC the quantification ranges over. These matters are explored in [4], where also infinitary axiomatisations of APALC and GALC are given. For our current purposes, though, the difference in the range of quantification does not play a role.

3 The Satisfiability Problem of QPALCs is Σ_1^1 -hard

We prove the Σ_1^1 -hardness of the satisfiability problem of QPALCs via a reduction from the recurring tiling problem [17].

Definition 3.1. Let C be a finite set of *colours*. A *tile* is a function $\tau : \{\text{north, south, east, west}\} \rightarrow C$. A finite set of tiles T is called an *instance* of the tiling problem. A *solution* to an instance of the tiling problem is a function² $f : \mathbb{N} \times \mathbb{N} \rightarrow T$ such that for all $(i, j) \in \mathbb{N} \times \mathbb{N}$,

$$f(i, j)(\text{north}) = f(i, j+1)(\text{south}) \text{ and } f(i, j)(\text{east}) = f(i+1, j)(\text{west}).$$

Definition 3.2. Let T be a finite set of tiles with a designated tile $\tau^* \in T$. The *recurring tiling problem* is the problem to determine whether there is a solution to instance T of the tiling problem such that τ^* appears *infinitely* often in the first column.

We assume without loss of generality that the designated tile τ^* occurs only in the first column.

²Throughout the paper we assume that $0 \in \mathbb{N}$.

3.1 Encoding a Tiling

For our construction we will require five propositional variables — north, south, east, west and centre — to designate the corresponding sides of tiles. Additionally, we will have designated propositional variables for each colour in C , and for each tile $\tau_i \in T$ there is a propositional variable p_i that represents this tile. Finally, we will use p^* for the special τ^* .

In our construction, we will represent each tile with (at least) five states: one for each of the four sides of a tile, and one for the centre. As for agents, we require only three of them for our construction. Agent s , for square, cannot distinguish states within the same tile. Agent v , for vertical, cannot distinguish between the northern part of one tile and the southern part of the tile above. Similarly, the horizontal agent h cannot distinguish between the eastern and western parts of adjacent tiles. See Figure 1 for the depiction of an intended grid-like model.




Figure 1: Left: a representation of a single tile τ_i , where agent s has the universal relation within the dashed square, relations h and v are equivalences, and reflexive arrows are omitted. Each state is labelled by a set of propositional variables that are true there. Right: an example of a grid-like model that we construct in our proof. Each tile τ has a similar structure as presented on the left of the figure.

Let an instance T of the recurring tiling problem be given. We start by construction of formula Ψ_T that will be satisfied in a given model if and only if the model is grid-like. We will build up Ψ_T step-by-step, defining useful subformulas along the way. Let Position be the following set $\text{Position} := \{\text{north}, \text{south}, \text{east}, \text{west}, \text{centre}\}$.

The first constraint, expressed by formula one_colour , is that each state is coloured by exactly one colour. To ensure that all five parts — north, south, east, west, and centre — are present in a current square, we state in all_parts that in all squares the square agent s has access to all five relevant states.

$$one_colour := \bigvee_{c \in C} \left(c \wedge \bigwedge_{d \in C \setminus \{c\}} \neg d \right) \quad all_parts := \Box_s \bigvee_{q \in \text{Position}} q \wedge \bigwedge_{q \in \text{Position}} \Diamond_s q$$

The formulas hor and $vert$ state that the relation h only allows us to move between east and west states, while v only allows movement between north and south states.

$$hor := \bigwedge_{q \in \{\text{north}, \text{south}, \text{centre}\}} (q \rightarrow \Box_h q) \quad vert := \bigwedge_{q \in \{\text{east}, \text{west}, \text{centre}\}} (q \rightarrow \Box_v q)$$

With one_pos we force each state to satisfy exactly one propositional variable from $Position$, and with one_tile we ensure that all states within the same tile are labelled by the tile proposition.

$$one_pos := \bigvee_{q \in Position} \left(q \wedge \bigwedge_{q' \in Position \setminus \{q\}} \neg q' \right) \quad one_tile := \bigvee_{\tau_i \in T} \left(p_i \wedge \bigwedge_{\tau_j \in T \setminus \{\tau_i\}} \neg p_j \right)$$

Next, we force each state in a tile to satisfy exactly one atom corresponding to their designated colour:

$$state_col := \bigvee_{\tau_i \in T} \left(p_i \rightarrow \bigwedge_{q \in Position \setminus \{centre\}} (q \rightarrow \tau_i(q)) \right),$$

where $\tau_i(q)$ is the colour of the tile τ_i on the side q (e.g. $\tau_i(\text{south})$ is the bottom colour of tile τ_i).

All the formulas considered so far deal with the representation of a single tile. We will use the following abbreviation:

$$\Psi_{tile} := one_colour \wedge all_parts \wedge hor \wedge vert \wedge one_pos \wedge one_tile \wedge state_col$$

Adjoining tiles are required to have the same colour on the sides facing each other, we simulate this by requiring that agents h and v consider a current colour in the top and right directions. In such a way we also ensure that the grid is infinite in the positive quadrant.

$$adj_tiles := \bigwedge_{c \in C} ((north \wedge c \rightarrow \diamond_v south \wedge \square_v c) \wedge (east \wedge c \rightarrow \diamond_h west \wedge \square_h c))$$

We are concerned with the reduction from the $\mathbb{N} \times \mathbb{N}$ recurring tiling problem, i.e. our grid will have left and bottom edges. We force the existence of a tile at position $(0,0)$ with the following formula:

$$init := \blacklozenge_{\{h,v,s\}} (\blacksquare_{\{v,s\}} (west \rightarrow \square_h west) \wedge \blacksquare_{\{h,s\}} (south \rightarrow \square_v south))$$

For the remaining formulas, it is useful to define two abbreviations. We use $\square_{up} \varphi$ to denote $\square_s (north \rightarrow \square_v (south \rightarrow \varphi))$, i.e., we first move, by agent s , to the state representing the northern quadrant of the tile, then we move, by agent v , to southern quadrant of the tile above, where we evaluate φ . Similarly, we use $\square_{right} \varphi$ to denote $\square_s (east \rightarrow \square_h (west \rightarrow \varphi))$. The duals \diamond_{up} and \diamond_{right} are defined as usual.

The next two formulas are used to guarantee that for every tile there are unique tiles, up to PALC-indistinguishability, above it and to its right.

$$up := [!](\diamond_{up} \diamond_s centre \rightarrow \square_{up} \diamond_s centre) \\ right := [!](\diamond_{right} \diamond_s centre \rightarrow \square_{right} \diamond_s centre)$$

Additionally, we use the following two formulas to establish a commutative property: going up and then $right$ results in a state that is PALC-indistinguishable from going $right$ and then up .

$$right\&up := [!](\diamond_{right} \diamond_{up} \diamond_s centre \rightarrow \square_{up} \square_{right} \diamond_s centre) \\ up\&right := [!](\diamond_{up} \diamond_{right} \diamond_s centre \rightarrow \square_{right} \square_{up} \diamond_s centre)$$

Finally, we make sure that any two states that are h or v related and that are in the same position are parts of indistinguishable tiles.

$$no_change := \bigwedge_{q,q' \in \text{Position}} [!)((q \wedge \diamond_s q') \rightarrow (\Box_h(q \rightarrow \diamond_s q') \wedge \Box_v(q \rightarrow \diamond_s q')))$$

The formula *hor* states that unless we are in a east or west position, we cannot go to a different position using *h*. Similarly, *vert* states that unless we are in a north or south position we can't use *v* to change position. The formula *no_change* then states that any move by relation *h* or *v* that does not change the position must lead to an indistinguishable tile.

We abbreviate formulas with quantifiers as

$$\Psi_{x\&y} := up \wedge right \wedge right\&up \wedge up\&right \wedge no_change$$

In our reduction, we are interested in grids where a special tile appears infinitely often in the first column of the grid. The following formula requires that the special tile appears only in the leftmost column:

$$tile_left := p^* \rightarrow \Box_s(\text{west} \rightarrow \Box_h \text{west})$$

All of this completes the necessary requirements for the grid. Now, by adding a common knowledge modality for all agents, we force all of the aforementioned formulas to hold everywhere in the grid.

$$\Psi_T := \blacksquare_{\{h,v,s\}} (\Psi_{tile} \wedge adj_tiles \wedge init \wedge \Psi_{x\&y} \wedge tile_left)$$

Observe that Ψ_T does not say anything about the special tile τ^* appearing infinitely often in the first column. The formula merely requires that if there is a special tile, then it should appear in the first column. We first show that Ψ_T forces a grid-like model, and only after that will we consider the (in)finite number of occurrences of the special tile.

Lemma 1. Let T be an instance of the recurring tiling problem. If T can tile $\mathbb{N} \times \mathbb{N}$, then Ψ_T is satisfiable.

Proof. Assume that there is a tiling of the $\mathbb{N} \times \mathbb{N}$ plane with a finite set of tiles T . We construct model $M = (S, \sim, V)$ satisfying Ψ_T directly from the given tiling. In particular,

- $S = \mathbb{N} \times \mathbb{N} \times \{\mathfrak{n}, \mathfrak{s}, \mathfrak{e}, \mathfrak{w}, \mathfrak{c}\}$,
- $\sim_s = \{(i, j, \mathfrak{l}), (i', j', \mathfrak{l}') \mid i = i' \text{ and } j = j'\}$
- \sim_v is the reflexive closure of $\{(i, j, \mathfrak{n}), (i, j + 1, \mathfrak{s})\}$
- \sim_h is the reflexive closure of $\{(i, j, \mathfrak{e}), (i + 1, j, \mathfrak{w})\}$
- for all $\tau_k \in T$, $V(p_k) = \{(i, j, \mathfrak{l}) \mid \tau_k \text{ is at } (i, j)\}$
- for all $c \in C$, $V(c) = \{(i, j, \mathfrak{l}) \mid \tau(\mathfrak{l}) = c\}$
- for all $l \in \text{Position}$, $V(l) = \{(i, j, \mathfrak{l}) \mid l \text{ corresponds to } \mathfrak{l}\}$

To argue that $M_{(0,0,\mathfrak{e})} \models \Psi_T$ we first notice that due to the fact that T tiles the $\mathbb{N} \times \mathbb{N}$ plane and by the construction of M , subformulas of Ψ_T that do not involve arbitrary announcements are straightforwardly satisfied.

Now, consider the formula *up*. For every (i, j, \mathfrak{l}) , there is at most one (i', j', \mathfrak{l}') that is reachable by taking an *s*-step to a north state followed by a *v*-step to a south state, namely $(i', j', \mathfrak{l}') = (i, j + 1, \mathfrak{s})$. Furthermore, this property is retained in any submodel of M . As a consequence, in any state of any submodel of M , $\diamond_{up}\chi$ implies $\Box_{up}\chi$, for every χ . In particular, it follows that $M_{(i,j,\mathfrak{l})} \models [!](\diamond_{up}\diamond_s \text{centre} \rightarrow \Box_{up}\diamond_s \text{centre})$, i.e., $M_{(i,j,\mathfrak{l})} \models up$.

Similar reasoning shows that (i, j, \mathfrak{l}) satisfies the other conjuncts of $\Psi_{x\&y}$. Hence $M_{(i,j,\mathfrak{l})} \models \Psi_{tile} \wedge adj_tiles \wedge init \wedge \Psi_{x\&y} \wedge tile_left$, for all (i, j, \mathfrak{l}) , and thus $M_{(0,0,\mathfrak{e})} \models \Psi_T$. \square

The more complex part of the reduction is to show that if Ψ_T is satisfiable, then a tiling exists.

Lemma 2. Let T be an instance of the recurring tiling problem. If Ψ_T is satisfiable, then T can tile $\mathbb{N} \times \mathbb{N}$.

Proof. Let M be such that $M_s \models \Psi_T$. The model M is partitioned by \sim_s , we refer to these partitions as grid points, and label these points as follows.

- The grid point containing s is labelled $(0, 0)$.
- If A and B are grid points, A is labelled (i, j) and there is a north-state in A that is v -indistinguishable to a south-state in B , then B is labelled $(i, j + 1)$.
- If A and B are grid points, A is labelled (i, j) and there is a east-state in A that is h -indistinguishable to a west-state in B , then B is labelled $(i + 1, j)$.

Note that a single grid point might have multiple labels. We say that (i, j) is tiled with τ_i if there is some grid point labelled with (i, j) that contains a state where p_i holds. We start by noting that because the main connective of Ψ_T is $\blacksquare_{\{h,v,s\}}$, the formula holds in every labelled grid point. For every labelled grid point X and every $x \in X$, we therefore have $M_x \models \psi_{tile}$. So X contains states for every direction, each labelled with exactly one colour that corresponds to the tile that holds on X . We continue by proving the following claim.

Claim 1: Let X, A and B be grid points where X is labeled (i, j) while A and B are both labeled $(i, j + k)$ by virtue of being k -steps to the north of X . Then A and B are PALC-indistinguishable, in the sense that for every $\chi \in \text{PALC}$, if there is an $a \in A$ such that $M_a \models \chi$ then there is a $b \in B$ such $M_b \models \chi$ (and vice versa).

Proof of Claim 1: By induction on k . As base case, let $k = 1$ and suppose towards a contradiction that, for some $\chi \in \text{PALC}$ and $a \in A$, $M_a \models \chi$ while for every $b \in B$, $M_b \not\models \chi$. Consider then the formula $\text{centre} \rightarrow \diamond_s \chi$. Every centre state in A satisfies this formula, while none of the centre states in B do. Hence, for every state $x \in X$, $M_x \models [\text{centre} \rightarrow \diamond_s \chi](\diamond_{up} \diamond_s \text{centre} \wedge \neg \square_{up} \diamond_s \text{centre})$. But that contradicts $M_x \models up$. From this contradiction, we prove the base case $k = 1$.

Now, suppose as induction hypothesis that $k > 1$ and that the claim holds for all $k' < k$. Again, suppose towards a contradiction that $M_a \models \chi$ while $M_b \not\models \chi$ for all $b \in B$. Let A' and B' be grid points that lie $k - 1$ steps to the north of X and one step to the south of A and B , respectively. Then for every $a' \in A'$ and $b' \in B'$, $M_{a'} \models \diamond_{up} \diamond_s \chi$ and $M_{b'} \models \diamond_{up} \neg \diamond_s \chi$. By the induction hypothesis, A' and B' are indistinguishable, so $M_{a'} \models \diamond_{up} \diamond_s \chi \wedge \diamond_{up} \neg \diamond_s \chi$. But then there are distinguishable grid points one step to the north of A' , contradicting the induction hypothesis. From this contradiction, we prove the induction step and thereby the claim.

Similar reasoning shows that any two grid points A, B that are labeled $(i + k, j)$ by virtue of being k steps to the right of the same grid point X are indistinguishable. Now, we can prove the next claim.

Claim 2: Let X, A and B be grid points, where X is labelled (i, j) , A is labelled $(i + 1, j + 1)$ by virtue of being above A' which is to the right of X , and B is labelled $(i + 1, j + 1)$ by virtue of being to the right of B' which is above B . Then A and B are PALC-indistinguishable.

Proof of claim 2: Suppose towards a contradiction that for some $\chi \in \text{PALC}$ and $a \in A$ we have $M_a \models \chi$, while $M_b \not\models \chi$ for all $b \in B$. Then for $x \in X$ we have $M_x \models [\text{centre} \rightarrow \diamond_s \chi](\diamond_{right} \diamond_{up} \diamond_s \text{centre} \wedge \diamond_{up} \diamond_{right} \neg \diamond_s \text{centre})$, contradicting $M_x \models right \& up$.

From Claim 1 it follows that any A and B that are labelled (i, j) by virtue of being i steps to the right and then j steps up from $(0, 0)$ are PALC-indistinguishable. Claim 2 then lets us commute the “up” and

“right” moves. Any path to (i, j) can be obtained from the path that first goes right i steps then up j steps by a finite sequence of such commutations. Hence any grid points A and B that are labelled (i, j) are PALC-indistinguishable.

The tile formulas p_i , for every $\tau_i \in T$, are PALC-formulas, so there is exactly one tile τ_i that is assigned to the grid point (i, j) . Furthermore, *state_col* then guarantees that each side of a grid point has the colour corresponding to the tile, and *adj_tiles* guarantees that the tile colours match. This shows that if Ψ_T is satisfiable, then T can tile $\mathbb{N} \times \mathbb{N}$. \square

3.2 Encoding the Recurring Tile

The final formula that is satisfied in a grid model if and only if a given tiling has a tile that occurs infinitely often in the first column would be

$$\Psi_T \wedge \blacksquare_{\{v,s\}} [\blacksquare_{\{h,s\}} \neg p^*] \neg \Psi_T.$$

In other words, the recurring tiling problem can be reduced to the APALC-satisfiability problem, where the reduction maps the instance (T, τ^*) of the recurring tiling problem to the satisfiability of $\Psi_T \wedge \blacksquare_{\{v,s\}} [\blacksquare_{\{h,s\}} \neg p^*] \neg \Psi_T$.

Intuitively, the formula states that if we remove all rows with the special tile, then our model is no longer a grid. See Figure 2, where on the left we have a grid with the special grey tile τ^* appearing infinitely often in the first column (every other tile in the first column is grey). Formula $\blacksquare_{\{h,s\}} \neg p^*$ holds only in those squares of the grid that lie on rows without the special tile. Thus, announcing $\blacksquare_{\{h,s\}} \neg p^*$ removes all rows that has the grey tile (see the right part of Figure 2). Since the grey tile appears infinitely often in the original grid, we have to remove an infinite number of rows after the announcement of $\blacksquare_{\{h,s\}} \neg p^*$, thus ensuring that what is left of the original model is not a grid.




Figure 2: Left: An original grid with a special grey tile τ^* appearing infinitely often in the first column. Right: The grid after the public announcement of $\blacksquare_{\{h,s\}} \neg p^*$. Crossed-out rows are not preserved after the announcement.

Theorem 1. Let T be an instance of the tiling problem with a special tile $\tau^* \in T$. Set T can tile $\mathbb{N} \times \mathbb{N}$ with τ^* appearing infinitely often in the first column if and only if $\Psi_T \wedge \blacksquare_{\{v,s\}} [\blacksquare_{\{h,s\}} \neg p^*] \neg \Psi_T$ is satisfiable.

Proof. First, let us can extend the labelling from the proof of Lemma 2 as follows:

- For every $q \in \text{Position}$, if A and B are grid points, A is labeled (i, j) and there is a q state in A that is v or h -indistinguishable from a q state in B , then B is labeled (i, j) .

It follows from *no_change* that this extended labelling retains the property that any two grid points with the same label are PALC-indistinguishable. Furthermore, from *hor* and *vert* it follows that every grid point that is reachable by *h*, *v* and *s* is now labelled with some coordinates (i, j) . Hence we can identify the $\{h, v, s\}$ -reachable grid points in any model of Ψ_T with $\mathbb{N} \times \mathbb{N}$.

Now, assume that set T cannot tile the $\mathbb{N} \times \mathbb{N}$ plane with a special tile $\tau' \in T$ appearing infinitely often in the first column. We argue that in this case, $\Psi_T \wedge \blacksquare_{\{v,s\}}([\blacksquare_{\{h,s\}} \neg p^*] \neg \Psi_T)$ is not satisfiable. The first conjunct is straightforward. If T cannot tile the $\mathbb{N} \times \mathbb{N}$ plane, then, by Lemma 2, Ψ_T is not satisfiable.

So suppose that T can tile the plane, but only in such a way that τ^* occurs finitely often. For every model $M_{(0,0,l)}$ of Ψ_T , there is then some $k \in \mathbb{N}$ that is the last row in which p^* is true. The formula $\blacksquare_{\{h,s\}} \neg p^*$ holds exactly on those rows where p^* does not hold in the first column. As a result, the update $[\blacksquare_{\{h,s\}} \neg p^*]$ does not remove any rows past row k . The grid points $\mathbb{N} \times \mathbb{N}_{>k}$ then still form a grid that is isomorphic to $\mathbb{N} \times \mathbb{N}$, and that is tiled. See Figure 3 for a depiction of the situation.

It follows that $M_{(0,k,l)} \not\models [\blacksquare_{\{h,s\}} \neg p^*] \neg \Psi_T$, and therefore $M_{(0,0,l)} \not\models \blacksquare_{\{v,s\}}[\blacksquare_{\{h,s\}} \neg p^*] \neg \Psi_T$. This is true for every model of Ψ_T , so $\Psi_T \wedge \blacksquare_{\{v,s\}}[\blacksquare_{\{h,s\}} \neg p^*] \neg \Psi_T$ is not satisfiable.

If, on the other hand, T can tile the plane in such a way that τ^* occurs infinitely often in the first column, then there is a model of Ψ_T where the modality $[\blacksquare_{\{h,s\}} \neg p^*]$ removes infinitely many rows, and therefore does not leave any infinite grid. So $\Psi_T \wedge \blacksquare_{\{v,s\}}[\blacksquare_{\{h,s\}} \neg p^*] \neg \Psi_T$ is satisfiable. \square




Figure 3: Left: An original grid with a special grey tile τ^* appearing finitely often in the first column. Right: The grid after the public announcement of $\blacksquare_{\{h,s\}} \neg p^*$. Crossed-out rows are not preserved after the announcement. A full $\mathbb{N} \times \mathbb{N}$ grid that is still available after the announcement is depicted with thick lines.

In the construction of Ψ_T and proofs of Lemmas 1 and 2, we used APALC quantifiers $[!]$. We can prove the similar results for GALC and CALC quantifiers by substituting $[!]$ with $[\{h, v, s\}]$ and $[\{\{h, v, s\}\}]$ correspondingly, and substituting PALC with $\text{PALC}^{\{h,v,s\}}$. We get the hardness result from the Σ_1^1 -completeness of the recurring tiling problem [18].

Theorem 2. The satisfiability problem of QPALCs is Σ_1^1 -hard.

The Σ_1^1 -hardness of the satisfiability problems of QPALCs together with the fact that the class of Σ_1^1 problems is strictly greater than the class of co-RE problems [23, Chapter 4] imply that the sets of validities of the logics are not RE, which, in turn, implies that QPALCs are not finitely axiomatisable.

Corollary 1. The set of valid formulas of QPALCs is neither RE nor co-RE.

Corollary 2. QPALCs do not have finitary axiomatisations.

4 Discussion

The existence of finitary axiomatisations of any of APAL, GAL, and CAL is a long-standing open problem. In this paper, we have showed that the satisfiability problem of the logics extended with common knowledge modality is Σ_1^1 -hard, and thus they do not admit of finitary axiomatisations. Table 1 contains the overview of the known results, including those shown in this paper, and open questions.

	APAL	GAL	CAL	APALC	GALC	CALC
Finitary axiomatisation	?	?	?	✗ (Cor. 2)	✗ (Cor. 2)	✗ (Cor. 2)
Infinitary axiomatisation	✓[6]	✓[1]	?	✓[4]	✓[4]	?

Table 1: Overview of the known results and open problems.

It is important to point out that the use of common knowledge is instrumental in our construction. Arguments from [14, 3] did not rely on common knowledge to enforce local grid properties globally, and instead the authors used an agent with the universal relation over the set of states. This approach is good enough if one wants to demonstrate the existence of a grid-like model. However, if we also require that the grid satisfies some property, like a special tile occurring infinitely often in the first column, then the presence of the global agent makes it harder to ensure this. The problem is that such an unrestrained relation may access other grids within the same model, and thus we may end up in the situation when the property is satisfied by a set of grids taken together and not by any single grid.

Our construction is ‘tighter’ than those in [14, 3]. In particular, our vertical and horizontal agents can ‘see’ only one step ahead. This guarantees that we stay within the same grid. In order to force grid properties globally, we use common knowledge operators that allow us to traverse a given grid-like model in all directions. It is not yet clear how to have a ‘tight’ grid and still be able to traverse the model without common knowledge. With this work, apart from showing that QPALCs are Σ_1^1 -hard, we also hope to have elucidated the exact obstacle one has to overcome in order to claim the same about QPALs.

Acknowledgements

We would like to thank the three anonymous reviewers for their encouraging comments and constructive suggestions, which helped us to improve the presentation of our result.

References

- [1] Thomas Ågotnes, Philippe Balbiani, Hans van Ditmarsch & Pablo Seban (2010): *Group announcement logic*. *Journal of Applied Logic* 8(1), pp. 62–81, doi:10.1016/j.jal.2008.12.002.
- [2] Thomas Ågotnes & Hans van Ditmarsch (2008): *Coalitions and announcements*. In Lin Padgham, David C. Parkes, Jörg P. Müller & Simon Parsons, editors: *Proceedings of the 7th AAMAS, IFAAMAS*, pp. 673–680.
- [3] Thomas Ågotnes, Hans van Ditmarsch & Tim French (2016): *The undecidability of quantified announcements*. *Studia Logica* 104(4), pp. 597–640, doi:10.1007/s11225-016-9657-0.
- [4] Thomas Ågotnes & Rustam Galimullin (2023): *Quantifying over information change with common knowledge*. *Autonomous Agents and Multi-Agent Systems* 37(19), p. 40, doi:10.1007/s10458-023-09600-1.
- [5] Robert J. Aumann (1961): *The core of a cooperative game without side payments*. *Transactions of the American Mathematical Society* 98(3), pp. 539–552, doi:10.2307/1993348.

- [6] Philippe Balbiani, Alexandru Baltag, Hans van Ditmarsch, Andreas Herzig, Tomohiro Hoshi & Tiago de Lima (2008): 'Knowable' as 'known after an announcement'. *Review of Symbolic Logic* 1(3), pp. 305–334, doi:10.1017/S1755020308080210.
- [7] Philippe Balbiani & Hans van Ditmarsch (2015): *A simple proof of the completeness of APAL*. *Studies in Logic* 8(2), pp. 65–78.
- [8] Robert L. Berger (1966): *The undecidability of the domino problem*. *Memoirs of the American Mathematical Society*, pp. 1–38.
- [9] Hans van Ditmarsch (2023): *To be announced*. *Information and Computation* 292, pp. 1–42, doi:10.1016/j.ic.2023.105026.
- [10] Hans van Ditmarsch, Wiebe van der Hoek & Barteld Kooi (2008): *Dynamic epistemic logic*. *Synthese Library* 337, Springer, doi:10.1007/978-1-4020-5839-4.
- [11] Hans van Ditmarsch, Wiebe van der Hoek, Barteld Kooi & Louwe B. Kuijer (2017): *Arbitrary arrow update logic*. *Artificial Intelligence* 242, pp. 80–106, doi:10.1016/j.artint.2016.10.003.
- [12] Ronald Fagin, Joseph Y. Halpern, Yoram Moses & Moshe Vardi (1995): *Reasoning about knowledge*. The MIT Press.
- [13] Jie Fan (2016): *Removing your ignorance by announcing group ignorance: a group announcement logic for ignorance*. *Studies in Logic* 9(4), pp. 4–33.
- [14] Tim French & Hans van Ditmarsch (2008): *Undecidability for arbitrary public announcement logic*. In Carlos Areces & Robert Goldblatt, editors: *Proceedings of the 7th AiML*, College Publications, pp. 23–42.
- [15] Rustam Galimullin (2019): *Coalition announcements*. Ph.D. thesis, University of Nottingham, UK.
- [16] Rustam Galimullin (2021): *Coalition and relativised group announcement logic*. *Journal of Logic, Language and Information* 30(3), pp. 451–489, doi:10.1007/s10849-020-09327-2.
- [17] David Harel (1985): *Recurring dominoes: making the highly undecidable highly understandable*. In Marek Karpinski & Jan van Leeuwen, editors: *Topics in the Theory of Computation, North-Holland Mathematics Studies* 102, North-Holland, pp. 51–71, doi:10.1016/S0304-0208(08)73075-5.
- [18] David Harel (1986): *Effective transformations on infinite trees, with applications to high undecidability, dominoes, and fairness*. *Journal of the ACM* 33(1), pp. 224–248, doi:10.1145/4904.4993.
- [19] Louwe B. Kuijer (2015): *Unsoundness of $R(\Box)$* . *Manuscript*. Available at https://personal.us.es/hvd/APAL_counterexample.pdf.
- [20] Louwe B. Kuijer (2017): *Arbitrary arrow update logic with common knowledge is neither RE nor co-RE*. In Jérôme Lang, editor: *Proceedings of the 16th TARK, EPTCS* 251, pp. 373–381, doi:10.4204/EPTCS.251.27.
- [21] John-Jules Ch. Meyer & Wiebe van der Hoek (1995): *Epistemic logic for AI and computer science*. *Cambridge Tracts in Theoretical Computer Science* 41, CUP, doi:10.1017/CB09780511569852.
- [22] Joseph S. Miller & Lawrence S. Moss (2005): *The undecidability of iterated modal relativization*. *Studia Logica* 79(3), pp. 373–407, doi:10.1007/s11225-005-3612-9.
- [23] Piergiorgio Odifreddi (1989): *Classical recursion theory*. *Studies in Logic and the Foundations of Mathematics* 142, Elsevier.
- [24] Marc Pauly (2002): *A modal logic for coalitional power in games*. *Journal of Logic and Computation* 12(1), pp. 149–166, doi:10.1093/logcom/12.1.149.
- [25] Jan Plaza (1989): *Logics of public communications*. In: *Proceedings of the 4th ISMIS*, Oak Ridge National Laboratory, pp. 201–216.