The expressivity of update logics

Louwe B. Kuijer
University of Groningen, Faculty of Philosophy
Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands
l.b.kuijer@rug.nl, +31503636153

July 30, 2014

Abstract

We prove two new results about logics involving updates and common knowledge. The first result is that the logic $\mathcal{L}_{\mathrm{AU}^*}$ using Arrow Common Knowledge is more expressive than the logic $\mathcal{L}_{\mathrm{AR}}$ using Relativised Common Knowledge. The second result is that the logic $\mathcal{L}_{\mathrm{AUC}}$ using Arrow Updates and normal Common Knowledge is equally expressive as $\mathcal{L}_{\mathrm{AU}^*}$.

Together with previously known results this fully determines the expressivity landscape of all logics involving any combination of normal Common Knowledge (C), Relativized Common Knowledge (R), Arrow Common Knowledge (U*), Public Announcements (P) and Arrow Updates (U).

Keywords: Expressivity, public announcements, arrow updates, relativized common knowledge, arrow common knowledge

1 Introduction

In this paper we consider all logics that can be obtained by adding a combination of common knowledge, relativised common knowledge [12], arrow common knowledge [11], public announcements [13, 4] and arrow updates [11] to a basic modal logic.

Generally we can use only one logic at a time. So if we have multiple logics we have to choose between them. As such it becomes interesting to compare them to each other. Usually every logic has its own strengths and weaknesses so we cannot conclude that one logic is simply better than another. We can however sometimes say that one logic is better than another in one particular aspect, so with respect to some specific criterion.

There are several such criteria that can be used to compare logics. We could for example look at the computational complexity for one of the decision problems associated with the logic, or at the succinctness of the logic. Here we want to compare logics by another criterion, namely that of their *expressivity* (or *expressive power*). As the word suggests the expressivity of a logic is a measure of what can be expressed in the logic. So if a logic $\mathcal{L}_2$ is at least as expressive as a logic $\mathcal{L}_1$ then everything that can be expressed in $\mathcal{L}_1$ can also be expressed in $\mathcal{L}_2$.

What this means is that for every $\mathcal{L}_1$ formula φ_1 there is a $\mathcal{L}_2$ formula φ_2 with the same meaning; so for every $\mathcal{L}_1$ formula there is an equivalent $\mathcal{L}_2$ formula. The formulas of $\mathcal{L}_1$ can then be seen as abbreviations for the formulas of $\mathcal{L}_2$, so everything that can be done using $\mathcal{L}_1$ can also be done using $\mathcal{L}_2$. If we only look at what a logic can do (and not at how efficiently it does so) there is then no reason to use $\mathcal{L}_1$, since $\mathcal{L}_2$ does at least as well in every situation.

Let us consider a well known example. Let $\mathcal{L}_{\text{prop}}$ be propositional logic with the connectives $\neg, \vee, \wedge, \rightarrow$ and $\leftrightarrow$, and let $\mathcal{L}_{\{\neg,\vee\}}$ be propositional logic with only the connectives $\neg$ and $\vee$. It is well known that for every $\mathcal{L}_{\text{prop}}$ formula there is an equivalent $\mathcal{L}_{\{\neg,\vee\}}$ formula, so $\mathcal{L}_{\{\neg,\vee\}}$ is at least as expressive as $\mathcal{L}_{\text{prop}}$.

Here we chart the expressivity landscape of all logics under consideration, so all 32 combinations of common knowledge, relativized common knowledge, arrow common knowledge, public announcements and arrow updates. The approach we take is very similar to the one in [10], where the expressivity landscape of a different (but partially overlapping) set of logics is charted.

For many of the logics the relative expressivity is already known. There are however two important new expressivity results introduced in this paper, as well as a number of results that follow from these two results. The first is that the logic using relativised common knowledge is not as expressive as the logic using arrow common knowledge. The second is that the logic using arrow updates and normal common knowledge is equally expressive as the logic using arrow common knowledge.

1.1 Overview

In Section 2 we briefly introduce and informally discuss some properties and applications of the different operators. Then, in Section 3, we give a number of definitions that are required to compare the expressivity of the logics under consideration. In Section 4 the expressivity landscape is shown and an overview is given of both the previously known results and the new results. In Section 5 a proof is given of the first new result, that the logic using relativized common knowledge is not as expressive as the logic using arrow common knowledge. In Section 6 a proof is given of the second new result, that the logic using arrow updates and normal common knowledge is as expressive as the logic using arrow common knowledge.

2 Introducing the Operators

Multi-agent Kripke models can be used to model the information states of agents. One important property of information states is that they can change. A common way to see information change is to consider it as *changes made to the Kripke model.*[1]

One important and very general approach to information change is to use *action models* (see for example [4, 2, 3, 8]). In a logic using action models every action $[\alpha]$ is associated with an action model M_α and performing $[\alpha]$ in a model $\mathcal{M}$ changes the model to a certain submodel of the product model $\mathcal{M} \times M_\alpha$. A notable consequence of this is that applying an action may increase the size of your model. Another very general approach to information (and factual) change is to use *Global Graph Modifiers*, see [1]. Global Graph Modifiers allow one to add worlds, add or remove arrows and change the valuations of propositional variables. Because worlds and arrows can be added the use of global graph modifiers can also increase the size of your model.

Here however we focus on a particular kind of information change, where only new information is acquired (and nothing forgotten or proven false) and the new information is made publicly available. This restricted kind of information change can be described using the general Action Models or Global Graph Modifiers, but there are simpler options. New public information can only remove access to alternatives that were previously considered possible, it can never add new alternatives. This allows us to restrict ourselves to model changing operators that go from a model to one of its submodels.

A model consists of a set of possible worlds, accessibility relations between the possible worlds and the valuation of the propositional variables on the worlds. We are modeling information change, not factual change, so the values of the propositional variables should remain unchanged. Since we want to go

[1] Another approach is to consider information change as a state transition inside a larger model. The two approaches are not fundamentally different; a change from model $\mathcal{M}_1$ to model $\mathcal{M}_2$ can be seen as a state transition in a larger model containing both $\mathcal{M}_1$ and $\mathcal{M}_2$. See [5] for a discussion of dynamic epistemic logic with the dynamic operations seen as state transitions.

from a model to one of its submodels this leaves us with the choice to let the information change operator remove either worlds or accessibility arrows.

The most commonly used choice is to remove certain possible worlds using *public announcements* (see for example [13, 4, 8]). A public announcement $[\varphi]$ removes all worlds where the formula φ does not hold from the model. A less commonly used alternative is to remove certain accessibility arrows. A very simple version of this is the variation on public announcements in [9, 10] where $[\varphi]$ does not remove the worlds where φ does not hold but merely the arrows to such worlds. Removing arrows in this way has the same result as removing the $\neg\varphi$ worlds, an inaccessible world might as well not exist.

A more powerful way to remove arrows is to use *arrow updates*, see [11]. An arrow update $[U]$ consists of a number of clauses, $U = \{(u_1, a_1, u'_1), \cdots, (u_n, a_n, u'_n)\}$, where we do not require that $a_i \neq a_j$ whenever $i \neq j$. An arrow satisfies a clause (u_i, a_i, u'_i) iff it is an arrow for agent a_i and it goes from a world that satisfies the start condition u_1 and to a world that satisfies the end condition u'_i.[2] The update removes those arrows that satisfy none of the clauses.[3]

Another operator that is often used in logics about information is the common knowledge operator C_B, where B is a group of agents. The formula $C_B\varphi$ holds in a world w iff φ holds in all worlds w' that are reachable from w by a "B-path" (that is, a sequence of arrows belonging to agents in B that connect w to a successor w_1 of w, w_1 to a successor w_2 or w_1 and so up to an arrow that connects w_n to a successor w' of w_n).

For both public announcements and arrow updates there is an associated variant of common knowledge. The common knowledge variant for public announcements is *relativized common knowledge*, defined in [12]. The formula $C_B(\varphi_1, \varphi_2)$ stands for φ_2 being common knowledge relative to φ_1. It holds in a world w iff φ_2 holds in all worlds w' that are reachable from w by a B-path *that only consists of φ_1 worlds.*

The corresponding variant of common knowledge for arrow updates is *arrow common knowledge*, defined in [11]. The formula $\{U\}^*\varphi$ stands for φ being common knowledge relative to the arrow update U. It holds in a world w iff φ holds in all worlds w' that are reachable from w by a path *that only consists of arrows that satisfy U.*[4]

[2] We use the slightly awkward terms "start condition" and "end condition", as opposed to "precondition" and "postcondition", in order to prevent confusion with the preconditions in action models.

[3] The operator $[U]$ has some similarities to the operator $[a_1 - (\varphi_1, \psi_1); \cdots; a_n - (\varphi_n, \psi_n)]$ from [1], but there are also two differences. The first difference is that $[a_1 - (\varphi_1, \psi_1); \cdots; a_n - (\varphi_n, \psi_n)]$ specifies the arrows that are to be removed whereas $[U]$ specifies the arrows that are to be retained. The second and more important difference is that the clauses from $[U]$ are considered *simultaneously*, while the clauses $a_1 - (\varphi_1, \psi_1); \cdots; a_n - (\varphi_n, \psi_n)$ are considered *sequentially*.

[4] The $*$ in $\{U\}^*$ is a Kleene star; the operator $\{U\}^*$ is an iterated version of $\{U\}$, which is defined by $\mathcal{M}, w \models \{U\}\varphi$ iff $\mathcal{M}, w' \models \varphi$ for all worlds that are accessible from w by an arrow satisfying a clause from U. Unfortunately $\{U\}$ is visually too similar to $\{U\}^*$ to be practical, we will write $\Box_U$ instead of $\{U\}$. We should be careful to distinguish $\{U\}^*$ from the operator $[U]^*$ (that is not used in this paper) that repeatedly applies an update.

The different building blocks discussed so far can be combined in different ways. We could for example define a logic $\mathcal{L}_{\mathrm{APU}^*}$ that uses basic modal logic together with public announcements and arrow common knowledge, or a logic $\mathcal{L}_{\mathrm{AR}}$ that uses basic modal logic together with relativized common knowledge.

2.1 Dynamic and Static Operators

The operators under consideration here can be split into two different kinds: dynamic operators and static operators. The difference between the two kinds is that dynamic operators change the model when they are interpreted while static operators do not.

The dynamic operators used here are public announcements and arrow updates, the static operators are the three types of common knowledge as well as the Boolean operators and $\Box_a$ of basic modal logic. Two of the static operators are combinations of common knowledge with a dynamic operator. It is worthwhile to spend a few moments to see what it means for a static operator to be related to a dynamic operator in such a way.

Let us start by considering a logic where we have the dynamic operators, the Boolean operators, the modal $\Box_a$ and a normal common knowledge operator C_B but not the two other common knowledge operators. Suppose that in this logic we use one of the dynamic operators, say a public announcement $[\varphi]$ in a pointed model $\mathcal{M}, w$. This announcement removes all $\neg\varphi$ worlds from $\mathcal{M}$, a process that cannot be undone. This means that in the updated model $\mathcal{M}_{[\varphi]}$ some of the information contained in the model $\mathcal{M}$ is lost. In particular we generally cannot determine from $\mathcal{M}_{[\varphi]}, w$ whether or not $\Box_a\varphi$ held before the update, so whether or not $\mathcal{M}, w \models \Box_a\varphi$. Arrow updates destroy information in the same way.

But occasionally we want to use something similar to the dynamic modalities but that does not destroy information. This can be done by adding static operators that correspond to the combination of a dynamic operator and another operator. Such a new static operator applies its update, performs its associated operation and then un-applies the update. Or, to put it another way, it temporarily *pretends* to apply a dynamic operator.

The static connective $\Box_\varphi$ is the combination of a $\Box$ operator and a public announcement. It first applies the announcement φ, then takes a step in the updated model with $\Box$ and finally it undoes the update. So we have $\mathcal{M}, w \models \Box_\varphi\psi$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are accessible from w in the updated model $\mathcal{M}_{[\varphi]}$. Likewise, $\mathcal{M}, w \models \Box_U\psi$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are accessible from w in the updated model $\mathcal{M}_{[U]}$.

The operators $\Box_\varphi$ and $\Box_U$ do not add expressivity, however. We have $\mathcal{M}, w \models \Box_\varphi\psi$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are accessible from w in the updated model $\mathcal{M}_{[\varphi]}$, so if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are accessible from w in $\mathcal{M}$ that satisfy φ, so if and only if $\mathcal{M}, w \models \Box(\varphi \to \psi)$. Formulating a formula equivalent to $\Box_U\psi$ is harder but it can also be done; we have $\mathcal{M}, w \models \Box_U\psi$ if and only if $\mathcal{M}, w \models \bigwedge_{(u_1,a,u_2)\in U}(u_1 \to \Box(u_2 \to \psi))$. That the operators $\Box_\varphi$ and $\Box_U$ do not add expressivity means they would not add anything fundamentally new to the logic. This does not mean that they are useless; the operator $\Box_U$ is in fact used in several of the proofs in this paper. But there is no need to take them as primitive, they can be seen as abbreviations.

Things get more complicated if we combine the dynamic operators not with $\Box$ but with common knowledge. Earlier we defined relativized common knowledge $C_B(\varphi, \psi)$ as meaning "ψ holds in all worlds that are reachable by a B-path that contains only φ worlds". Note that this is equivalent to saying that $\mathcal{M}, w \models C_B(\varphi, \psi)$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are reachable from w by a B-path in the updated model $\mathcal{M}_{[\varphi]}$. So relativized common knowledge is indeed the static operator corresponding to the combination of common knowledge and a public announcement.

Likewise, we defined $\{U\}^*\psi$ as meaning "ψ holds in all worlds that are reachable by a path that only uses arrows that satisfy U". This is equivalent to saying that $\mathcal{M}, w \models \{U\}^*\psi$ if and only if $\mathcal{M}, w' \models \psi$ for all worlds w' that are reachable from w by a path in the updated model $\mathcal{M}_{[U]}$. So arrow common knowledge is, as the name suggests, the static operator corresponding to the combination of common knowledge and an arrow update.

The relativized common knowledge and arrow common knowledge operators are both rather complicated, and they were introduced mainly for technical reasons. Still, as [6] points out there is an informal reading of relativized common knowledge that, while not simple, can provide some intuition behind the operator. A formula $[\varphi]C_B\psi$, which contains a public announcement and a normal common knowledge formula, can be read as "if ψ is announced then it will become common knowledge (among B) that ψ is the case". A formula $C_B(\varphi, \psi)$ on the other hand can be read as "if φ is announced then it will become common knowledge (among B) that ψ *used to be the case* before the announcement". Likewise, $[U]C_B\psi$ can be read as "if U is announced it will become common knowledge (among B) that ψ is the case" whereas $\{U\}^*\psi$ can be read as "if U is announced it will become common knowledge (among B) that ψ *used to be the case* before the announcement".

Unlike $\Box_\varphi$ and $\Box_U$ there is no obvious way to express $C_B(\varphi, \psi)$ and $\{U\}^*\psi$ without using one of the new static operators. In fact, in [6] it was shown that the logic relativized common knowledge adds expressivity to a logic with $\wedge, \neg, \Box_a, C_B$ and $[\varphi]$ operators. So relativized common knowledge adds something fundamentally new to such a logic.

Arrow common knowledge is to common knowledge and arrow updates as relativized common knowledge is to common knowledge and public announcements. As such the result in [6] suggested that arrow common knowledge would probably add expressivity to a logic with $\wedge, \neg, \Box_a, C_B$ and $[U]$ operators. Here we prove that, surprisingly, this is not the case; for any formula using $\wedge, \neg, \Box_a, C_B, [U]$ and $\{U\}^*$ there is an equivalent formula using only $\wedge, \neg, \Box_a, C_B$ and $[U]$. This means that $\{U\}^*$ does not add anything fundamentally new to such a logic and that it could in theory be used as an abbreviation. It is not very practical to consider $\{U\}^*$ in this way however, as the translation from a formula with $\{U\}^*$ to one without is extremely complicated and causes an enormous increase in formula size.

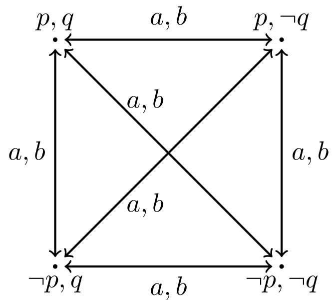

(a) $\mathcal{M}$

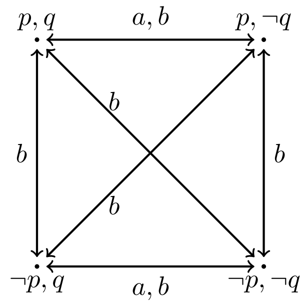

(b) $\mathcal{M}_{[U]}$

Figure 1: A model $\mathcal{M}$ representing a simple game and a model $\mathcal{M}_{[U]}$ representing the same game after a has looked at her card. Reflexive arrows are omitted.

2.2 Public announcements and arrow updates

Public announcements are quite widely used so we assume that the reader has encountered them before. Arrow updates on the other hand are not very commonly used so it seems worthwhile to give a short introduction to arrow updates, and especially the difference between arrow updates and public announcements.

The first thing to note is that everything that can be done using public announcements can also be done using arrow updates. In other words, arrow updates are at least as expressive as public announcements. To see why this is the case consider any public announcement $[\varphi]$. This announcement removes all worlds that do not satisfy φ. With arrow updates we cannot remove any worlds, but we can do something with the same effect: we can remove all arrows to $\neg\varphi$ worlds. A world that is not reachable in any way might as well not exist at all, so this has the same effect as removing all $\neg\varphi$ worlds.

But arrow updates can also be used in situations where public announcements cannot. Let us look at a simple example, loosely based on an example given in [11] (which was in turn based on an example in [7]).

Example 1. In a very simple card game there are two players, player a and player b. Both players are dealt a single card, face down. Player a either has the ace of spades (p) or the king of spades ($\neg p$), player b either has the ace of diamonds (q) or the king of diamonds ($\neg q$). At this point neither player knows which card either one of them holds. The situation as described so far can be modeled as shown in Figure 1a.

But then suppose that a (openly) looks at her card without showing it to b. This action is public, because a openly looks at her card. But it still creates some private information for a, namely which card she holds. This private information makes it impossible to model the event using a public announcement. We can however model it quite simply using an arrow update $[U]$.

In every p world agent a learns that p is true, so in those worlds she no longer holds $\neg p$ worlds possible. Likewise, in every $\neg p$ world she no longer holds p worlds possible. This means that a-arrows should only be retained if they go from a p world to a p world or from a $\neg p$ world to a $\neg p$ world. We can do this by including the clauses (p, a, p) and $(\neg p, a, \neg p)$ to U. Agent b on the other hand learns no new information that would allow him to distinguish between worlds that he could not previously distinguish, so all b-arrows should be retained. We can do this by including a $(\top, b, \top)$ clause in U.

In the end this gives us the update $U = (\top, b, \top), (p, a, p), (\neg p, a, \neg p)$. And indeed, if applied to $\mathcal{M}$ this gives us the model $\mathcal{M}_{[U]}$, shown in Figure 1b, which is a faithful representation of the game after a has looked at her card.

The most important property of the update in the above example is that the information a learns differs per world. In p worlds a learns that she holds the ace, while in $\neg p$ worlds she learns that she holds the king. This world-dependence makes it impossible to fully eliminate either the p worlds or the $\neg p$ worlds so public announcements cannot model the new information. Arrow updates on the other hand can model the new information just fine, by removing some (but not all) arrows between p and $\neg p$ worlds.

3 Definitions

Let us now define the different logics that we want to compare. In order to compare the expressivity of the different kinds of updates and common knowledges it is convenient to first define a logic $\mathcal{L}_\mathcal{T}$ that contains all the logics we consider. We can then compare the logics as fragments of $\mathcal{L}_\mathcal{T}$. The advantage of doing this is that it allows us to combine formulas from the different logics. For example, $[U]C_\mathcal{A}\varphi \leftrightarrow \{U\}^*[U]\varphi$ is only a well formed formula if we have one logic that contains all of the connectives $[U], C_\mathcal{A}, \leftrightarrow$ and $\{U\}^*$. We do have such a logic, namely $\mathcal{L}_\mathcal{T}$.

Let $\mathcal{A}$ be a finite nonempty set of agents and $\mathcal{P}$ a countable set of propositional variables.

Definition 1 (The language of $\mathcal{L}_\mathcal{T}$)**.** *The* formulas *of $\mathcal{L}_\mathcal{T}$ are given by*

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \Box_a\varphi \mid C_B\varphi \mid C_B(\varphi, \varphi) \mid [\varphi]\varphi \mid [U]\varphi \mid \{U\}^*\varphi$$
$$U ::= (\varphi, a, \varphi) \mid (\varphi, a, \varphi), U$$

where $p \in \mathcal{P}$, $B \subseteq \mathcal{A}$ and $a \in \mathcal{A}$. Let $\Phi_\mathcal{T}$ be the set of formulas of $\mathcal{L}_\mathcal{T}$.

We use $\wedge, \bigvee, \bigwedge, \rightarrow, \leftrightarrow, \top, \bot$ and $\Diamond_a$ in the usual way as abbreviations, omit parenthesis where this should not cause confusion and write a for $\{a\}$. We also slightly abuse notation by identifying an update $U = (u_1, a_1, u_1'), \cdots, (u_k, a_k, u_k')$ with the set $U = \{(u_1, a_1, u_1'), \cdots, (u_k, a_k, u_k')\}$. Furthermore, if $B \subseteq \mathcal{A}$ we write $\Box_B\varphi$ for $\bigwedge_{a\in B}\Box_a\varphi$ and $(\varphi_1, B, \varphi_2)$ for $\{(\varphi_1, a, \varphi_2) \mid a \in B\}$. Finally, we write $\Box$ for $\Box_\mathcal{A}$.

The models for $\mathcal{L}_\mathcal{T}$ are the standard Kripke models. It should be noted that although we speak of (common) *knowledge* we do not assume any of the frame conditions usually associated with epistemic logic.

Definition 2 (The models of $\mathcal{L}_\mathcal{T}$)**.** *An* $\mathcal{L}_\mathcal{T}$ model $\mathcal{M}$ *is a triple* $\mathcal{M} = (W, R, v)$ *where* W *is a set of worlds,* $R : \mathcal{A} \to \wp(W \times W)$ *assigns to each agent an accessibility relation on* W *and* $v : \mathcal{P} \to \wp(W)$ *is a valuation that assigns to each propositional variable a subset of the worlds.*

We say that w is a world of $\mathcal{M} = (W, R, v)$ iff $w \in W$. We can now define the semantics of $\mathcal{L}_\mathcal{T}$.

Definition 3 (The semantics of $\mathcal{L}_\mathcal{T}$)**.** *Given an* $\mathcal{L}_\mathcal{T}$ *model* $\mathcal{M} = (W, R, v)$*, a world* w *of* $\mathcal{M}$ *and* φ, ψ *formulas of* $\mathcal{L}_\mathcal{T}$ *define the* satisfaction relation $\models$ *by*

$$
\begin{array}{lll}
\mathcal{M}, w \models p & \textit{if} & w \in v(p), \\
\mathcal{M}, w \models \neg\varphi & \textit{if} & \mathcal{M}, w \not\models \varphi, \\
\mathcal{M}, w \models \varphi \vee \psi & \textit{if} & \mathcal{M}, w \models \varphi \textit{ or } \mathcal{M}, w \models \psi, \\
\mathcal{M}, w \models \Box_a \varphi & \textit{if} & \mathcal{M}, w' \models \varphi \textit{ for all } w' \textit{ such that } (w, w') \in R(a), \\
\mathcal{M}, w \models [\psi]\varphi & \textit{if} & \mathcal{M}, w \models \psi \textit{ implies } \mathcal{M}_{[\psi]}, w \models \varphi, \\
\mathcal{M}, w \models [U]\varphi & \textit{if} & \mathcal{M}_{[U]}, w \models \varphi, \\
\mathcal{M}, w \models C_B \varphi & \textit{if} & \mathcal{M}, w' \models \varphi \textit{ for all } w' \textit{ such that } (w, w') \in R(B)^*, \\
\mathcal{M}, w \models C_B(\psi, \varphi) & \textit{if} & \mathcal{M}, w' \models \varphi \textit{ for all } w' \textit{ such that } (w, w') \in R_{[\psi]}(B)^*, \\
\mathcal{M}, w \models \{U\}^* \varphi & \textit{if} & \mathcal{M}, w' \models \varphi \textit{ for all } w' \textit{ such that } (w, w') \in R^*_{[U]}
\end{array}
$$

where

- $W_{[\varphi]} = \{w \in W \mid \mathcal{M}, w \models \varphi\}$,
- $R_{[\varphi]}(a) = R(a) \cap (W_{[\varphi]} \times W_{[\varphi]})$ *for* $a \in \mathcal{A}$,
- $v_{[\varphi]}(p) = v(p) \cap W_{[\varphi]}$ *for* $p \in \mathcal{P}$,
- $\mathcal{M}_{[\varphi]} = (W_{[\varphi]}, R_{[\varphi]}, v_{[\varphi]})$,
- $R_{[U]}(a) = \{(w_1, w_2) \in R(a) \mid \exists (u, a, u') \in U : \mathcal{M}, w_1 \models u$ *and* $\mathcal{M}, w_2 \models u'\}$ *for* $a \in \mathcal{A}$,
- $\mathcal{M}_{[U]} = (W, R_{[U]}, v)$,
- $R(B)^*$ *is the reflexive transitive closure of* $\bigcup_{a \in B} R(a)$,
- $R_{[\varphi]}(B)^*$ *is the reflexive transitive closure of* $\bigcup_{a \in B} R_{[\varphi]}(a)$,
- $R^*_{[U]}$ *is the reflexive transitive closure of* $\bigcup_{a \in \mathcal{A}} R_{[U]}(a)$.

We write $\mathcal{M} \models \varphi$ *if* $\mathcal{M}, w \models \varphi$ *for every world* w *of* $\mathcal{M}$ *and* $\models \varphi$ *if* $\mathcal{M} \models \varphi$ *for every model* $\mathcal{M}$.

Most of the semantics are as usual, although there are a two things worth pointing out. The first is that the common knowledge operators take the *reflexive* transitive closure of the relevant relation. This is not very unusual, but neither is taking the transitive closure instead. The second is that a public announcement formula $[\psi]\varphi$ is automatically true in every world $\neg\psi$ world. Again, this is not unusual but there are other options. The results presented in this paper also hold for the alternative semantics, only some very minor changes in the proofs would be required.

What we are really interested in is not $\mathcal{L}_{\mathcal{T}}$ but certain fragments of it. We define these fragments as in [10].

Definition 4 (Fragments of $\mathcal{L}_{\mathcal{T}}$)**.** *Let*

- *A, representing 'agents', stand for the connectives* $\neg, \vee$ *and* $\Box_a$
- *C, representing 'common knowledge', stand for the connective* C_B
- *R, representing 'relativised common knowledge', stand for the connective* $C_B(\psi, \varphi)$
- *P, representing 'public announcement', stand for the connective* $[\varphi]$
- *U, representing 'arrow updates', stand for the connective* $[U]$
- U^**, representing 'arrow common knowledge', stand for the connective* $\{U\}^*$*.*

The logic $\mathcal{L}_X$ *for a finite string X is the logic* $\mathcal{L}_{\mathcal{T}}$ *with the language restricted to only those connectives that belong to a letter in X. Let* Φ_X *be the set of formula of* $\mathcal{L}_X$*.*

So for example the logic $\mathcal{L}_{\mathrm{AUC}}$ is the logic using the connectives $\neg, \vee, \Box_a, C_B$ and $[U]$. We also sometimes denote the logic $\mathcal{L}_{\mathrm{X}}$ by the string X, so AUC is the logic $\mathcal{L}_{\mathrm{AUC}}$. We can easily define the relative expressivity of such fragments.

We write $\models$ for the satisfaction relation of the fragments as well as for the satisfaction relation of $\mathcal{L}_{\mathcal{T}}$. There is no risk of confusion as the different satisfaction relations coincide whenever multiple ones are defined.

Definition 5. *Let* $\mathcal{L}_1$ *and* $\mathcal{L}_2$ *be fragments of* $\mathcal{L}_{\mathcal{T}}$*. Then* $\mathcal{L}_2$ *is* at least as expressive as $\mathcal{L}_1$*, denoted* $\mathcal{L}_1 \preceq \mathcal{L}_2$*, if for each* $\mathcal{L}_1$ *formula* φ_1 *there is an* $\mathcal{L}_2$ *formula* φ_2 *such that*

$$\models \varphi_1 \leftrightarrow \varphi_2.$$

We say that $\mathcal{L}_2$ *is* more expressive than $\mathcal{L}_1$*, denoted* $\mathcal{L}_1 \prec \mathcal{L}_2$*, if* $\mathcal{L}_1 \preceq \mathcal{L}_2$ *and* $\mathcal{L}_2 \not\preceq \mathcal{L}_1$*. We say that* $\mathcal{L}_2$ *and* $\mathcal{L}_1$ *are* equally expressive*, denoted* $\mathcal{L}_1 \equiv \mathcal{L}_2$*, if* $\mathcal{L}_1 \preceq \mathcal{L}_2$ *and* $\mathcal{L}_2 \preceq \mathcal{L}_1$*.*

We can coherently write $\models \varphi_1 \leftrightarrow \varphi_2$ even though φ_1 and φ_2 are in different logics because both logics are fragments of $\mathcal{L}_{\mathcal{T}}$. Note that the relation $\preceq$ is reflexive and transitive. The relation $\equiv$ inherits the reflexivity and transitivity of $\preceq$ and is also symmetric, so it is an equivalence relation.

It will also be useful to define the depth of a formula.

Definition 6. *For* $\varphi \in \Phi_{\mathcal{T}}$ *define the* depth $d(\varphi)$ *of* φ *recursively by*

- $d(p) = 0$ *for* $p \in \mathcal{P}$*,*
- $d(\neg\varphi_1) = d(\varphi_1)$*,*

- $d(\varphi_1 \vee \varphi_2) = \max(d(\varphi_1), d(\varphi_2))$,
- $d(\Box_a \varphi_1) = d(C_B \varphi_1) = d(\varphi_1) + 1$,
- $d(C_B(\varphi_1, \varphi_2)) = d([\varphi_1]\varphi_2) = \max(d(\varphi_1), d(\varphi_2)) + 1$,
- $d([U]\varphi_1) = d(\{U\}^* \varphi_1) = \max(d(\varphi_1), d(U)) + 1$,
- $d(U) = \max_{(u,a,u') \in U}(d(u), d(u') + 1)$.

We say that φ if of pure depth *n if $d(\varphi) = n$ and there is no subformula φ' of φ such that $d(\varphi') = n$.*

The only clause that may be somewhat surprising is that of $d(U)$. The reason for adding an extra $+1$ to the depth of end conditions is that they are evaluated in the next world, and thus reach one world further than a start condition of the same depth.

The concept of pure depth is useful to restrict the number of possibilities for the form of a formula; an AC formula of pure depth 1 for example must be either of the form $\Box_a \varphi'$ or of the form $C_B \varphi'$. The formulas of depth n are the Boolean combinations of the formulas of pure depth at most n.

Finally, it is useful to have a flexible definition of a path.

Definition 7. *Given an $\mathcal{L}_\mathcal{T}$ model $\mathcal{M} = (W, R, v)$ and two worlds w_1 and w_n of $\mathcal{M}$ a path π from w_1 to w_n is an ordered set of triples*

$$\pi = ((w_1, a_1, w_2), (w_2, a_2, w_3), \cdots, (w_{n-1}, a_{n-1}, w_n))$$

where $n \in \mathbb{N}$, and $a_i \in \mathcal{A}$ and $(w_i, w_{i+1}) \in R(a_i)$ for $1 \leq i \leq n-1$.

Let $B \subseteq \mathcal{A}$, φ a formula and U an update. The path π is a B-path if $a_i \in B$ for $1 \leq i \leq n-1$, a φ-path if $\mathcal{M}, w_i \models \varphi$ for $1 \leq i \leq n$ and a U-path if for $1 \leq i \leq n-1$ there is a clause $(u, a, u') \in U$ such that $\mathcal{M}, w_i \models u$, $a = a_i$ and $\mathcal{M}, w_{i+1} \models u'$. Conditions can be combined, π is an $(X_1, \cdots, X_k)$-path if it is an X_j-path for all $1 \leq j \leq k$.

4 The logics under consideration

Using different combinations of A, C, R, P, U and U* we could define $2^6 = 64$ different fragments of $\mathcal{L}_\mathcal{T}$. Not all these fragments are interesting, however.

In this paper we do not consider logics that do not have the A connectives. This is not because logics without some or all of the A connectives are guaranteed to be uninteresting; logics with one of the common knowledge operators but not the $\Box_a$ operator are for example comparable to temporal logics with “future” but not “next”. But such logics without A are outside the scope of this paper and left for further work.

This leaves us with 32 fragments. There are however a few easy reductions of some connectives to other ones that allow us to reduce that number further.

Lemma 1. *For any $\mathcal{L}_{\mathcal{T}}$ formulas φ, ψ and any $B \subseteq \mathcal{A}$ we have $\models [\psi]\varphi \leftrightarrow (\psi \to [(\psi, \mathcal{A}, \psi)]\varphi)$, $\models C_B\varphi \leftrightarrow C_B(\top, \varphi)$ and $\models C_B(\psi, \varphi) \leftrightarrow \{(\psi, B, \psi)\}^*\varphi$.*

The proof should be immediately clear and is left to the reader. Lemma 1 allows us to restrict ourselves to fragments having at most one of the update connectives $[\varphi]$ or $[U]$ and at most one of the common knowledge connectives C_B, $C_B(\cdot,\cdot)$ and $\{U\}^*$; if more than one of these connectives occurs in a logic only the 'strongest' one is relevant.[5]

This leaves 12 logics that can be ordered two dimensionally, with the update connective (if any) on one axis and the common knowledge connective (if any) on the other. The logics and their relative expressivities are shown in Figure 2.

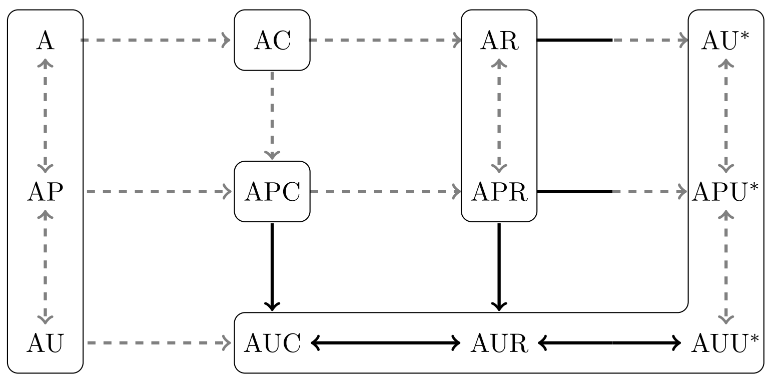

Figure 2: The landscape of logics using basic modal logic (A) and a combination of public announcements (P), arrow updates (U), common knowledge (C), relativised common knowledge (R) and arrow update common knowledge (U*). Arrows X $\longrightarrow$ Y indicate that X $\prec$ Y. Double arrows X $\longleftrightarrow$ Y indicate that X $\equiv$ Y. Dashed gray arrows indicate previously established results, solid black arrows indicate new results. For arrows that are part gray and dashed and part black and solid the result in one direction was previously established but the result in the other direction is new. Boundaries around nodes indicate equivalence classes of logics that are equally expressive. For reasons of clarity not all arrows are drawn, but the omitted arrows all follow from the drawn ones by transitivity.

Note that although not all arrows are drawn the arrows in Figure 2 are sufficient to know the relative expressivity of any of the logics by transitivity and reflexivity of $\preceq$.

4.1 Overview of previously known results

The arrows in Figure 2 that are drawn dashed and in gray were previously known. That A $\equiv$ AP was shown in [13]. That AC $\prec$ APC was shown in [4]. That AR $\equiv$ APR was shown in [12]. In [6] it was shown that APC $\prec$ AR,

[5] Note that among other things this implies that $\mathcal{L}_{\mathcal{T}} \equiv \mathcal{L}_{\mathrm{AUU}^*}$.

which implies that APC $\prec$ APR. Finally, in [11] it was shown that AU $\equiv$ AP, AUU* $\equiv$ AU* and AR $\preceq$ AU*. The remaining dashed arrows in Figure 2 are either trivial or follow from other dashed arrows by transitivity.

4.2 New expressivity results

The arrows that are drawn solid and in black in Figure 2 are new results. There are ten such new results, each corresponding to one half of an arrow. They all follow by transitivity from two results, however.

The first result is that AU* $\not\preceq$ AR. This result is proven in Section 5. This result is not very surprising, it was already predicted in [11]. The second result is that AU* $\preceq$ AUC. This result is proven in Section 6. Unlike the previous result this result is rather surprising, considering that the difference between AUC and AU* is very similar to the difference between APC and AR and we have AR $\not\preceq$ APC.

The proof that AU* $\preceq$ AUC is very complicated and technical. We therefore only give an overview of the proof in this paper itself and provide the details as supplementary data. But before looking at even an overview of the proof it might be worthwhile to consider what the difference between AUC and APC is that causes APC to be less expressive than AR while AUC is as expressive as AU*. This difference is the ability to *store information*.

Recall that the informal readings of the operators $C_B(\varphi, \psi)$ and $\{U\}^*\psi$ look to the past. The formula $C_B(\varphi, \psi)$ holds if, after announcing $[\varphi]$, it will be common knowledge that ψ *used to* hold before the announcement. Likewise, $\{U\}^*\psi$ holds if, after announcing $[U]$, it will be common knowledge that ψ *used to* hold before the announcement. So if we want to simulate $C_B(\varphi, \psi)$ in APC or $\{U\}^*\psi$ in AUC we have to find a way to store information about the current model in such a way that it is not destroyed by the update. In APC we cannot do this, so AR is more expressive than APC. But arrow updates are more powerful than public announcements, and this extra power allows AUC to store just enough information to simulate $\{U\}^*\psi$. The ways in which AUC stores information are many and varied, but let us briefly consider one of them.

Suppose we want to apply an update $U = (u_1, a_1, u_1'), \cdots, (u_n, a_n, u_n')$ but first store information about which worlds satisfy a given formula φ. Then instead of updating with $[U]$ we could update with a $[U'] = [(u_1 \wedge \neg\varphi, a_1, u_1'), \cdots, (u_n \wedge \neg\varphi, a_n, u_n')]$. After this update every world that *used to* satisfy φ now satisfies $\Box\bot$, allowing us to recognize it. Of course there are many complications to this method. In particular, there might be worlds that satisfy $\neg\varphi \wedge [U]\Box\bot$. Such worlds will satisfy $\neg\varphi \wedge [U']\Box\bot$, so they will be false positives. But, through a lot of technical work, we can exclude that possibility. The important thing to note is that there is no similar way to store information using APC: there is no public announcement $[\alpha]$ that guarantees that every φ world satisfies $\neg[\alpha]\bot \wedge [\alpha]\Box\bot$ because there is no formula α that identifies the worlds that are accessible from a φ world (as opposed to the $\Diamond\varphi$ worlds from which a φ world is accessible).

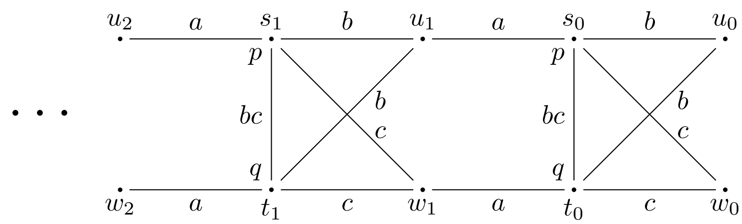

Figure 3: The model $\mathcal{M}_c$. This is an S5 models so every world has a reflexive arrow for every agent, but these are not drawn for reasons of clarity. For every $n \in \mathbb{N}$ the worlds s_n and u_{n+1} are connected by a, as are t_n and w_{n+1}.

5 AU* is more expressive than AR

The proof we give of the fact that AU* is more expressive than AR is very similar to the usual proofs of such results. We want to show that AU* $\not\leq$ AR, so that there are AU* formulas for which there is no equivalent AR formula. The most straightforward way to do this is to construct an AU* formula ξ and show that there is no AR formula equivalent to it.

At this point we would like to proceed by constructing a model $\mathcal{M}_c$ with two points w^+ and w^- such that ξ distinguishes between $\mathcal{M}_c, w^+$ and $\mathcal{M}_c, w^-$ even though there is no AR formula that distinguishes between the two worlds. This would be sufficient to show that there is no AR formula equivalent to ξ. Unfortunately this is too hard, we cannot find such a model. So we do something slightly more complicated.

Instead of creating a model where ξ distinguishes between two particular worlds w^+ and w^- we construct a model where ξ distinguishes between two worlds u_{2i} and u_{2i+1} for every $i \in \mathbb{N}$. For each i there will be some AR formula φ_i that distinguishes between u_{2i} and u_{2i+1}, but we will show that for such φ_i it must hold that $d(\varphi_i) > i$. Every AR formula has a fixed and finite depth, so while for each i there is an AR formula φ_i that distinguishes u_{2i} from u_{2i+1} there is no AR formula φ_∞ that distinguishes u_{2i} from u_{2i+1} for every $i \in \mathbb{N}$. This shows that there is no AR formula equivalent to the AU* formula ξ.

Now let us construct the model $\mathcal{M}_c$ and the formula ξ. Let $\mathcal{M}_c$ be the model shown in Figure 3. Furthermore, let $\xi = \{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}^*(\Diamond_a(p \vee q) \vee \Diamond_b(p \vee q))$. Note that the accessibility relations on $\mathcal{M}_c$ are reflexive, transitive and Euclidean so $\mathcal{M}_c$ is an S5 model.

Lemma 2. *For every $i \in \mathbb{N}$ we have $\mathcal{M}_c, u_{2i} \models \xi$ and $\mathcal{M}_c, u_{2i+1} \not\models \xi$.*

Proof. First let us look at the subformula $\Diamond_a(p \vee q) \vee \Diamond_b(p \vee q)$ of ξ. Every world other than u_0 and w_0 satisfies $\Diamond_a(p \vee q)$. Furthermore, u_0 satisfies $\Diamond_b(p \vee q)$. So w_0 is the only world in $\mathcal{M}_c$ that does not satisfy $\Diamond_a(p \vee q) \vee \Diamond_b(p \vee q)$. As a result, any world in $\mathcal{M}_c$ satisfies ξ if and only if w_0 is not reachable from that world by a $\{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}$-path.

Now let us look at the set $\{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}$ of clauses.

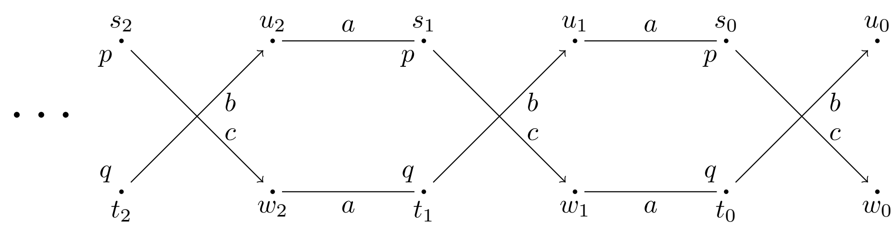

Figure 4: The arrows in $\mathcal{M}_c$ that satisfy $\{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}$. Reflexive arrows are omitted.

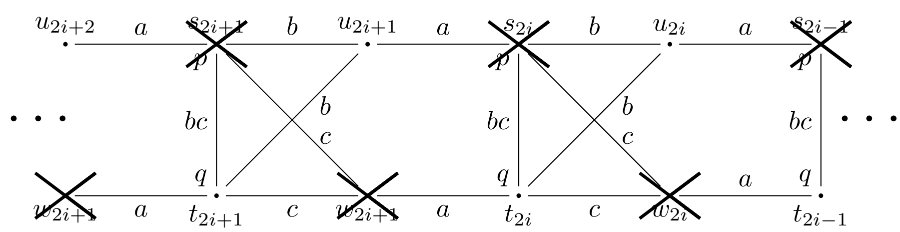

Figure 5: An illustration of why $C_B(\varphi_1, \varphi_2)$ cannot distinguish between u_j and u_{j+1} if w_j and s_j do not satisfy φ_1. Worlds not satisfying φ_1 are crossed out.

The clause $(\top, a, \top)$ is satisfied by all a arrows. Most b- and c-arrows do not satisfy any of the clauses however; only the b-arrows from t_i to u_i and the c-arrows from s_i to w_i satisfy $(q, b, \neg p \wedge \neg q)$ and $(p, c, \neg p \wedge \neg q)$ respectively. The arrows in $\mathcal{M}_c$ that satisfy any of the clauses are therefore as shown in Figure 4. The world w_0 is reachable from u_j by a $\{(\top, a, \top), (p, c, \neg p \wedge \neg q), (q, b, \neg p \wedge \neg q)\}$-path if and only if j is odd. This implies that $\mathcal{M}_c, u_{2i} \models \xi$ and $\mathcal{M}_c, u_{2i+1} \not\models \xi$ for every $i \in \mathbb{N}$. $\square$

Now we should show that no AR formula φ of depth $n := d(\varphi)$ can distinguish between u_{2i} and u_{2i+1} for any $i \geq n$. In order to make the induction step in the proof proceed smoothly it is convenient to prove a slightly stronger lemma.

Lemma 3. *Take any $n \in \mathbb{N}$, any $i \geq n$, any $x \in \{s, t, u, w\}$ and any $\varphi \in \Phi_{\text{AR}}$ such that $d(\varphi) \leq n$. Then $\mathcal{M}_c, x_{2i} \models \varphi$ if and only if $\mathcal{M}_c, x_{2i+1} \models \varphi$.*

Proof. By induction on n. As base case suppose that $n = 0$. Then φ is a Boolean combination of propositional variables. The worlds x_{2i} and x_{2i+1} have the same values for all propositional variables so this φ cannot distinguish between them.

Suppose then as induction hypothesis that $n > 0$ and that the lemma holds for all $n' < n$. The proof for $x = s$ is completely analogous to the proof for $x = t$ and the proof for $x = u$ is likewise completely analogous to the proof for $x = w$. We therefore omit the proofs for $x \in \{t, w\}$ and only show that the lemma holds for $x \in \{s, u\}$.

Suppose towards a contradiction that φ distinguishes between x_{2i} and x_{2i+1}. If a Boolean combination of formulas distinguishes between two worlds then so does at least one of the combined formulas. We can therefore assume without loss of generality that φ if of pure depth. Furthermore, if $d(\varphi) < n$ it follows immediately from the induction hypothesis that φ does not distinguish between x_{2i} and x_{2i+1} so we can also assume that $d(\varphi) = n$.

So φ is of pure depth $n > 0$, which implies that it must be of the form $\Box_\alpha \varphi_1$ or $C_B(\varphi_1, \varphi_2)$ for some $\alpha \in \mathcal{A}$, $B \subseteq \mathcal{A}$ and $\varphi_1, \varphi_2 \in \Phi_{\text{AR}}$ such that $d(\varphi_1), d(\varphi_2) < n$.

First let us suppose that $\varphi = \Box_\alpha \varphi_1$. We have $2i \geq 2n > 0$ so for every arrow from x_{2i} to y_k (with $y \in \{s, t, u, w\}$ and $2i + 1 \geq k \geq 2i - 1$) there is a corresponding arrow from x_{2i+1} to y_{k+1} and vice versa. This implies that in order for φ to distinguish between x_{2i} and x_{2i+1} it is necessary that φ_1 distinguishes between some y_k and y_{k+1} with $y \in \{s, t, u, w\}$ and $2i + 1 \geq k \geq 2i - 1$. But $k \geq 2i - 1 \geq 2n - 1 > 2(n - 1) \geq 2d(\varphi_1)$ so φ distinguishing between y_k and y_{k+1} would contradict the induction hypothesis. The formula φ therefore cannot distinguish between x_{2i} and x_{2i+1}.

Let us then suppose that $\varphi = C_B(\varphi_1, \varphi_2)$. If B does not contain both a and at least one of b, c then (B, φ)-paths cannot take us far in this model: if $a \notin B$ then at most the worlds s_{2i}, u_{2i}, t_{2i} and w_{2i} are B-reachable from x_{2i} and if $b, c \notin B$ then at most the worlds s_{2i}, u_{2i+1} are B-reachable from s_{2i}, t_{2i} and w_{2i+1} from t_{2i}, u_{2i} and s_{2i-1} from u_{2i} or w_{2i} and t_{2i-1} from w_{2i}. In each case there are counterparts to these worlds reachable from x_{2i+1} and by the induction hypothesis φ_2 cannot distinguish between these counterpart worlds. So $C_B(\varphi_1, \varphi_2)$ cannot distinguish between x_{2i} and x_{2i+1}.

So in order to distinguish between x_{2i} and x_{2i+1} the set B must contain both a and at least one of b, c. If x_{2i} and x_{2i+1} are reachable from each other by a (B, φ_1)-path then, independent of φ_2, we have $\mathcal{M}_c, x_{2i} \models C_B(\varphi_1, \varphi_2)$ if and only if $\mathcal{M}_c, x_{2i+1} \models C_B(\varphi_1, \varphi_2)$. So in order to distinguish the worlds there must be at least one $\neg\varphi_1$ world on every B-path from x_{2i} to x_{2i+1}. Furthermore, this $\neg\varphi_1$ world cannot be either x_{2i} or x_{2i+1}, since then $C_B(\varphi_1, \varphi_2)$ would reduce to φ_2 on both worlds and φ_2 cannot distinguish between them.

The model $\mathcal{M}_c$ is constructed in such a way that if every B-path from x_{2i} to x_{2i+1} contains at least one $\neg\varphi_1$ world then so does every B-path from either of those worlds to any world further than two steps away. Exactly which worlds must satisfy $\neg\varphi_1$ depends on x and B though.

Suppose $x = u$ and $B \cap \{a, b, c\} = \{a, b, c\}$. Then in order for every B-path between u_{2i} and u_{2i+1} to contain at least one $\neg\varphi_1$ world we have to have $\mathcal{M}_c, s_j \not\models \varphi_1$ and either $\mathcal{M}_c, t_j \not\models \varphi_1$ or $\mathcal{M}_c, w_{2i+1} \not\models \varphi_1$. But, since φ_1 is of depth $\leq n-1$ it follows from the induction hypothesis that it cannot distinguish between s_{2i+2}, s_{2i+1}, s_{2i} and s_{2i-1}, between w_{2i+2}, w_{2i+1}, w_{2i} and w_{2i-1} or between t_{2i+2}, t_{2i+1}, t_{2i} and t_{2i-1}. But then every (B, φ_1)-path from u_{2i} can contain at most the worlds u_{2i} and t_{2i}. By the induction hypothesis none of these worlds can be distinguished from their counterpart by φ_2, so $C_B(\varphi_1, \varphi_2)$ cannot distinguish between u_{2i} and u_{2i+1}. See Figure 5 for an illustration of the (B, φ_1)-paths if w_{2i} and s_{2i} do not satisfy φ_1.

The situations for the other options for x and B are similar. The following table shows which worlds have to satisfy $\neg\varphi_1$ and which worlds (B,φ_1)-paths from x_{2i} can contain, at the most. Note that only maximal (with respect to set inclusion) sets are given. For example, one of the options for $\neg\varphi_1$ worlds in the case $x = u$ and $\{a,b,c\}$ is s_k and t_k (for all $2i+2 \geq k \geq 2i-1$) in which case paths from x_{2i} can contain only the world x_{2i} itself. But the other option is for s_k and w_k to satisfy $\neg\varphi_1$ in which case paths from x_{2i} could contain both u_{2i} and t_{2i}.

x	$B \cap \{a,b,c\}$	Must satisfy $\neg\varphi_1$ $(\forall k \in \{2i-1, 2i, 2i+1, 2i+2\})$	Paths from x_{2i} can contain at most
u	$\{a,b,c\}$	s_k and t_k or w_k	$\{u_{2i}, t_{2i}\}$
u	$\{a,b\}$	s_k	$\{u_{2i}, t_{2i}, w_{2i+1}\}$
u	$\{a,c\}$	s_k, t_k or w_k	$\{u_{2i}, s_{2i-1}, t_{2i-1}\}$ or $\{u_{2i}, s_{2i-1}, w_{2i-1}\}$
s	$\{a,b,c\}$	u_k and w_k or t_k	$\{s_{2i}, t_{2i}\}$ or $\{s_{2i}, w_{2i}\}$
s	$\{a,b\}$	u_k	$\{s_{2i}, t_{2i}, w_{2i+1}\}$
s	$\{a,c\}$	t_k or w_k	$\{s_{2i}, u_{2i+1}, t_{2i}\}$

In every one of these cases it follows from the induction hypothesis that φ_2 cannot distinguish between the path from x_{2i} and the path from x_{2i+1}. So $C_B(\varphi_1, \varphi_2)$ does not distinguish between x_{2i} and x_{2i+1}.

For every possible form of φ we have now shown that φ does not distinguish between x_{2i} and x_{2i+1}, contradicting the assumption that φ does distinguish between them and thereby completing the proof. $\square$

Recall that we did not assume any of the frame conditions often used in epistemic logic, so we use any K-model as opposed to only KD45-, S4- or S5-models. The model $\mathcal{M}_c$ that we used is the proof of Lemma 3 is an S5 model (and therefore also a KD45- and S4-model) though. As a result the proof does not depend on the fact that we use K-models, allowing us to conclude that AU* is more expressive than AR not only over K but also over KD45, S4 and S5.

6 AUC is equally expressive as AU*

In this section we show that AU$^* \preceq$ AUC. Unfortunately the proof is very long and technical. We therefore give only an overview of the proof here and leave the full proof as Appendices A, B and C in the supplementary data.

6.1 Notation

Due to the technical nature of the proof even the overview is made easier by introducing some more notation. First let us define some abbreviations regarding $\Box$ and $\Diamond$.

Definition 8. *For any $\varphi \in \Phi_{\mathcal{T}}$, $B \subseteq \mathcal{A}$ and U an arrow update let*

- $\boxdot_B\varphi$ *stand for* $\varphi \wedge \Box_B\varphi$ *and* $\Diamond_B\varphi$ *stand for* $\varphi \vee \Diamond_B\varphi$,
- $\Box_U\varphi$ *stand for* $\bigwedge_{(u_1,a,u_2)\in U}(u_1 \to \Box_a(u_2 \to \varphi))$,
- $\Diamond_U\varphi$ *stand for* $\bigvee_{(u_1,a,u_2)\in U}(u_1 \wedge \Diamond_a(u_2 \wedge \varphi))$.

The formulas $\Box_U\varphi$ and $\Diamond_U\varphi$ thus state that φ holds in every/at least one U-successor.[6]

It is also convenient to be able to specify certain arrows that are *not* to be retained. We do this by overlining the clauses that specify arrows that should be removed. Let $U_1 = \{(u_1, a_1, u_1'), \cdots, (u_k, a_k, u_k')\}$ and $U_2 = \{\overline{(u_{k+1}, a_{k+1}, u_{k+1}')}, \cdots, \overline{(u_{k+l}, a_{k+l}, u_{k+l}')}\}$. An arrow is retained by the update $U = U_1 \cup U_2$ iff it satisfies at least one of the clauses of U_1 and none of the clauses of U_2. We can define updates with overlined clauses as abbreviations of updates without overlined clauses. The trick is to note that an arrow from w_1 to w_2 for agent a satisfies none of the clauses in U_2 if and only if for every clause $(u, a, u') \in U_2$ either $\mathcal{M}, w_1 \not\models u$ or $\mathcal{M}, w_2 \not\models u'$. In such a case we can partition U_2 into U' and $U_2 \setminus U'$ where w_1 does not satisfy the start conditions of U' and w_2 does not satisfy the end conditions of $U_2 \setminus U'$. So, for any $U' \subseteq U_2$, if $(u_i, a_i, u_i') \in U_1$ then $\left(u_i \wedge \bigwedge_{\overline{(u_j,a_i,u_j')}\in U'} \neg u_j, a_i, u_i' \wedge \bigwedge_{\overline{(u_j,a_i,u_j')}\in U_2\setminus U'} \neg u_j'\right)$ satisfies a clause from U_1 and none of the clauses from U_2.

Definition 9. *For* $U_1 = \{(u_1, a_1, u_1'), \cdots, (u_k, a_k, u_k')\}$ *and* $U_2 = \{\overline{(u_{k+1}, a_{k+1}, u_{k+1}')}, \cdots, \overline{(u_{k+l}, a_{k+l}, u_{k+l}')}\}$ *let* $U_1 \cup U_2$ *stand for*

$$\left\{\left(u_i \wedge \bigwedge_{\overline{(u_j,a_i,u_j')}\in U'} \neg u_j, a_i, u_i' \wedge \bigwedge_{\overline{(u_j,a_i,u_j')}\in U_2\setminus U'} \neg u_j'\right) \mid (u_i, a_i, u_i') \in U_1, U' \in \wp(U_2)\right\}.$$

We use $\overline{U}$ as shorthand for $\{\top, \mathcal{A}, \top\} \cup \{\overline{(u, a_i, u')} \mid (u, a_i, u') \in U\}$. The formulas $\Box_{\overline{U}}\varphi$ and $\Diamond_{\overline{U}}\varphi$ thus state that φ holds in every/at least one world that is a successor but not a U-successor.

We also need notation for two more concepts about formulas.

Definition 10. *For* $\varphi \in \Phi_{\mathcal{T}}$ *let* $\mathrm{Pvar}(\varphi)$ *be the set of propositional variables that occur in* φ.

Definition 11. *For* $Q \subseteq \mathcal{P}$ *and* $n \in \mathbb{N}$ *let* $\Phi_Q^n := \{\varphi \in \Phi_{\mathrm{AUC}} \mid d(\varphi) \leq n$ *and* $\mathrm{Pvar}(\varphi) \subseteq Q\}$.

The main use of Φ_Q^n will be in conjunctions $\bigwedge_{\varphi\in\Phi_Q^n} \psi_\varphi$. Strictly speaking this is of course not a formula, as it contains an infinite number of conjuncts. However, if Q is finite—as it will be when we use it—the set Φ_Q^n contains only a finite number of *mutually non-equivalent* formulas. We can then consider $\bigwedge_{\varphi\in\Phi_Q^n} \psi_\varphi$ to be a conjunction over some maximal choice of non-equivalent formulas in Φ_Q^n.

[6]Note that $\Box_U$ is the static operator associated with $\Box$ and $[U]$ and that $\Diamond_U$ is the static operator associated with $\Diamond$ and $[U]$, as discussed in Section 2.1.

6.2 Variable use

In the proof that $\mathrm{AU}^* \preceq \mathrm{AUC}$ a large number of formulas are defined. While the names given to the formulas do not, strictly speaking, matter there is a pattern in the naming, and knowing this pattern may aid in understanding the proof. The proof finds a AUC formula α that is equivalent to $\{U\}^*\varphi$ by using a case distinction.

A formula δ_i is a AU^* formula corresponding to case i. A formula β_i is an AUC formula that is both necessary and sufficient for being in case i, so $\models \delta_i \leftrightarrow \beta_i$. A formula α_i finally is a AUC formula that is equivalent to $\{U\}^*\varphi$ given that we are in case i, so $\models (\delta_i \wedge \{U\}^*\varphi) \leftrightarrow (\beta_i \wedge \alpha_i)$.[7]

6.3 The main strategy

Fix any arrow update U containing only AUC formulas and any AUC formula φ and let $\chi := \{U\}^*\varphi$. If we can find a AUC formula α such that $\models \alpha \leftrightarrow \chi$ that would suffice to show that AUC is at least as expressive as AU^*.

What we need then is a strategy to find such α. This poses two challenges. First, given any pointed model $\mathcal{M}, w$ we must identify the worlds that are U-reachable from w. Second, we must check whether φ holds in all of those worlds.

The most straightforward method to identify the U-reachable worlds is to update with $[U]$; the U-reachable worlds in $\mathcal{M}$ are exactly the reachable worlds in $\mathcal{M}_{[U]}$. Unfortunately the update $[U]$ may destroy information, so given a world w' of $\mathcal{M}$ it may be impossible to determine from $\mathcal{M}_{[U]}$ whether $\mathcal{M}, w' \models \varphi$. By using this simple method to solve the first problem we would make it impossible to solve the second problem.

So in order to solve both problems we need to update with a different arrow update U'. This U' will be very similar to U so the reachable worlds in $\mathcal{M}_{[U']}$ are mostly the U-reachable worlds in $\mathcal{M}$. But in addition to most arrows from U the update U' will retain just enough structure to create *witnesses* for the existence of certain worlds in $\mathcal{M}$. The question then is what worlds we want to create witnesses for and how we want to use them. Again there is a straightforward choice, namely to create witnesses for $\neg\varphi$ worlds. But, again, this straightforward choice runs into trouble. So instead we create witnesses for worlds that are "on the boundary" of the U-reachable area, so those worlds reached by a U-arrow from which a $\overline{U}$-arrow departs. So we take U' in such a way that in $\mathcal{M}_{[U']}$ every maximal path ends in a witness world.

We then make one final change. Let $U'' = U' \cup \{\overline{(\neg\varphi, \mathcal{A}, \top)}\}$. This change cuts all paths that contain a $\neg\varphi$ world. So in $\mathcal{M}_{[U'']}$ maximal paths end in a witness world if and only if they do not contain a world that was a $\neg\varphi$ world in $\mathcal{M}$. So $\mathcal{M}, w \models \{U\}^*\varphi$ if and only if every maximal path from $\mathcal{M}_{[U'']}$ ends in a witness world. There are of course several complications, but those can be dealt with. Let us look at a very simple example.

[7] As the choice of the symbols α, β and δ suggests there are also formulas γ_i. These only occur in the detailed proof in the supplementary data though, where γ_i is a necessary condition for being in case i, so $\models \delta_i \rightarrow \gamma_i$.

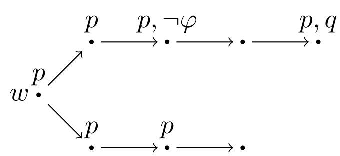

(a) The model $\mathcal{M}$.

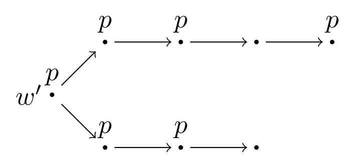

(b) The model $\mathcal{M}'$.

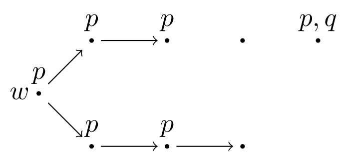

(c) The model $\mathcal{M}_{[U'']}$.

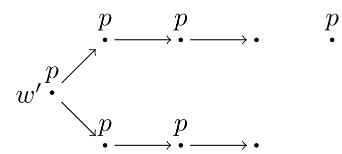

(d) The model $\mathcal{M}'_{[U'']}$.

Figure 6: An example of using $\neg p$ worlds as witnesses. The formula φ holds everywhere except where noted otherwise.

Example 2. Suppose $U = (p, \mathcal{A}, p)$. Then we can create witnesses for U-reachable worlds from which a $\overline{U}$-arrow departs by retaining arrows from p worlds to $\neg p$ worlds, so by taking $U' = (p, \mathcal{A}, p), (p, \mathcal{A}, \neg p)$. We then have $U'' = (p, \mathcal{A}, p), (p, \mathcal{A}, \neg p), \overline{(\neg\varphi, \mathcal{A}, \top)}$. Consider the pointed models $\mathcal{M}, w$ and $\mathcal{M}', w'$ as shown in Figure 6. We want to check whether every path from w (resp. w') in $\mathcal{M}_{[U'']}$ (resp. $\mathcal{M}'_{[U'']}$) ends in a witness world, so we check for $C_{\mathcal{A}}(\Box\bot \to \neg p)$. We have $\mathcal{M}_{[U'']}, w \not\models C_{\mathcal{A}}(\Box\bot \to \neg p)$ and $\mathcal{M}'_{[U'']}, w \models C_{\mathcal{A}}(\Box\bot \to \neg p)$ so $\mathcal{M}, w \not\models \{U\}^*\varphi$ and $\mathcal{M}', w' \models \{U\}^*\varphi$.

But even with an example as simple as $U = (p, \mathcal{A}, p)$ complications can occur if we look at different models. For example, the method as described above will not work if the origin world w satisfies either $\neg p$ or $C_{\mathcal{A}}p$. Both these cases can be easily dealt with however, by checking for them before applying $[U'']$ and treating them separately.

A more interesting kind of complication is if $\mathcal{M}, w \not\models C_{\mathcal{A}}p$ but there are some branches that are reachable from w that do satisfy $C_{\mathcal{A}}p$. If we simply apply $[U'']$ these branches will end in a p world and therefore look just like branches that were cut short due to the presence of a $\neg\varphi$ world. So we risk getting false positives for the detection of $\neg\varphi$ worlds. The solution to this complication is to cut off all $C_{\mathcal{A}}p$ branches unless they contain a $\neg\varphi$ world. This means we have to modify U'' to $U''' = U'' \cup \{\overline{(\top, \mathcal{A}, C_{\mathcal{A}}(p \wedge \varphi))}\}$.

These are all the complications we can encounter for this simple U however. The formula $\{(p, \mathcal{A}, p)\}^*\varphi$ is equivalent to

$$\varphi \wedge (C_{\mathcal{A}}p \to C_{\mathcal{A}}\varphi) \wedge (\neg C_{\mathcal{A}}p \to$$
$$[(p, \mathcal{A}, p), (p, \mathcal{A}, \neg p), \overline{(\neg\varphi, \mathcal{A}, \top)}, \overline{(\top, \mathcal{A}, C_{\mathcal{A}}(p \wedge \varphi))}]C_{\mathcal{A}}(\Box_{\mathcal{A}}\bot \to \neg p)),$$

where the first two conjuncts take care of the two degenerate cases and the third uses witnesses as described above.

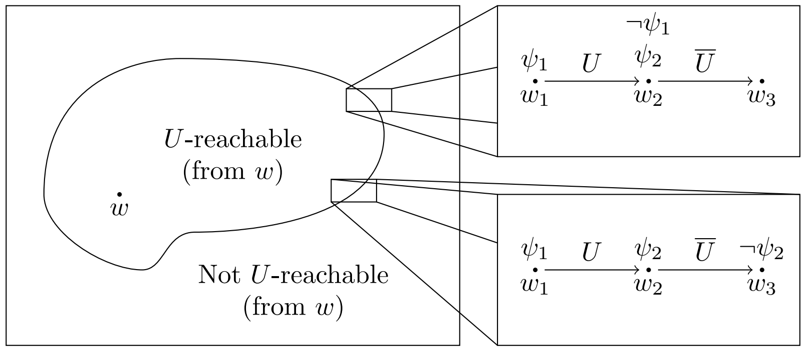

Figure 7: Two possibilities for a $\overline{U}$-arrow following a U-arrow.

6.4 Creating witnesses

The main strategy requires us to create witnesses for worlds that are on the boundary of a U-area. We leave the details of how to do this to the supplementary data, but here we do present a global overview of why it is always possible to create such a witness.

The important realization is that whenever there is a world reached by a U-arrow and from which a $\overline{U}$-arrow departs then there must be a *simple* difference between two *nearby* objects. We call this difference between the two objects a *boundary condition*.[8]

Here *nearby* means "reachable in at most $d(\{U\}^*\varphi)$ steps". What it means for a difference to be *simple* is a little more complicated. Let us take a closer look at the situation. We have a world, call it w_2, that is reached by a U-arrow and from which a $\overline{U}$ arrow departs. Let w_1 be the source of the U-arrow and w_3 the destination of the $\overline{U}$-arrow.

Let us focus on the arrows from w_1 to w_2 and from w_2 to w_3 for a moment. The arrow from w_1 to w_2 is a U-arrow so there is a clause $(\psi_1, a, \psi_2) \in U$ that the arrow satisfies. The arrow from w_2 to w_3 is not a U-arrow so in particular it does not satisfy (ψ_1, a, ψ_2). Then there are three possibilities: the first possibility is that the arrow from w_2 to w_3 is not an a-arrow. The second possibility is that $\mathcal{M}, w_2 \not\models \psi_1$. The third possibility is that $\mathcal{M}, w_3 \not\models \psi_2$. For the latter two possibilities see Figure 7.

If the first possibility holds the nearby objects that differ are arrows, and the simple difference between them is that they belong to different agents. So the first boundary condition is that there are nearby arrows belonging to different agents. Let us suppose then that there are no two nearby arrows that belong to different agents. So we are working in a single-agent part of the model.

Let us take a closer look at the second and third possibilities. In both cases there is a U-reachable world w_i satisfying $\psi_i \wedge \Diamond\neg\psi_i$ with $i \in \{1, 2\}$. This

[8]Note that boundary conditions are not called so because they only occur on the boundary. Instead they are called so because they have to occur on every boundary.

difference between the world w_i satisfying ψ_i and the world w_{i+1} satisfying $\neg\psi_i$ has to “come from somewhere”.

One possible cause for the difference between w_i and w_{i+1} is the existence of two nearby worlds w' and w'' and a propositional variable p such that $\mathcal{M}, w' \models p$ and $\mathcal{M}, w'' \not\models p$. This is the second kind of boundary condition; a difference in the value of a propositional variable in two worlds.

Suppose then that we are in a situation where there are no nearby arrows for different agents and no nearby difference in propositional variables. Then the only possible cause for the difference between w_i and w_i' is the existence of one or more worlds satisfying $\Box\bot$ (and some satisfying $\Diamond\top$). This is the third kind of simple difference, some worlds having a successor and other worlds having none.

For technical reasons it is convenient to split this simple difference up unto two boundary conditions. The first boundary condition is when there is a nearby world satisfying $\Diamond\psi' \wedge \Diamond\neg\psi'$, with ψ' a formula of depth lower than ψ_i. We are still in the situation where there are no different agents and no different propositional variables, so either the ψ' or the $\neg\psi'$ branch must contain a nearby $\Box\bot$ world. But it is not necessary that both branches contain such a world. The other boundary condition is when w_i is near a dead end. In that case $\mathcal{M}, w_i \models \Diamond^k\Box\bot \wedge \Box^{k+1}\bot$ for some $k \leq n$.

So the four kinds of boundary condition are:

1. Two arrows that belong to different agents.
2. Two worlds that have a different value for some propositional variable.
3. A world satisfying $\Diamond\psi' \wedge \Diamond\neg\psi'$ for some ψ'.
4. The world w_i being one step further away from a dead end than the world w_{i+1}.

In addition to these four conditions there are also two “degenerate boundary conditions”. These represent situations where there is no boundary at all.

5. There are reachable U-arrows but no reachable $\overline{U}$-arrows.
6. There are no departing U-arrows at all.

These are the only possibilities. This should be intuitively clear, but we also provide a full proof in Section C in the supplementary data.

For each of the first four boundary conditions we can create a witness, and in the last two we do not need a witness because $\{U\}^*\varphi$ reduces to $C_{\mathcal{A}}\varphi$ or φ respectively. The witness for the first boundary condition is simply the arrow that belongs to a different agent. The witness for the second boundary condition is a world with a different value for the variable. The witness for the fourth boundary condition is the dead end itself. The only difficult boundary condition is the third, where we have $\Diamond\psi' \wedge \Diamond\neg\psi'$. In that case we cut off either the branch starting at the ψ' successor or the branch starting at the $\neg\psi'$ successor. This results in a world satisfying $\Diamond\Box\bot \wedge \Diamond\Diamond\top$, which we use as witness.

6.5 The case distinction

We have six different boundary conditions (two of which are degenerate). Unfortunately each of those is solved in a different way. This can be problematic, because there may be different branches with different boundary conditions and we can only solve one of them at a time. The solution is to make a case distinction that allows us to solve the different conditions one at a time.

We have our four types of boundary condition and two types of degenerate boundary condition. Based on these types we want to make a case distinction. To every world in any model we assign one of six cases. Case number i is associated with boundary condition type i, but not in an entirely straightforward way: a world is in case i if and only if there is at least one U-reachable world satisfying boundary condition type i and no U-reachable world satisfying boundary condition type $j < i$. Do recall that boundary conditions may "coincidentally" hold in worlds that are not on the boundary of the U-area. These "coincidental" boundary conditions do count for which case we are in.

Defining the cases in this way gives us three important properties for our case distinction. Firstly the cases are mutually exclusive, since case i requires a lack of worlds with condition $j < i$. Secondly the cases are exhaustive, there is always an accessible branch with one of the (possibly degenerate) boundary conditions. Finally, if a world w is in case i and w' is U-reachable from w then w' is either in case i or in a case $j > i$. After all, if a world with boundary condition $k < i$ would be U-reachable from w' then the same branch would be U-reachable from w contradicting the assumption that w is in case i.

These three properties allow us to solve all cases by "working backwards". We first find a formula α_{i+1} that is equivalent to $\{U\}^*\varphi$ on worlds in case $i+1$. Then we use the fact that we have already solved case $i+1$ to solve case i. The process is best explained with the help of a series of example figures, so consider Figures 8–11.

Example 3. First consider the model $\mathcal{M}_t$ as shown in Figure 8. If a world satisfies boundary condition i then the world is labeled b_i. There could be boundary conditions of many different types but in order to keep the example simple we consider an example where only boundary conditions of types 1, 2 and 3 occur.

Note that some worlds satisfy multiple boundary conditions. In particular, every world that is the endpoint of a branch, and therefore satisfies $\Box_U\bot$, satisfies condition 6 in addition to any other conditions it may satisfy. This does not matter though; if a world w satisfies conditions i and $j > i$ then any world that can reach w is either in case i or in some case $k < i$. Also note that there is one "coincidental" boundary condition of type 1 in the lowest branch of the model. Some of the worlds are labeled w_j for some $1 \leq j \leq 5$. We will be taking a closer look at those worlds. The worlds w_1 and w_2 are in case 3, the world w_3 is in case 2 and the worlds w_4 and w_5 are in case 1.

In order to determine on which worlds in this model the formula $\{U\}^*\varphi$ holds we start by considering those worlds that are in case 3. There are no boundary conditions of type 4, 5 or 6 in this model so we know how to solve

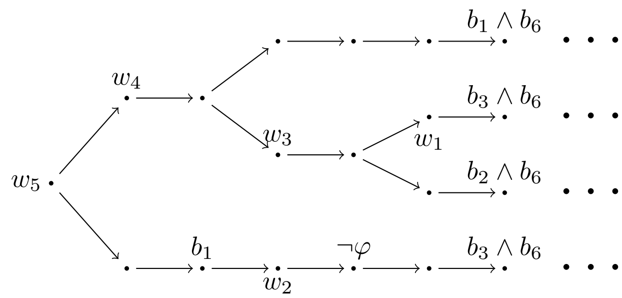

Figure 8: The U-reachable part of an example model $\mathcal{M}_t$ used to illustrate the method of working backward through the cases. The dots at the end of each branch represent the fact that the model continues, but with non-U arrows.

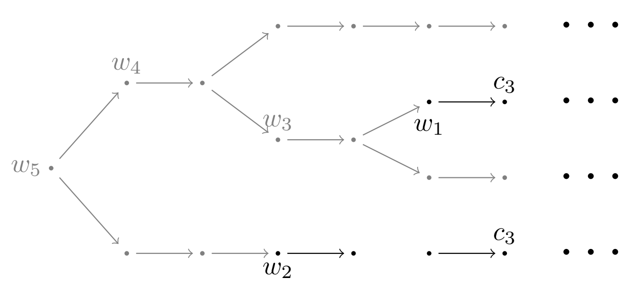

Figure 9: The model $\mathcal{M}_{t[U_3'']}$. Worlds not in case 3 are grayed out. A witness for condition 3 is represented by c_3.

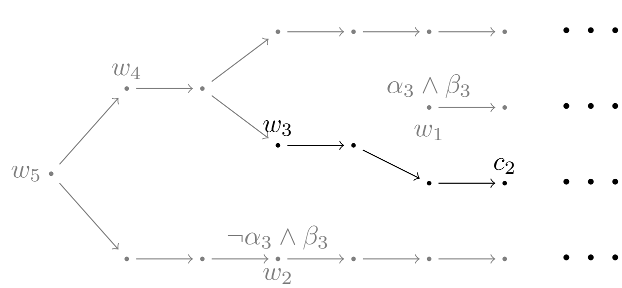

Figure 10: The model $\mathcal{M}_{t[U_2'' \cup \overline{(\top, \mathcal{A}, \alpha_3 \wedge \beta_3)}]}$. Worlds not in case 2 are grayed out. A witness for condition 2 is represented by c_2.

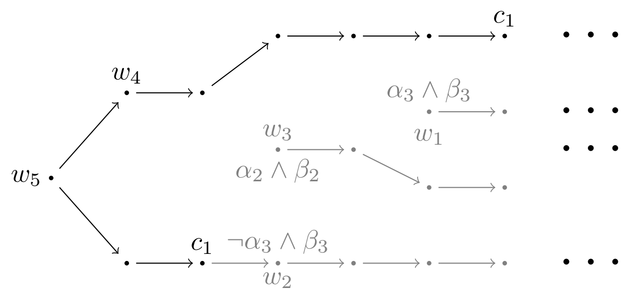

Figure 11: The model $\mathcal{M}_{t[U_3'' \cup \overline{(\top,\mathcal{A},\alpha_2\wedge\beta_2)} \cup \overline{(\top,\mathcal{A},\alpha_3\wedge\beta_3)}]}$. Worlds not in case 1 are grayed out. A witness for condition 1 is represented by c_1.

case 3; we use an update $[U_3'']$ to create witnesses for worlds satisfying boundary condition 3 and additionally remove arrows to $\neg\varphi$ worlds. See Figure 9. We can then see that every path from w_1 in $\mathcal{M}_{t[U_3'']}$ ends in a witness world, so $\mathcal{M}_t, w_1 \models \{U\}^*\varphi$. There is however a path from w_2 in $\mathcal{M}_{t[U_3'']}$ that does not end in a witness world so $\mathcal{M}_t, w_2 \not\models \{U\}^*\varphi$.

Now we can use what we have learned about case 3 to solve case 2. Just like in case 3 we use an update U_2'' that creates witnesses for boundary condition 2 and that removes arrows to $\neg\varphi$ worlds. But we add one additional clause. Recall that we use β_3 for the formula that identifies the case 3 worlds and α_3 for the formula that is equivalent to $\{U\}^*\varphi$ in case 3 worlds. This means that the clause $\overline{(\top,\mathcal{A},\alpha_3\wedge\beta_3)}$ removes arrows to case 3 worlds if and only if they satisfy $\{U\}^*\varphi$.[9] The resulting model $\mathcal{M}_{t[U_2'' \cup \overline{(\top,\mathcal{A},\alpha_3\wedge\beta_3)}]}$ is shown in Figure 10. Note that all paths from w_3 end in a witness world, so $\mathcal{M}_t, w_3 \models \{U\} * \varphi$.

Finally we can use both previous cases to solve case 1. Again we start with an update U_1'' to create witnesses and remove arrows to $\neg\varphi$ worlds. But now we add two extra clauses, $\overline{(\top,\mathcal{A},\alpha_2\wedge\beta_2)}$ and $\overline{(\top,\mathcal{A},\alpha_3\wedge\beta_3)}$. See Figure 11 for the resulting model $\mathcal{M}_{t[U_3'' \cup \overline{(\top,\mathcal{A},\alpha_2\wedge\beta_2)} \cup \overline{(\top,\mathcal{A},\alpha_3\wedge\beta_3)}]}$. We have $\mathcal{M}_t, w_2 \not\models \alpha_3$, so the arrow to w_2 is not removed by the update. As a result there is a path from w_5 that does not end in a witness world, so $\mathcal{M}_t, w_5 \not\models \{U\}^*\varphi$. Note that, unlike the arrow to w_2, the arrow to w_3 is removed, because $\mathcal{M}_t, w_3 \models \alpha_2 \wedge \beta_2$. As a result every path from w_4 ends in a witness world, so $\mathcal{M}_t, w_4 \models \{U\}^*\varphi$.

6.6 Further Complications

The method detailed above allows us to deal with the most important complication, namely that there could be different branches with different boundary conditions. Unfortunately there are several remaining complications. Notable examples include the fact that the $\overline{U}$-arrows retained to create witnesses might connect one U-area to another, and the fact that for one of the witness types

[9] If there were branches in cases 4, 5 or 6 we would have to add additional clauses $\overline{(\top,\mathcal{A},\alpha_4\wedge\beta_4)}$, $\overline{(\top,\mathcal{A},\alpha_5\wedge\beta_5)}$ and $\overline{(\top,\mathcal{A},\alpha_6\wedge\beta_6)}$ as well.

it is impossible to tell whether a given path ends in (as opposed to contains) a witness world. We also need a large number of subcases in addition to the six main cases.

These complications only arise in the detailed proof however, so the ways to deal with these complications are also given there.

6.7 Formulas Representing the Cases

Above the different cases were described informally. But before proving that $\mathrm{AU}^* \preceq \mathrm{AUC}$ we should find formal descriptions of the cases. Which case we are in is determined by the existence or nonexistence of certain kinds of worlds within a certain distance of a U-reachable world. This kind of condition is easily phrased in AU^* but not in AUC so let us first describe the conditions in English and AU^*.

Recall that $\chi = \{U\}^*\varphi$ where φ is a AUC formula and U contains only AUC formulas. Let $n := d(\chi)$. The nearby difference must then be within distance n of a U-reachable world. For technical reasons we check for some differences up to a distance of a multiple of n though. This gives us the following cases.

1. We are in the first case if $\Diamond_U\top$ and there are at least two agents a_1 and a_2 that have an arrow within distance $3n$ of a U-reachable world. The AU^* representation of this case is

$$\delta_1 := \Diamond_U\top \wedge \bigvee_{a_1\neq a_2\in\mathcal{A}} (\neg\{U\}^*\boxdot^{3n}\Box_{a_1}\bot \wedge \neg\{U\}^*\boxdot^{3n}\Box_{a_2}\bot).$$

2. We are in the second case if $\Diamond_U\top$, we are not in the first case and there is a propositional variable $p \in \mathrm{Pvar}(\chi)$ such that both p and $\neg p$ hold within distance $3n$ of a U-reachable world. The AU^* representation of this case is

$$\delta_2 := \Diamond_U\top \wedge \neg\delta_1 \wedge \bigvee_{p\in\mathrm{Pvar}(\chi)} (\neg\{U\}^*\boxdot^{3n}p \wedge \neg\{U\}^*\boxdot^{3n}\neg p).$$

3. We are in the third case if $\Diamond_U\top$, we are not in the first or second case and there is a world w_2 within distance n of a U-reachable world such that the successors of w_2 are distinguishable by a AUC formula of depth at most $2n$ using only the propositional variables in χ. The AU^* representation of this case is

$$\delta_3 := \Diamond_U\top \wedge \neg\delta_1 \wedge \neg\delta_2 \wedge \bigvee_{\psi\in\Phi^{2n}_{\mathrm{Pvar}(\chi)}} \neg\{U\}^*\neg\Diamond\!\!\cdot^{\,n}(\Diamond\psi \wedge \Diamond\neg\psi)$$

4. We are in the fourth case if $\Diamond_U\top$, we are not in one of the previous cases and there is a U-reachable world where $\Diamond_{\overline{U}}\top$ holds. The AU^* representation of this case is

$$\delta_4 := \Diamond_U\top \wedge \neg\delta_1 \wedge \neg\delta_2 \wedge \neg\delta_3 \wedge \neg\{U\}^*\neg\Diamond_{\overline{U}}\top$$

5. We are in the fifth case if $\Diamond_U \top$, we are not in one of the previous cases and there is no U-reachable world where $\Diamond_{\overline{U}} \top$ holds. The AU* representation of this case is

$$\delta_5 := \Diamond_U \top \wedge \neg\delta_1 \wedge \neg\delta_2 \wedge \neg\delta_3 \wedge \neg\delta_4 \wedge \{U\}^* \neg \Diamond_{\overline{U}} \top$$

6. We are in the sixth case if $\neg\Diamond_U \top$. The AU* representation of this case is

$$\delta_6 := \neg\Diamond_U \top.$$

In order to construct the AUC formula α that is equivalent to the AU* formula χ we first have to find AUC formulas $\beta_1, \cdots, \beta_6$ such that $\models \beta_i \leftrightarrow \delta_i$ for $i \in \{1, \cdots, 6\}$. Then we find α_6 such that $\models (\delta_6 \wedge \chi) \leftrightarrow (\beta_6 \wedge \alpha_6)$, use this α_6 to find α_5 such that $\models (\delta_5 \wedge \chi) \leftrightarrow (\beta_5 \wedge \alpha_5)$ and so on until we have $\alpha_1, \cdots, \alpha_6$ such that

$$\alpha = (\beta_1 \to \alpha_1) \wedge (\beta_2 \to \alpha_2) \wedge (\beta_3 \to \alpha_3) \wedge (\beta_4 \to \alpha_4) \wedge (\beta_5 \to \alpha_5) \wedge (\beta_6 \to \alpha_6).$$

6.8 AU* $\preceq$ AUC

Lemma 4. *There are AUC formula $\beta_1, \cdots, \beta_6$ such that $\models \beta_i \leftrightarrow \delta_i$ for all $1 \leq i \leq 6$.*

The proof of this is included as Section A in the supplementary data.

Lemma 5. *There are AUC formulas $\alpha_6, \cdots, \alpha_1$ such that $\models (\delta_i \wedge \chi) \leftrightarrow (\delta_i \wedge \alpha_i)$ for all $6 \geq i \geq 1$.*

The proof of this is included as Section B in the supplementary data.

Lemma 6. *There is an AUC formula α such that $\models \chi \leftrightarrow \alpha$.*

Proof. The cases $\delta_1, \cdots, \delta_6$ are exhaustive, so $\models \chi \leftrightarrow ((\delta_1 \to \chi) \wedge \cdots \wedge (\delta_6 \to \chi))$. Then by Lemma 5 there are AUC formulas $\alpha_1, \cdots, \alpha_6$ such that $\models \chi \leftrightarrow ((\delta_1 \to \alpha_1) \wedge \cdots \wedge (\delta_6 \to \alpha_6))$. Furthermore, by Lemma 4 there are AUC formulas $\beta_1, \cdots, \beta_6$ such that $\models \chi \leftrightarrow ((\beta_1 \to \alpha_1) \wedge \cdots \wedge (\beta_6 \to \alpha_6))$. This proves the lemma with $\alpha = (\beta_1 \to \alpha_1) \wedge \cdots \wedge (\beta_6 \to \alpha_6)$. □

The formula χ was taken $\chi = \{U\}^* \varphi$ with any update U containing only AUC formulas and any AUC formula φ. Lemma 6 therefore allows us to eliminate all occurrences of operators $\{U\}^*$ in any AU* formula by first eliminating the innermost occurrences and working outward. We therefore have the following corollary.

Theorem 1. AUC $\equiv$ AU*

Proof. It was shown in [11] that AUU* $\equiv$ AU*, which together with Lemma 1 shows that AUC $\preceq$ AU*. Furthermore, it follows from Lemma 6 that AU* $\preceq$ AUC. We therefore have AUC $\equiv$ AU*. □

7 Conclusion

We have demonstrated two new expressivity results about logics using arrow updates. The first result is that the logic $\mathcal{L}_{\mathrm{AU}^*}$ using arrow common knowledge is strictly more expressive than the logic $\mathcal{L}_{\mathrm{AR}}$ using relativised common knowledge. This result is not surprising and had in fact been predicted in [11] where AU* was introduced.

The second result is that the logic $\mathcal{L}_{\mathrm{AUC}}$ using arrow updates and common knowledge is as expressive as $\mathcal{L}_{\mathrm{AU}^*}$. This result is rather surprising considering that the logic $\mathcal{L}_{\mathrm{AR}}$ was shown in [6] to be strictly more expressive than the logic $\mathcal{L}_{\mathrm{APC}}$ using public announcements and common knowledge, and the difference between $\mathcal{L}_{\mathrm{AUC}}$ and $\mathcal{L}_{\mathrm{AU}^*}$ is comparable to the difference between $\mathcal{L}_{\mathrm{APC}}$ and $\mathcal{L}_{\mathrm{AR}}$.

These two new results together with results from [13], [4], [12], [6] and [11] fully determine the expressivity landscape of all logics using any combination of common knowledge, relativised common knowledge, public announcements, arrow updates and arrow common knowledge. This expressivity landscape is shown in Figure 2.

One interesting property of this landscape is that there are no logics in it with incomparable expressivity; if $\mathcal{L}_{\mathrm{X}}$ and $\mathcal{L}_{\mathrm{Y}}$ are logics of the kind under consideration then either $\mathcal{L}_{\mathrm{X}} \preceq \mathcal{L}_{\mathrm{Y}}$ or $\mathcal{L}_{\mathrm{Y}} \preceq \mathcal{L}_{\mathrm{X}}$.

A remaining open question is the succinctness of the different logics. In particular, the translation from $\mathcal{L}_{\mathrm{AU}^*}$ to $\mathcal{L}_{\mathrm{AUC}}$ demonstrated in this paper has an extremely high growth in formula size. Whether this is necessarily so or there is an efficient translation is not currently known.

Acknowledgments

I would like to thank Barteld Kooi, Allard Tamminga and two anonymous referees for carefully reading an earlier version of this paper (including the highly technical supplementary data) and for providing very helpful feedback.

References

[1] Guillaume Aucher, Philippe Balbiani, Luis Fariñas del Cerro, and Andreas Herzig. Global and local graph modifiers. In Carlos Areces and Stéphane Demri, editors, *Proceedings of the 5th Workshop on Methods for Modalities (M4M5 2007)*, pages 293–307, 2009.

[2] Alexandru Baltag. A logic of epistemic actions. In W. van der Hoek, J-J.C. Meyer, and C. Witteveen, editors, *Proceedings of ESSLLI'99 Workshop on Foundations and Applications of Collective Agent Based Systems*, 1999.

[3] Alexandru Baltag and Lawrence S. Moss. Logics for epistemic programs. *Synthese*, 139:165–224, 2004.

[4] Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. The logic of public announcements, common knowledge and private suspicions. Technical Report TR534, Indiana University Bloomington, 1999.

[5] Johan van Benthem, Jelle Gerbrandy, Tomohiro Hoshi, and Eric Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, 38(5):491–526, 2009.

[6] Johan van Benthem, Jan van Eijck, and Barteld Kooi. Logics of communication and change. *Information and Computation*, 204:1620–1662, 2006.

[7] Hans van Ditmarsch. *Knowledge Games*. PhD thesis, University of Groningen, 2000.

[8] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer, 2004.

[9] Jelle Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, University of Amsterdam, 1999.

[10] Barteld Kooi. Expressivity and completeness for public update logics via reduction axioms. *Journal of Applied Non-Classical Logics*, 17:231–253, 2007.

[11] Barteld Kooi and Bryan Renne. Arrow update logic. *The Review of Symbolic Logic*, 4:536–559, 2011.

[12] Barteld Kooi and Johan van Benthem. Reduction axioms for epistemic actions. In Renate Schmidt, Ian Pratt-Hartmann, Mark Reynolds, and Heinrich Wansing, editors, *AiML-2004: Advances in Modal Logic*, volume UMCS-04-9-1 of *Technical Report Series*, pages 197–211. University of Manchester, 2004.

[13] Jan Plaza. Logics of public communication. In M.L. Emrich, M.S. Phifer, M. Hadzikadic, and Z.W. Ras, editors, *Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems, Poster Session Program*, pages 201–216. Oak Ridge National Laboratory, ORNL/DSRD-24, 1989. Reprinted as [14].

[14] Jan Plaza. Logics of public communication. *Synthese*, 158:165–179, 2007.

A Constructing $\beta_1, \cdots, \beta_6$

In order to prove Lemma 4 we have to show that there are $\beta_1, \cdots, \beta_6$ such that $\models \delta_i \leftrightarrow \beta_i$ for all $1 \leq i \leq 6$. Here we construct the β_i.

A.1 Constructing β_1

Recall that case 1 is the case where there are at least two agents for which there is an arrow departing from a world within $3n$ steps of a U-reachable world. We have

$$\delta_1 = \Diamond_U \top \wedge \bigvee_{a_1 \neq a_2 \in \mathcal{A}} (\neg\{U\}^* \boxdot^{3n} \Box_{a_1} \bot \wedge \neg\{U\}^* \boxdot^{3n} \Box_{a_2} \bot).$$

A.1.1 Subcases of case 1

There are several subcases of case 1. Let $B_1, \cdots, B_{2^{|\mathcal{A}|} - |\mathcal{A}| - 1}$ be all the subsets of $\mathcal{A}$ with at least two elements, ordered in such a way that if $B_i \subset B_j$ then $i > j$ and let $\mathcal{A} = \{a_1, \cdots, a_{|\mathcal{A}|}\}$. The subcases of case 1 are the cases $1.i.$-1 with $1 \leq 1 \leq 2^{|\mathcal{A}|} - |\mathcal{A}| - 1$ and the cases $1.i.j.k$ with $0 \leq i, j \leq |\mathcal{A}|$, $i \neq j$ and $0 \leq k \leq 3n$.

The case $1.i.$-1 corresponds to the case where there are U-paths from w that contain multiple agents (so if there are multiple agents for which a U-arrow departs from a U-reachable world) and B_i is exactly the set of agents for which there is such an arrow. The case $1.i.j.k$ corresponds to the case where the U-paths contain only agent a_i, j is the smallest number other than i such that an a_j arrow departs from a world within distance $3n$ from a U-reachable world and k is the shortest distance from a U reachable world to a world from which an a_k arrow departs.

The minimality conditions serve to make the different subcases easy to order. The cases $1.1. - 1$ to $1.2^{|\mathcal{A}|} - |\mathcal{A}| - 1. - 1$ followed by the cases 1.1.2.0 to $1.|\mathcal{A}|.|\mathcal{A} - 1|.3n$ are mutually exclusive, exhaustive of case 1, and taking a U-arrow can take you from one case to a later one but never to a previous one.

A.1.2 Constructing $\beta_{1.i.-1}$

For $1 \leq i \leq 2^{|\mathcal{A}|} - |\mathcal{A}| - 1$ let

$$\gamma_{1.i.-1} := \bigwedge_{a \in B_i} [U] \neg C_{\mathcal{A}} \Box_a \bot$$

and

$$\beta_{1.i.-1} := \gamma_{1.i.-1} \wedge \bigwedge_{j<i} \neg\gamma_{1.j.-1}.$$

The formula $\beta_{1.i.-1}$ holds iff the agents in B_i all occur in U-paths and there is no superset of B_i for which this is the case, so if B_i is exactly the agents that occur in U-paths, which is case $1.i. - 1$. This already implies $\Diamond_U \top$ so we don't have to include it explicitly.

A.1.3 Formulas useful for case $1.i.j.k$

Let
$$U_a^0 := U \cup \{(\top, a, \top)\}.$$
Furthermore, for $a \in \mathcal{A}$ let
$$U_a^{i+1} := U_a^i \cup \{(\Diamond^{i+1}\Diamond_a\top, \mathcal{A}, \Diamond^i\Diamond_a\top)\}.$$
Now let
$$\gamma_{1.a.i} := [U_a^i]\neg C_{\mathcal{A}}\Box_a\bot.$$
The arrow update U_a^i retains exactly the arrows that are U-arrows, a-arrows or arrows leading towards a $\Diamond_a\top$ world within distance i.

Suppose now that there are no a-arrows departing from any world within i steps of a U-reachable world. Then the only reachable arrows that are retained are the U-arrows, none of which is an a-arrow by assumption. We then have $\neg\gamma_{1.a.i}$.

Suppose on the other hand that there is an a arrow departing from a world within i steps of a U-reachable world. Then because of the $(\Diamond^{j+1}\Diamond_a\top, \mathcal{A}, \Diamond^j\Diamond_a\top)$ clauses the path to this a arrow will be retained and because of the $(\top, a, \top)$ clause the a arrow itself will be retained. We then have $\gamma_{1.a.i}$.

We thus have $\models \gamma_{1.a.i} \leftrightarrow \neg\{U\}^*\boxdot^i\Box_a\bot$.

A.1.4 Constructing $\beta_{1.i.j.k}$

For $1 \leq i, j \leq |\mathcal{A}|$, $i \neq j$ and $0 \leq k \leq 3n$ let
$$\gamma_{1.i.j.k} := \Diamond_{a_i}\top \wedge \bigwedge_{l \neq i}[U]C_{\mathcal{A}}\Box_{a_j}\bot \wedge \gamma_{1.a_j.k}.$$
Now let
$$\beta_{1.i.j.0} := \gamma_{1.i.j.0} \wedge \bigwedge_{j'<j, j'\neq i} \neg\gamma_{1.i.j'.3n}.$$
and for $k > 0$
$$\beta_{1.i.j.k} := \gamma_{1.i.j.k} \wedge \neg\gamma_{1.i.j.k-1} \wedge \bigwedge_{j'<j, j'\neq i} \neg\gamma_{1.i.j'.3n}.$$
The formula $\gamma_{1.i.j.k}$ holds iff all of the following hold: (in order of the conjunct that guarantees the property)

- the U-paths contain only agent a_i but there is an a_j arrow departing from a world within k steps of a U-reachable world
- there is no a_j arrow departing from a world within $k-1$ steps of a U-reachable world
- there is no $j' < j$ such that $j \neq i$ and there is an $a_{j'}$ arrow departing from a world within $3n$ steps of a U-reachable world.

The formula $\beta_{1.i.j.k}$ therefore holds exactly in case $1.i.j.k$.

A.1.5 Constructing β_1

Let us take

$$\beta_1 := \bigvee_{1\leq i\leq 2^{|\mathcal{A}|}-|\mathcal{A}|-1} \beta_{1.i.-1} \vee \bigvee_{1\leq i\leq|\mathcal{A}|} \bigvee_{1\leq j\leq|\mathcal{A}|, j\neq i} \bigvee_{1\leq k\leq 3n} \beta_{1.i.j.k}.$$

Because the subcases are exhaustive of case 1 we then have $\models \delta_1 \leftrightarrow \beta_1$.

A.2 Constructing β_2

Recall that case 2 is the case where there is a propositional variable $p \in \mathrm{Pvar}(\chi)$ such that both p and $\neg p$ hold in some world within $3n$ steps of a U-reachable world. We have

$$\delta_2 = \Diamond_U \top \wedge \neg\delta_1 \wedge \bigvee_{p\in\mathrm{Pvar}(\chi)} (\neg\{U\}^* \boxdot^{3n} p \wedge \neg\{U\}^* \boxdot^{3n} \neg p).$$

A.2.1 Subcases of case 2

Let $\mathrm{Pvar}(\chi) = \{p_1, \cdots, p_{|\mathrm{Pvar}(\chi)|}\}$. There are also several subcases of case 2, cases $2.i.j$ for $1 \leq i \leq |\mathrm{Pvar}(\chi)|$ and $0 \leq j \leq 3n$. The case $2.i.j$ is the case where both p_i and $\neg p_i$ occur within distance j of a U-reachable world but not within distance $j-1$ and there is no $i' < i$ such that both $p_{i'}$ and $\neg p_{i'}$ occur within distance $3n$ of a U-reachable world.

A.2.2 Constructing $\beta_{2.i.j}$

Let

$$U^+_{2.i.0} := U,$$
$$U^-_{2.i.0} := U,$$
$$U^+_{2.i.j+1} := U^+_{2.i.j} \cup \{(\Diamond^{j+1} p_i, a, \Diamond^j p_i)\}$$

and

$$U^-_{2.i.j+1} := U^-_{2.i.j} \cup \{(\Diamond^{j+1} \neg p_i, a, \Diamond^j \neg p_i)\}.$$

The update $[U^+_{2.i.j}]$ retains all arrows in U and all arrows to p_i worlds that can be reached within j steps. Likewise, the update $[U^-_{2.i.j}]$ retains all arrows in U and all arrows to $\neg p_i$ worlds that can be reached within j steps. As such, if there is any p_i world within j steps of a U-reachable world the formula $\neg[U^+_{2.i.j}]C_{\mathcal{A}}\neg p_i$ will hold and if there is any $\neg p_i$ world within j steps of a U-reachable world the formula $\neg[U^-_{2.i.j}]C_{\mathcal{A}} p_i$ will hold.

For $1 \leq i \leq |\mathrm{Pvar}(\chi)|$ and $0 \leq j \leq 3n$ let

$$\gamma_{2.i.j} := \neg[U^-_{2.i.j}]C_{\mathcal{A}} p_i \wedge \neg[U^+_{2.i.j}]C_{\mathcal{A}} \neg p_i$$

and

$$\beta_{2.i.j} := \Diamond_U \top \wedge \neg\beta_1 \wedge \gamma_{2.i.j} \wedge \neg\gamma_{2.i.j-1} \wedge \neg \bigvee_{i'<i} \gamma_{2.i'.3n}.$$

The formula $\beta_{2.i.j}$ thus holds iff both p_i and $\neg p_i$ occur within j steps of a U-reachable world but not within $j-1$ steps of a U-reachable world and there is no $i' < i$ such that both $p_{i'}$ and $\neg p_{i'}$ occur within $3n$ steps of a U-reachable world.

A.2.3 Constructing β_2

Let us take

$$\beta_2 := \bigvee_{1\leq i\leq |\mathrm{Pvar}(\chi)|} \bigvee_{0\leq j\leq 3n} \beta_{2.i.j}.$$

Because the subcases are exhaustive of case 2 we then have $\models \delta_2 \leftrightarrow \beta_2$.

A.3 Constructing β_3

Recall that case 3 is the case where there is a world w_1 near a U-reachable world such that the successors of w_1 are distinguishable by a short formula, so there is a short ψ such that $\mathcal{M}, w_1 \models \Diamond\psi \wedge \Diamond\neg\psi$. We have

$$\delta_3 = \Diamond_U \top \wedge \neg\delta_1 \wedge \neg\delta_2 \wedge \bigvee_{\psi\in\Phi^{2n}_{\mathrm{Pvar}(\chi)}} \neg\{U\}^*\neg\hat{\Diamond}^n(\Diamond\psi \wedge \Diamond\neg\psi)$$

A.3.1 Subcases of case 3

The subcases of case 3 are the cases $3.i$ for $0 \leq i \leq n$. The case $3.i$ is the case where the closest world w_1 with successors distinguishable by formulas in $\Phi^{2n}_{\mathrm{Pvar}(\chi)}$ is at distance i from a U-reachable world.

A.3.2 Introduction to case 3

Case 3 is, unfortunately, significantly more complicated than the previous cases. The main idea is the same, we use an update $U^{\psi}_{3.i}$ that retains U-arrows and arrows that go towards a $\Diamond\psi\wedge\Diamond\neg\psi$ world within i steps. The difficult part is to make sure that we can recognize in $\mathcal{M}_{[U^{\psi}_{3.i}]}$ whether or not a world w_1 satisfied $\Diamond\psi \wedge \Diamond\neg\psi$ in $\mathcal{M}$.

The main instrument for doing this will be a formula we give the name ϑ,

$$\vartheta := \Diamond\Box\bot \wedge \Diamond\Diamond\top.$$

The idea is that we guarantee that $\neg\psi$ worlds have successors in $\mathcal{M}_{[U^{\psi}_{3.i}]}$ while ψ worlds do not. This way, if $\mathcal{M}, w_1 \models \Diamond\psi \wedge \Diamond\neg\psi$ then $\mathcal{M}_{[U^{\psi}_{3.i}]}, w_1 \models \vartheta$. In order guarantee that $\neg\psi$ worlds have successors we need another case distinction. Let

$$\tau^{\psi}_i := \Diamond_U \top \vee \hat{\Diamond}^i(\Diamond\psi \wedge \Diamond\neg\psi).$$

We now construct two updates, $[U^{\psi,+}_{3.i}]$ and $[U^{\psi,-}_{3.i}]$. The update $[U^{\psi,+}_{3.i}]$ will give the right result in case the $\neg\psi$ successor of w_1 satisfies τ^{ψ}_i, the update $[U^{\psi,-}_{3.i}]$ will give the right result in case it does not.

A.3.3 Constructing $U^{\psi.\pm}_{3.i}$

First the + case. Let

$$U^{\psi.+}_{3.0} := U \cup \{(\Diamond\psi \wedge \Diamond\neg\psi, \mathcal{A}, \top), \overline{(\psi, \mathcal{A}, \top)}\}$$

and

$$U^{\psi.+}_{3.i+1} := U^{\psi.+}_{3.i} \cup \{(\Diamond^{i+1}(\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{i}(\Diamond\psi \wedge \Diamond\neg\psi))\}.$$

The $(\Diamond^{j+1}(\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{j}(\Diamond\psi \wedge \Diamond\neg\psi))$ clauses make sure that if there is a world w_1 within i steps of a U-reachable world with $\mathcal{M}, w_1 \models \Diamond\psi \wedge \Diamond\neg\psi$ then this world w_1 will be reachable. The $(\Diamond\psi \wedge \Diamond\neg\psi, \mathcal{A}, \top)$ clause makes sure that both a $\neg\psi$ successor w_2 and a ψ successor w_3 of w_1 are reachable. If w_2 satisfies τ^{ψ}_{i} then either $U \subseteq U^{\psi.+}_{3.i}$ or $(\Diamond^{j+1}(\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{j}(\Diamond\psi \wedge \Diamond\neg\psi))$ will make sure that w_2 has a successor in $\mathcal{M}_{[U^{\psi.+}_{3.i}]}$. Finally, the $\overline{(\psi, \mathcal{A}, \top)}$ clause makes sure that w_3 has no successors in $\mathcal{M}_{[U^{\psi.+}_{3.i}]}$. We therefore have $\mathcal{M}_{[U^{\psi.+}_{3.i}]}, w_1 \models \vartheta$ and w_1 is reachable in $\mathcal{M}_{[U^{\psi.+}_{3.i}]}$.

The only problem that may arise is if there is a ψ world on the path to w_1. This possibility will be dealt with at a later stage.

Now let us consider the $-$ case. Let

$$\begin{aligned} U^{\psi.-}_{3.0} := &\{(u_1, a, u_2 \wedge \tau^{\psi}_{0}) \mid (u_1, a, u_2) \in U\} \cup \\ &\{(\neg\tau^{\psi}_{0}, \mathcal{A}, \top), (\Diamond\psi \wedge \Diamond\neg\psi, \mathcal{A}, \top), \overline{(\psi, \mathcal{A}, \top)}\} \end{aligned}$$

and for $0 < i \leq n$ let

$$\begin{aligned} U^{\psi.-}_{3.i} := &\{(u_1, a, u_2 \wedge \tau^{\psi}_{i}) \mid (u_1, a, u_2) \in U\} \cup \{(\neg\tau^{\psi}_{i}, \mathcal{A}, \top), (\Diamond\psi \wedge \Diamond\neg\psi, \mathcal{A}, \top), \\ &(\Diamond^{i}(\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)), \cdots, \\ &(\Diamond(\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond\psi \wedge \Diamond\neg\psi), \overline{(\psi, \mathcal{A}, \top)}\}. \end{aligned}$$

The $-$ case works much like the + case. The exception is that in this case the $\neg\psi$ successor w_2 of the $\Diamond\psi \wedge \Diamond\neg\psi$ world w_1 is assumed to satisfy $\neg\tau^{\psi}_{i}$, so it would not have a successor in $\mathcal{M}_{[U^{\psi.+}_{3.i}]}$. In this case the $(\neg\tau^{\psi}_{0}, \mathcal{A}, \top)$ clause however guarantees that as long as $\mathcal{M}, w_2 \models \Diamond\top$ we have $\mathcal{M}_{[U^{\psi.-}_{3.i}]}, w_2 \models \Diamond\top$. (The case $\mathcal{M}, w_2 \models \Box\bot$ will be dealt with later.)

This would cause problems because $(\neg\tau^{\psi}_{i}, \mathcal{A}, \top)$ could make worlds reachable that are neither U-reachable nor on the path to a $\Diamond\psi \wedge \Diamond\neg\psi$ world, but this is averted by putting the extra τ^{ψ}_{i} end condition in the clauses from U. While every arrow from a $\neg\tau^{\psi}_{i}$ world is retained the $\neg\tau^{\psi}_{i}$ worlds themselves are only reachable from $\Diamond\psi \wedge \Diamond\neg\psi$ worlds.

Again, problems may arise if there is a ψ world on the path to w_1 and in this case problems may also arise because of ϑ worlds appearing due to cutting links because of the new τ^{ψ}_{i} end condition. Both these problems will be dealt with in a later stage.

A.3.4 Constructing $\beta_{3.i}$

For $1 \leq i \leq n$ let

$$\beta_{3.0} := \Diamond_U \top \wedge \neg\beta_1 \wedge \neg\beta_2 \wedge \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} \left(\neg[U^{\psi.+}_{3.0}]C_{\mathcal{A}}\neg\vartheta \vee \neg[U^{\psi.-}_{3.0}]C_{\mathcal{A}}\neg\vartheta\right),$$

$$\gamma_{3.i} := \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} \left(\neg[U^{\psi.+}_{3.i}]C_{\mathcal{A}}\neg\vartheta \vee \neg[U^{\psi.-}_{3.i}]C_{\mathcal{A}}\neg\vartheta\right),$$

and

$$\beta_{3.i} := \Diamond_U \top \wedge \neg\beta_1 \wedge \neg\beta_2 \wedge \gamma_{3.i} \wedge \neg\gamma_{3.i-1}.$$

This $\beta_{3.i}$ is exactly what we were looking for.

Lemma 7. *For any model $\mathcal{M}$, any world w of $\mathcal{M}$ and any $0 < i \leq n$ we have that*

$$\mathcal{M}, w \models \neg\beta_1 \wedge \neg\beta_2 \wedge \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} \neg\{U\}^*\neg(\Diamond\psi \wedge \Diamond\neg\psi) \Leftrightarrow \mathcal{M}, w \models \beta_{3.0}$$

and

$$\mathcal{M}, w \models \neg\beta_1 \wedge \neg\beta_2 \wedge \neg\beta_{3.i-1} \wedge \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} \neg\{U\}^*\neg\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi) \Leftrightarrow \mathcal{M}, w \models \beta_{3.i}$$

To see why this is the case, first recall that a disjunct $\neg[U^{\psi.+}_{3.i}]C_{\mathcal{A}}\neg\vartheta \vee \neg[U^{\psi.-}_{3.i}]C_{\mathcal{A}}\neg\vartheta$ of $\beta_{3.i}$ has the same value as the disjunct $\neg\{U\}^*\neg\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$, unless one of two problems occurs.

We now show that both problems are solved by taking the disjunction over all formulas in $\Phi^{2n}_{\mathrm{Pvar}(\chi)}$. Consider the second of the problems, that the τ^ψ_i end condition in the $(u_1, a, u_2 \wedge \tau^\psi_i)$ clauses of $U^{\psi.-}_{3.i}$ can cause a ϑ world to come into existence in $\mathcal{M}_{[U^{\psi.-}_{3.i}]}$ where no $\Diamond\psi \wedge \Diamond\neg\psi$ world existed in $\mathcal{M}$. As a result, we could have $\neg[U^{\psi.-}_{3.i}]C_{\mathcal{A}}\neg\vartheta$ but $\neg\neg\{U\}^*\neg\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$.

Suppose we are in the situation where this happens. Then there is no U-reachable $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ world, so all reachable arrows that are retained by $U^{\psi.-}_{3.i}$ are U-arrows. We have a $U^{\psi.-}_{3.i}$-reachable world w_1 such that $\mathcal{M}_{[U^{\psi.-}_{3.i}]}, w_1 \models \vartheta$. Then w_1 has $U^{\psi.-}_{3.i}$-successors w_2 and w_3 such that $\mathcal{M}_{[U^{\psi.-}_{3.i}]}, w_2 \models \Diamond\top$ and $\mathcal{M}_{[U^{\psi.-}_{3.i}]}, w_3 \models \Box\bot$. The only clauses that could keep w_2 and w_3 reachable from w_1 are $(u_1, a, u_2 \wedge \tau^\psi_i)$ clauses, so in particular $\mathcal{M}, w_2 \models \tau^\psi_i$ and $\mathcal{M}, w_3 \models \tau^\psi_i$. The $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ disjunct of τ^ψ_i cannot hold so the other disjunct $\Diamond_U\top$ must hold in both worlds.

The successor of w_2 must also satisfy τ^ψ_i and therefore $\Diamond_U\top$. The world w_3 also has at least one U-successor but that arrow is not retained so all U-successors of w_3 satisfy $\neg\tau^\psi_i$ and therefore $\neg\Diamond_U\top$. But then $\mathcal{M}, w_2 \models \Diamond_U\Diamond_U\top$ and $\mathcal{M}, w_3 \models \neg\Diamond_U\Diamond_U\top$. The formula $\Diamond_U\Diamond_U\top$ is of depth at most $2n$ so

$\mathcal{M}, w_1 \models \Diamond\psi' \wedge \Diamond\neg\psi'$ for some $\psi' \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}$, so while we don't have $\mathcal{M}, w \models \neg\{U\}^*\neg\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ we do have $\mathcal{M}, w \models \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} \neg\{U\}^*\neg(\Diamond\psi \wedge \Diamond\neg\psi)$.

Consider then the first of the problems, that it is possible to have $\mathcal{M}, w \models \neg\{U\}^*\neg\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ without $\mathcal{M}, w \models \neg[U^{\psi,+}_{3.0}]C_{\mathcal{A}}\neg\vartheta \vee \neg[U^{\psi,-}_{3.0}]C_{\mathcal{A}}\neg\vartheta$ if there is a ψ world on the path from w to the $\Diamond\psi \wedge \Diamond\neg\psi$ world w_1.

It can be shown that in such a case there is a ψ' such that the world w_1 also satisfies $\Diamond\psi' \wedge \Diamond\neg\psi'$ and ψ' does not occur on the path from w to w'. So while we don't have $\mathcal{M}, w \models \neg[U^{\psi,+}_{3.0}]C_{\mathcal{A}}\neg\vartheta \vee \neg[U^{\psi,-}_{3.0}]C_{\mathcal{A}}\neg\vartheta$ we do have $\mathcal{M}, w \models \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} \left(\neg[U^{\psi,+}_{3.i}]C_{\mathcal{A}}\neg\vartheta \vee \neg[U^{\psi,-}_{3.i}]C_{\mathcal{A}}\neg\vartheta\right)$.

The proof of the existence of such a ψ' is rather long and technical and is therefore included in Section C of the Appendix as Lemma 21.

A.3.5 Constructing β_3

We can then simply take

$$\beta_3 := \bigvee_{0 \leq i \leq n} \beta_{3.i}.$$

The subcases are exhaustive of case 3, so $\models \delta_3 \leftrightarrow \beta_3$.

A.4 Constructing β_4

Recall that we are in case 4 if $\Diamond_U\top$, we are not in one of the previous cases and there is a U-reachable world where $\Diamond_{\overline{U}}\top$ holds. We have

$$\delta_4 = \Diamond_U\top \wedge \neg\delta_1 \wedge \neg\delta_2 \wedge \neg\delta_3 \wedge \neg\{U\}^*\neg\Diamond_{\overline{U}}\top$$

In this case, as mentioned in Section 6.4 we must have $\mathcal{M}, w_1 \models \psi \wedge \Diamond\neg\psi$ for some U-reachable world w_1, and this difference between w_1 world and its $\neg\psi$ successor w_2 must have one of the following causes:

1. there are two agents b_1, b_2 and two worlds w_3, w_4 such that $b_1 \neq b_2$, $\mathcal{M}, w_3 \models \Diamond_{b_1}\top$, $\mathcal{M}, w_4 \models \Diamond_{b_2}\top$ and both w_3 and w_4 are reachable from w_1 in at most $d(\psi)$ steps.

2. there are a propositional variable $p \in \mathrm{Pvar}(\psi)$ and two worlds w_3, w_4 such that $\mathcal{M}, w_3 \models p$, $\mathcal{M}, w_4 \models \neg p$ and both w_3 and w_4 are reachable from w_1 in at most $d(\psi) + 1$ steps.

3. there are a formula $\psi' \in \Phi^{d(\psi)}_{\mathrm{Pvar}(\psi)}$ and a world w_3 such that $\mathcal{M}, w_3 \models \Diamond\psi' \wedge \Diamond\neg\psi'$, w_3 is reachable from w_1 in at most k steps and $k + d(\psi') \leq d(\psi)$.

4. there is a $k \leq d(\psi)$ such that $\mathcal{M}, w_1 \models \Diamond^k\Box\bot \wedge \Box^{k+1}\bot$.

The first three possibilities mentioned correspond to the first three cases of our case distinction. If we are in case 4 we are however by definition not in one of these case so the fourth possibility must hold.

Let

$$\beta_4 := \Diamond_U\top \wedge \neg\beta_1 \wedge \neg\beta_2 \wedge \neg\beta_3 \wedge \neg C_{\mathcal{A}}\Box_{\overline{U}}\bot \wedge C_{\mathcal{A}}(\Diamond_{\overline{U}}\top \to \bigvee_{1\leq i\leq n} \Diamond^i\Box\bot \wedge \Box^{i+1}\bot).$$

Then $\models \delta_4 \leftrightarrow \beta_4$. To see why this is the case note that everywhere where a U-arrow is followed by an $\overline{U}$-arrow we have $\Diamond^i\Box\bot \wedge \Box^{i+1}\bot$ for some $1 \leq i \leq n$ and that every successor of a $\Diamond^i\Box\bot \wedge \Box^{i+1}\bot$ world satisfies $\Diamond^{i-1}\Box\bot \wedge \Box^i\bot$.

A.5 Constructing β_5 and β_6

Finding an AUC description for the last two cases is trivial. Case 5 is the case where there are no $\overline{U}$-arrows departing from U-reachable worlds, so

$$\delta_5 = \Diamond_U\top \wedge \neg\delta_1 \wedge \neg\delta_2 \wedge \neg\delta_3 \wedge \neg\delta_4 \wedge \{U\}^*\neg\Diamond_{\overline{U}}\top$$

We can take

$$\beta_5 := \Diamond_U\top \wedge C_{\mathcal{A}}\neg\Diamond_{\overline{U}}\top.$$

Case 6 is the case where $\neg\Diamond_U\top$ holds,

$$\delta_6 := \neg\Diamond_U\top.$$

The formula δ_6 is itself already an AUC formula so we can simply take

$$\beta_6 := \delta_6.$$

B Constructing $\alpha_6, \cdots, \alpha_1$

In order to prove Lemma 5 we have to show that there are $\alpha_6, \cdots, \alpha_1$ such that $\models (\delta_i \wedge \chi) \leftrightarrow (\delta_i \wedge \alpha_i)$ for all $6 \geq i \geq 1$. Here we construct the α_i.

Cases 6, 5 and 4 can be solved quite easily and directly. Cases 3, 2 and 1 are more difficult and are solved in the way described in Section 6.3: by cutting at $\neg\varphi$ worlds and then checking whether the witnesses are still reachable.

It is convenient to have a formula representing "the solution for all later cases". The letter ζ (with various indices) will be used to represent this.

B.1 Constructing α_6 and α_5

Cases 6 and 5 are very simple cases, mostly because in both cases it is impossible to go to a different case by taking a U-arrow.

Case 6 is the case where there are no outgoing U-arrows worlds. We can therefore take

$$\alpha_6 := \varphi.$$

Case 5 is the case where there is no U-reachable world with an $\overline{U}$ arrow departing from it. Every reachable world is therefore U-reachable, so we can take

$$\alpha_5 := C_{\mathcal{A}}\varphi.$$

B.2 Constructing α_4

In case 4 the only possible cause of two worlds being distinguishable (by a formula of depth at most n) is that one of them satisfies $\Diamond^j\Box\bot \wedge \Box^{j+1}\bot$ with $j < n$ and the other does not. For $0 \leq i < n$ let $\sigma_i := \Diamond^i\Box\bot \wedge \Box^{i+1}\bot$ and let $\sigma_n := \bigwedge_{0\leq i<n} \neg\sigma_i$. Note that arrows can only go either from a σ_n world to another σ_n world or from a σ_{j+1} world to a σ_j world.

The depth of φ is less than n, so whether it holds in a world is fully determined by which of the σ_i holds. Furthermore, since we are in case 4 we know that there are both a reachable $\Diamond_U\top$ world and a reachable $\Diamond_{\overline{U}}$ world, so there must be a U-reachable world satisfying σ_i with $i < n$.

Let w be a world and let k be the index such that $\mathcal{M}, w \models \sigma_k$. It is fixed by U and φ at which indices there is no outgoing U-arrow and at which indices $\neg\varphi$ holds, so k fully determines whether χ holds. We can take

$$\alpha_4 := \bigwedge_{0\leq i\leq n} (\sigma_i \to \lambda_i)$$

where for each $0 \leq i \leq n$ the formula λ_i is either $\top$ or $\bot$, as determined by U and φ.

B.3 Constructing α_3

The detecting formula β_3 was the most complicated of the detecting formulas. Likewise, α_3 is the most complicated of the solving formulas. Here for the first time we need the fact that we are working backward through the cases so we can use solutions to later cases in earlier ones. Let

$$\zeta_{3.n+1} := (\beta_6 \to \alpha_6) \wedge (\beta_5 \to \alpha_5) \wedge (\beta_4 \to \alpha_4).$$

Now suppose that for some $1 \leq i \leq n$ we are in case $\beta_{3.i}$ and the later cases have been solved, so $\zeta_{3.i+1}$ is already defined.

We want to find a formula $\alpha_{3.i}$ that detects whether there are U-reachable $\neg\varphi$ worlds. In order to do this we also consider whether there are U-reachable $\neg\varphi'$ worlds with $\varphi' := \varphi \wedge \Box_U\varphi$. Doing this will allow us to ignore "side paths" that are only a single world long in stage $\alpha_{3.i.3}$, because if there is a $\neg\varphi$ world on a single world "side path" then there is also a $\neg\varphi'$ world on the "main path". We can safely consider φ' instead of φ since $\models \{U\}\varphi \leftrightarrow \{U\}\varphi'$.

We split this into three parts; a formula $\alpha_{3.i.1}$ that detects whether there is a $\neg\varphi'$ world on every path towards a $\Diamond\psi\wedge\Diamond\neg\psi$ world, a formula $\alpha_{3.i.2}$ that detects whether there are $\neg\varphi'$ worlds on some but not all paths towards $\Diamond\psi\wedge\Diamond\neg\psi$ worlds and a formula $\alpha_{3.i.3}$ that detect whether there are U-reachable worlds that are in a later case and satisfy $\neg\varphi$ (note the lack of a $'$ on the φ here). Unfortunately we need to split the first of these cases into even more subcases, depending on whether the $\neg\psi$ successor of the $\Diamond\psi \wedge \Diamond\neg\psi$ world satisfies φ.

B.3.1 Constructing $\alpha_{3.i.1}$

For $0 \leq i \leq n$ let

$$U^{\psi.+}_{3.i.1.1} := U^{\psi.+}_{3.i} \cup \{\overline{(\neg\varphi' \wedge \neg\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \top)}\},$$

$$U^{\psi.+}_{3.i.1.2} := U^{\psi.+}_{3.i} \cup \{\overline{(\varphi' \wedge \neg\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \neg\varphi' \wedge \neg\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi))}\}$$

and

$$U^{\psi.-}_{3.i.1} := U^{\psi.-}_{3.i} \cup \{\overline{(u_1 \wedge \neg\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi), a, u_2 \wedge \neg\varphi')} \mid (u_1, a, u_2) \in U\}.$$

Having defined these updates let

$$\alpha_{3.i.1} := \varphi' \wedge \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} \left(\neg[U^{\psi.+}_{3.i.1.1}]C_{\mathcal{A}}\neg\vartheta \vee \neg[U^{\psi.+}_{3.i.1.2}]C_{\mathcal{A}}\neg\vartheta \vee \neg[U^{\psi.-}_{3.i.1}]C_{\mathcal{A}}\neg\vartheta\right).$$

If we are in case 3.i then the formula $\alpha_{3.i.1}$ is equivalent to there being at least one U-path to a $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ world that does not pass through a $\neg\varphi'$ world, as can be seen from the following two lemmas.

Lemma 8. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{3.i}$ we have*

$$\mathcal{M}, w \models \chi \Rightarrow \mathcal{M}, w \models \alpha_{3.i.1}$$

Proof. We are in case $3.i$ so there is a nearby world where $\Diamond\psi \wedge \Diamond\neg\psi$ holds. There are three possibilities for this world, see Figure 12.

Let w_5 be the U-reachable world with $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$, w_1 the $\Diamond\psi \wedge \Diamond\neg\psi$ world, w_2 a $\neg\psi$ successor of w_1, w_3 a ψ successor of w_1 and w_4 a successor or w_2.

In all cases the arrows up to and including the one departing from w_1 are not cut by $U^{\psi.+}_{3.i.1.1}$, $U^{\psi.+}_{3.i.1.2}$ or $U^{\psi.-}_{3.i.1}$. In order to see why this is the case note that the worlds up to w_5 satisfy φ' by the assumption that the U-path is $\neg\varphi'$ free and the worlds from the successor of w_5 to w_1 satisfy $\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)$. The updates $U^{\psi.+}_{3.i}$ and $U^{\psi.-}_{3.i}$ do not cut these arrows and the new clauses cannot apply due to the worlds satisfying $\neg\varphi'$ or $\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)$. The only arrow we need to retain that is in danger of being cut is the one from w_2 to w_4.

The first case, see Figure 12a, is if the arrow from w_2 to w_4 is a U-arrow and $\mathcal{M}, w_2 \models \varphi'$. The arrow from w_2 to w_4 is in this case retained by $[U^{\psi.+}_{3.i.1.1}]$ because the new clause in that update only removes arrows starting from $\neg\varphi'$ worlds.

The second case, see Figure 12b, is if the arrow from w_2 to w_4 is a U-arrow and $\mathcal{M}, w_2 \models \neg\varphi'$. The arrow from w_2 to w_4 is in this case retained by $[U^{\psi.+}_{3.i.1.2}]$ because the new clause in that update only removes arrows starting from φ' worlds.

The third case, see Figure 12c, is if the arrow from w_2 to w_4 is not a U-arrow. The arrow from w_2 to w_4 is in this case retained by $[U^{\psi.-}_{3.i.1}]$ because the new clause in that update only removes U-arrows.

In any case at least one of the updates in $\alpha_{3.i.1}$ will cause w_1 to become a reachable ϑ world, so $\alpha_{3.i.1}$ holds. $\square$

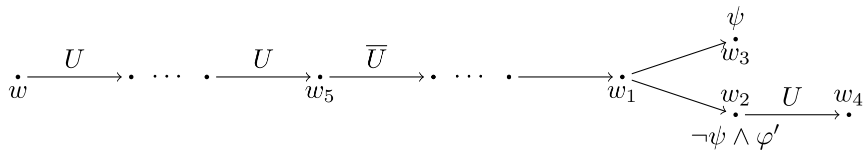

(a) Possibility 1, solved by $U^{\psi.+}_{3.i.1.1}$

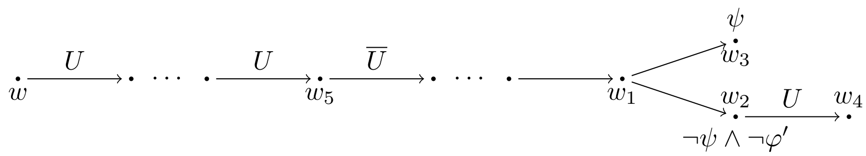

(b) Possibility 2, solved by $U^{\psi.+}_{3.i.1.2}$

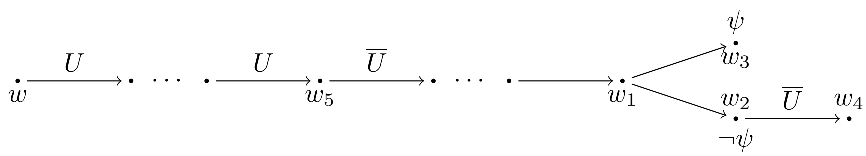

(c) Possibility 3, solved by $U^{\psi.-}_{3.i.1}$

Figure 12: Three possibilities for a $\Diamond\psi \wedge \Diamond\neg\psi$ world.

Lemma 9. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{3.i}$ we have that if every U-path to a $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ world passes through a $\neg\varphi'$ world then $\mathcal{M}, w \not\models \alpha_{3.i.1}$.*

Proof. It should be immediately clear from the definitions of $U^{\psi.+}_{3.i.1.1}$, $U^{\psi.+}_{3.i.1.2}$ or $U^{\psi.-}_{3.i.1}$ that updating with any of them will make the $\Diamond\psi \wedge \Diamond\neg\psi$ worlds unreachable.[10] The only possibility for $\alpha_{3.i.1}$ to hold would therefore be if one of the updates would cause a new ϑ world to come into existence by cutting at $\neg\varphi'$ worlds.

This however would require a U-reachable world w' with $\mathcal{M}, w' \models \Diamond[V]\top \wedge \Diamond\neg[V]\top$ with V a singleton update of one of the new clauses. But such an update $[V]$ is of depth less than $2n$ so this conflicts with us being in case $3.i$ if $i > 0$ and with no $\Diamond\psi \wedge \Diamond\neg\psi$ being U-reachable without passing a $\neg\varphi'$ world if $i = 0$. $\square$

[10] There is one exception to this, if the only $\neg\varphi'$ world on the path is w itself. This is excluded by the φ' conjunct of $\alpha_{3.i.1}$, however.

B.3.2 Constructing $\alpha_{3.i.2}$ for $i > 0$

For $1 \leq i \leq n$ let

$$
\begin{aligned}
U^+_{3.i.2} := {} & \{(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i}) \mid (u_1, a, u_2) \in U\} \cup \{\overline{(\beta_{3.i} \wedge \neg\alpha_{3.i.1}, \mathcal{A}, \top)}, \\
& (\beta_{3.i} \wedge \Box_U \neg\beta_{3.i} \wedge \Diamond^i \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{i-1} \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi))\}, \\
U^-_{3.i.2} := {} & \{(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i}) \mid (u_1, a, u_2) \in U\} \cup \\
& \{(\beta_{3.i} \wedge \Box_U \neg(\beta_{3.i} \wedge \alpha_{3.i.1}) \wedge \Diamond^i \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi), \\
& \mathcal{A}, \Diamond^{i-1} \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi))\}
\end{aligned}
$$

and

$$\alpha_{3.i.2} := [U^+_{3.i.2}]C_{\mathcal{A}}\neg\vartheta \wedge [U^-_{3.i.2}]C_{\mathcal{A}}\neg\vartheta.$$

Before discussing why $\alpha_{3.i.2}$ detects whether there are $\neg\varphi'$ worlds on some but not all paths to $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ worlds, it is important to note that $\alpha_{3.i.2}$ works very differently from $\alpha_{3.i.1}$. In $\alpha_{3.i.1}$ an update is used that guarantees that there is *at least one reachable ϑ world* if χ holds, whereas $\alpha_{3.i.2}$ uses an update that guarantees that there is *no reachable ϑ world* if χ holds.

Now, to see why $\alpha_{3.i.2}$ works. There are two parts to this. The first is that if there is no U-reachable $\neg\varphi'$ world then $\alpha_{3.i.2}$ holds. The second is that if there are $\neg\varphi'$ worlds on some but not all U-paths to $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ worlds then $\alpha_{3.i.2}$ does not hold.

Lemma 10. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{3.i}$ we have*

$$\mathcal{M}, w \models \chi \Rightarrow \mathcal{M}, w \models \alpha_{3.i.2}$$

Proof. First, note that the $(\beta_{3.i} \wedge \Box_U \neg\beta_{3.i} \wedge \Diamond^i \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{i-1} \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi))$ clause in the $+$ update and the corresponding clause in the $-$ update only retain arrows from $\beta_{3.i}$ to $\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)$ worlds. Such target worlds are too close to $\Diamond\psi \wedge \Diamond\neg\psi$ to be $\beta_{3.i}$ worlds; they could be $\beta_{3.i-1}$ at the most.

The $(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i})$ clauses only retain U-arrows from $\beta_{3.i}$ worlds to $\beta_{3.i}$ worlds. Every worlds that is still reachable after the update was therefore originally either U- and $\beta_{3.i}$-reachable or the successor of such a world.

From $\mathcal{M}, w \models \chi$ it follows that the conjunct $\Box_U \neg\beta_{3.i}$ in the start condition of the final clause of the $+$ update and the conjunct $\Box_U \neg(\beta_{3.i} \wedge \alpha_{3.i.1})$ in the start condition of the final clause of the $-$ update hold in the same U-reachable worlds, so the two clauses retain the same arrows (when restricting ourselves to the relevant parts of the model, the parts that are still connected to w after the update). Also, by Lemma 8 and the fact that $\mathcal{M}, w \models \chi$ we have that the start condition of the clause $\overline{(\beta_{3.i} \wedge \neg\alpha_{3.i.1}, \mathcal{A}, \top)}$ of the + update cannot hold in any relevant world.

In order to show that the Lemma holds it therefore suffices to show that for

$$\begin{aligned} U'_{3.i.2} := \{&(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i}) \mid (u_1, a, u_2) \in U\} \cup \\ &\{(\beta_{3.i} \wedge \Box_U \neg\beta_{3.i} \wedge \Diamond^i \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{i-1} \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi))\} \end{aligned}$$

we have $\mathcal{M}, w \models [U'_{3.i.2}]C_{\mathcal{A}}\neg\vartheta$.

Every $\beta_{3.i}$ world either has a $\beta_{3.i}$ successor or is a $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ world for some $\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}$. If it has a $\beta_{3.i}$ successor the arrow to that successor is retained by the $(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i})$ clauses. If it does not have a $\beta_{3.i}$ successor then the

$$(\beta_{3.i} \wedge \Box_U \neg\beta_{3.i} \wedge \Diamond^i \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{i-1} \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi))$$

clause retains the arrow to a $\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)$ successor. As mentioned before this $\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)$ is not itself a $\beta_{3.i}$ world since it is too close to a $(\Diamond\psi \wedge \Diamond\neg\psi)$ world so none of its outgoing arrows are retained.

The result is that each reachable $\Box\bot$ world in $\mathcal{M}_{[U'_{3.i.2}]}$ is a $\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)$ world, so each $\Diamond\Box\bot$ world is a $\beta_{3.i} \wedge \Box_U \neg\beta_{3.i}$ world and therefore a $\Box\Box\bot$. A $\Box\Box\bot$ world cannot satisfy ϑ, so this proves the Lemma. $\Box$

Now for the other part. Here we need an extra assumption in the lemma, namely that $\mathcal{M}, w \models \alpha_{3.i.2}$. This assumption is harmless: if $\mathcal{M}, w \not\models \alpha_{3.i.2}$ then we already know that $\mathcal{M}, w \not\models \chi$.

Lemma 11. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{3.i}$ we have that if*

- *there is a U-reachable world w' with $\mathcal{M}, w' \models \beta_{3.i} \wedge \neg\varphi'$ and*
- $\mathcal{M}, w \models \alpha_{3.i.1}$

then $\mathcal{M}, w \models \neg\alpha_{3.i.2}$.

Proof. Fix any world w' satisfying the condition of the Lemma. From $\mathcal{M}, w' \models \beta_{3.i} \wedge \neg\varphi'$ it follows that $\mathcal{M}, w' \models \neg\alpha_{3.i.1}$. Let w_1 be the first world on a U-path from w to w' that is a $\neg\alpha_{3.i.1}$ world. In particular this implies that there are no $\neg\varphi'$ worlds on the path before w_1.

We have $\mathcal{M}, w \models \alpha_{3.i.1}$ and $\mathcal{M}, w_1 \not\models \alpha_{3.i.1}$ so $w \neq w_1$. This guarantees the existence of a predecessor w_2 of w_1 on the U-path.

Now there are two possibilities for the predecessor w_2 of w_1 on the path. The first is that w_2 has a successor w_3 with $\mathcal{M}, w_3 \models \alpha_{3.i.1}$. The second possibility is that w_2 has no such successor.

The formula $\alpha_{3.i.1}$ holds if and only if there is a $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ world that is reachable without passing a $\neg\varphi'$ world. So if w_2 is a $\alpha_{3.i.2}$ world but has no $\alpha_{3.i.2}$ successor then it must itself be a $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ world. In this case w_2

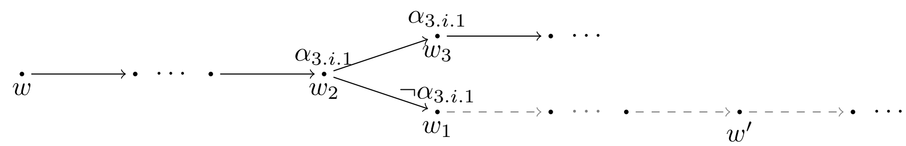

(a) Possibility 1: w_1 has a successor w_3 with $\mathcal{M}, w_3 \models \alpha_{3.i.1}$. We have $\mathcal{M}, w \not\models [U^+_{3.i.2}]C_{\mathcal{A}}\neg\vartheta$.

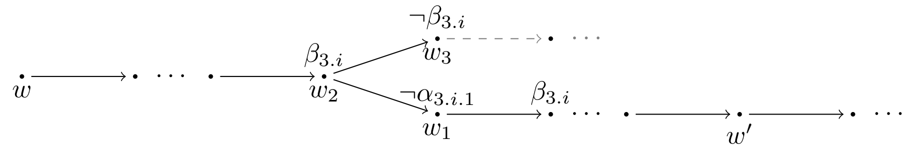

(b) Possibility 2: $\mathcal{M}, w_1 \models \Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ for some $\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}$. We have $\mathcal{M}, w \not\models [U^-_{3.i.2}]C_{\mathcal{A}}\neg\vartheta$.

Figure 13: The two possibilities for the conditions of Lemma 11 to hold.

has a successor w_3 with $\mathcal{M}, w_3 \models \Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)$ and therefore $\mathcal{M}, w_3 \not\models \beta_{3.i}$. The two cases are shown in Figure 13. There may be more arrows than the ones shown in the figure but such arrows do not matter as long as the arrows that are shown exist.

In the first case consider the update $U^+_{3.i.2}$ as shown in Figure 13a. Arrows that are not retained are drawn in gray and dashed. The arrows from w_2 to w_1 and w_3 are retained because they are $(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i})$ arrows. The arrow from w_3 to at least one of its successors is also retained by some clause—which one depends on whether w_3 has a $\alpha_{3.i.1}$ successor. The arrows from w_2 to its successors are not retained, because of the $\overline{(\beta_{3.i} \wedge \neg\alpha_{3.i.1}, \mathcal{A}, \top)}$ clause. We therefore have $\mathcal{M}_{[U^+_{3.i.2}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.2}$.

In the second case, shown in Figure 13b, consider the update $U^-_{3.i.2}$. The arrows from w_2 to w_1 and from w_1 to its successor are retained because of the $(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i})$ clauses. The arrow from w_2 to w_3 is retained because it is a $(\beta_{3.i} \wedge \Box_U \neg(\beta_{3.i} \wedge \alpha_{3.i.1}) \wedge \Diamond^i \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{i-1} \bigvee_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} (\Diamond\psi \wedge \Diamond\neg\psi)$ arrow. We therefore have $\mathcal{M}_{[U^+_{3.i.2}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.2}$. $\square$

B.3.3 Constructing $\alpha_{3.0.2}$

The formula $\alpha_{3.0.2}$ is very similar to $\alpha_{3.i.2}$, except that there is one more special case. In $\alpha_{3.i.2}$ we could quite easily guarantee that any $\alpha_{3.i.1}$ has a successor after the update by allowing arrows from $\Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ to $\Diamond^{i-1}(\Diamond\psi \wedge \Diamond\neg\psi)$ in case there is no $\beta_{3.i}$ successor. Doing the same for $\alpha_{3.0.2}$ and $\alpha_{3.0.1}$ is not possible. This essentially forces us to use two cases for what was the + case in $\alpha_{3.i.2}$. For similar reasons we split the $-$ case into two cases, both of which are also indexed by ψ.

Let

$$
\begin{aligned}
U_{3.0.2}^{+.1} := \{&(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0}) \mid (u_1, a, u_2) \in U\} \cup \{\overline{(\beta_{3.0} \wedge \neg\alpha_{3.0.1}, \mathcal{A}, \top)}, \\
&\overline{(\Diamond_U(\beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0}) \wedge \Diamond_U(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U\beta_{3.0}), \mathcal{A}, \beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0})}\}, \\
U_{3.0.2}^{+.2} := \{&(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0}) \mid (u_1, a, u_2) \in U\} \cup \{(\beta_{3.0} \wedge \neg\alpha_{3.0.1}, \mathcal{A}, \top), \\
&\overline{(\Diamond_U(\beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0}) \wedge \Diamond_U(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U\beta_{3.0}), \mathcal{A}, \beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0})}\}, \\
U_{3.0.2}^{-.\psi.1} := \{&(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0}) \mid (u_1, a, u_2) \in U\} \cup \\
&\{\overline{(\Diamond_U(\beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0}) \wedge \Diamond_U(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U\beta_{3.0}), \mathcal{A}, \beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0})} \\
&(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U\neg\alpha_{3.0.1}, \mathcal{A}, \neg\psi), \overline{(\beta_{3.0} \wedge \neg\alpha_{3.0.1} \wedge \psi, \mathcal{A}, \top)}\}, \\
U_{3.0.2}^{-.\psi.2} := \{&(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0}) \mid (u_1, a, u_2) \in U\} \cup \\
&\{\overline{(\Diamond_U(\beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0}) \wedge \Diamond_U(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U\beta_{3.0}), \mathcal{A}, \beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0})} \\
&(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U\neg\alpha_{3.0.1}, \mathcal{A}, \neg\psi), (\beta_{3.0} \wedge \neg\alpha_{3.0.1} \wedge \psi, \mathcal{A}, \top)\}
\end{aligned}
$$

and

$$\alpha_{3.0.2} := [U_{3.0.2}^{+.1}]C_{\mathcal{A}}\neg\vartheta \wedge [U_{3.0.2}^{+.2}]C_{\mathcal{A}}\neg\vartheta \wedge \bigwedge_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} ([U_{3.0.2}^{-.\psi.1}]C_{\mathcal{A}}\neg\vartheta \wedge [U_{3.0.2}^{-.\psi.2}]C_{\mathcal{A}}\neg\vartheta).$$

Lemma 12. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{3.0}$ we have*

$$\mathcal{M}, w \models \chi \Rightarrow \mathcal{M}, w \models \alpha_{3.0.2}$$

Proof. As in the $\alpha_{3.i.2}$ case it follows from $\mathcal{M}, w \models \chi$ that terms containing $\neg\alpha_{3.0.2}$ cannot be relevant. All four updates then simplify to

$$
\begin{aligned}
U'_{3.0.1} := \{&(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0}) \mid (u_1, a, u_2) \in U\} \cup \\
&\{\overline{(\Diamond_U(\beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0}) \wedge \Diamond_U(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U\beta_{3.0}), \mathcal{A}, \beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0})}\}.
\end{aligned}
$$

Suppose now that there is a world w_1 that is reachable after the update and that satisfies $\mathcal{M}_{[U'_{3.0.2}]}, w_1 \models \vartheta$. The update retains only arrows that were U-arrows from and to $\beta_{3.0}$ worlds, so $\mathcal{M}, w_1 \models \beta_{3.0}$. Now consider the successors of w_1, worlds w_2 and w_3 that are reachable from w_1 after the update such that w_2 has a successor in $\mathcal{M}_{[U'_{3.0.2}]}$ and w_3 does not.

The arrow from w_2 to its successor must be a U-arrow from a $\beta_{3.0}$ world to a $\beta_{3.0}$ world, so $\mathcal{M}, w_1 \models \Diamond_U(\beta_{3.0} \wedge \Diamond_U\beta_{3.0})$. The arrow from w_3 to its successors on the other hand cannot be U-arrows to a $\beta_{3.0}$ world as they would then be retained by the update. We therefore have $\mathcal{M}, w_3 \models \beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0}$. But then $\mathcal{M}, w_1 \models \Diamond_U(\beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0})$, so the arrow from w_1 to w_3 is cut by the $\overline{(\Diamond_U(\beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0}) \wedge \Diamond_U(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U\beta_{3.0}), \mathcal{A}, \beta_{3.0} \wedge \neg\Diamond_U\beta_{3.0})}$ clause.

This contradicts w_3 being reachable from w_1 after the update, so such a world w_1 cannot exist which proves the Lemma. ☐

Now for the other part. Again, we need a harmless extra condition, namely that $\mathcal{M}, w \models \alpha_{3.0.1}$.

Lemma 13. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{3.0}$ we have that if*

- *there is a U-reachable world w' with $\mathcal{M}, w' \models \beta_{3.0} \wedge \neg\varphi'$ and*
- *$\mathcal{M}, w \models \alpha_{3.0.1}$*

then $\mathcal{M}, w \models \neg\alpha_{3.0.2}$.

Proof. Fix any world w' satisfying the condition of the Lemma. From $M, w' \models \beta_{3.0} \wedge \neg\varphi'$ it follows that $\mathcal{M}, w' \models \neg\alpha_{3.0.1}$. Let w_1 be the first world on a U-path from w to w' that is a $\neg\alpha_{3.0.1}$ world. In particular this implies that there are no $\neg\varphi'$ worlds on the path before w_1.

Now let w_2 be the predecessor of w_1 on the path. There are four possibilities for the situation around w_2. The first possibility is that w_2 has a U-successor w_3 satisfying $\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U \beta_{3.0}$. The second possibility is that w_2 has no successor of the kind in case 1, but does have a successor w_3 satisfying $\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \neg\Diamond_U \beta_{3.0}$.

In the third and fourth possibilities w_2 has no successor satisfying $\beta_{3.0} \wedge \alpha_{3.0.1}$. From $\mathcal{M}, w_2 \models \alpha_{3.0.1}$ it follows that there is some $\Diamond\psi \wedge \Diamond\neg\psi$ world that is reachable from w_2 without passing over a $\neg\varphi$ world. If none of the successors of w_2 satisfy $\alpha_{3.0.1}$ this implies that w_2 must itself be a $\Diamond\psi \wedge \Diamond\neg\psi$ world for some $\psi \in \Phi^{3n}_{\mathrm{Pvar}(\chi)}$.

By negating this it if necessary we can take this ψ such that $\mathcal{M}, w_1 \models \psi$. Let $U'_{3.0.2} := U^{-.\psi.1}_{3.0.2} \setminus \{\overline{(\beta_{3.0} \wedge \neg\alpha_{3.0.1} \wedge \psi, \mathcal{A}, \top)}\} = U^{-.\psi.2}_{3.0.2} \setminus \{(\beta_{3.0} \wedge \neg\alpha_{3.0.1} \wedge \psi, \mathcal{A}, \top)\}$. The difference between the third and fourth case now is whether the $\neg\psi$ successor w_3 of w_2 has a successor in $\mathcal{M}_{[U'_{3.0.2}]}$. If it does we are in case 3, if it does not we are in case 4.

Note that because $\mathcal{M}, w_3 \models \neg\psi$ the $\overline{(\beta_{3.0} \wedge \neg\alpha_{3.0.1} \wedge \psi, \mathcal{A}, \top)}$ and $(\beta_{3.0} \wedge \neg\alpha_{3.0.1} \wedge \psi, \mathcal{A}, \top)$ clauses cannot apply to arrows from w_3. This means that if we are in case 3 then w_3 has a successor in $\mathcal{M}_{[U^{-.\psi.1}_{3.0.2}]}$ and if we are in case 4 then w_3 has no successor in $\mathcal{M}_{[U^{-.\psi.2}_{3.0.2}]}$.

The four different cases are shown in Figure 14. There may be more arrows than the ones shown in the figure but such arrows do not matter as long as the arrows that are shown exist. Arrows that are not retained are drawn in gray and dashed.

In the first case consider the update $U^{+.1}_{3.0.2}$ as shown in Figure 14a. The arrows from w_2 to w_1 and w_3 and the arrow from w_3 to its successor are retained because they are $(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0})$ arrows and, because $\mathcal{M}, w_2 \models \alpha_{3.0.1}$ and $\mathcal{M}, w_3 \models \alpha_{3.0.1}$, not $\overline{(\beta_{3.0} \wedge \neg\alpha_{3.0.1}, \mathcal{A}, \top)}$ arrows. The arrows from w_1 to its successors are not retained because they are $\overline{(\beta_{3.0} \wedge \neg\alpha_{3.0.1}, \mathcal{A}, \top)}$ arrows. We therefore have $\mathcal{M}_{[U^{+.1}_{3.0.2}]}, w \models \neg C_{\mathcal{A}} \neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.0.2}$.

In the second case consider the update $U^{+.2}_{3.0.2}$ as shown in Figure 14b. The arrows from w_2 to w_3 and w_1 are retained because they are $(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0})$ arrows. Arrows from w_1 to its successors (which must exist because $\beta_{3.0}$ holds in every world on the path to w', and therefore in particular on w_1) are retained because they are $(\beta_{3.0} \wedge \neg\alpha_{3.0.1}, \mathcal{A}, \top)$ arrows. Arrows from w_3 to its successors are not retained; they are not $(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0})$ arrows because the successors of w_3 are not $\beta_{3.0}$ worlds and they are not $(\beta_{3.0} \wedge \neg\alpha_{3.0.1}, \mathcal{A}, \top)$ arrows because $\mathcal{M}, w_3 \models \alpha_{3.0.1}$. We therefore have $\mathcal{M}_{[U^{+.2}_{3.0.2}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.0.2}$.

In the third case consider the update $U^{-.\psi.1}_{3.0.2}$ as shown in Figure 14c. The arrow from w_2 to w_1 is retained because it is an $(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0})$ arrow. The arrow from w_2 to w_3 is retained because it is an $(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U \neg\alpha_{3.0.1}, \mathcal{A}, \neg\psi)$ arrow. The arrow from w_3 to at least one of its successors is retained because by assumption it has a successor in $\mathcal{M}_{[U^{-.\psi.1}_{3.0.2}]}$. The arrows from w_1 to its successors are not retained because they are $(\beta_{3.0} \wedge \neg\alpha_{3.0.1} \wedge \psi, \mathcal{A}, \top)$ arrows. We therefore have $\mathcal{M}_{[U^{-.\psi.1}_{3.0.2}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.0.2}$.

In the fourth case consider the update $U^{-.\psi.2}_{3.0.2}$ as shown in Figure 14d. The arrow from w_2 to w_1 is retained because it is an $(u_1 \wedge \beta_{3.0}, a, u_2 \wedge \beta_{3.0})$ arrow. The arrow from w_1 to w_3 is retained because it is an $(\beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U \neg\alpha_{3.0.1}, \mathcal{A}, \neg\psi)$ arrow. The arrows from w_3 to its successors are not retained because by assumption it has no successors in $\mathcal{M}_{[U^{-.\psi.2}_{3.0.2}]}$. The arrows from w_1 to its successors are retained because they are $(\beta_{3.0} \wedge \neg\alpha_{3.0.1} \wedge \psi, \mathcal{A}, \top)$ arrows. We therefore have $\mathcal{M}_{[U^{-.\psi.2}_{3.0.2}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.0.2}$.

These four cases are exhaustive so this proves the Lemma. $\square$

B.3.4 Constructing $\alpha_{3.i.3}$

The formula $\alpha_{3.i.3}$ should find $\neg\varphi$ worlds that are in cases later than $3.i$, with the possible exception of $\neg\varphi$ worlds that are successors of $\beta_{3.i}$ worlds, as the predecessors of these $\neg\varphi$ worlds have already been detected as $\neg\varphi'$ worlds by $\alpha_{3.i.1}$ or $\alpha_{3.i.2}$.

For $0 \leq i \leq n$ and $\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}$ let

$$\begin{aligned} U^{\psi}_{3.i.3} := \{&(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i}) \mid (u_1, a, u_2) \in U\} \cup \\ &\{(\Diamond_U(\beta_{3.i} \wedge \neg\Diamond_U \beta_{3.i}) \wedge \Diamond_U(\beta_{3.i} \wedge \Diamond_U \beta_{3.i}), \mathcal{A}, \beta_{3.i} \wedge \neg\Diamond_U(\beta_{3.i} \vee \zeta_{3.i+1}))\} \cup \\ &\{(\beta_{3.i} \wedge \Diamond_U(\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1}), \mathcal{A}, \top), (\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1} \wedge \psi, \mathcal{A}, \top)\}\end{aligned}$$

and

$$\alpha_{3.i.3} := \bigwedge_{\psi \in \Phi^{2n}_{\mathrm{Pvar}(\chi)}} [U^{\psi}_{3.i.3}] C_{\mathcal{A}}\neg\vartheta.$$

So like with $\alpha_{3.i.2}$ we create ϑ worlds in case there are U-reachable $\neg\varphi$ worlds.

Lemma 14. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{3.i}$ we have*

$$\mathcal{M}, w \models \chi \Rightarrow \mathcal{M}, w \models \alpha_{3.i.3}$$

Proof. If $\mathcal{M}, w \models \chi$ then only $(u_1 \wedge \beta_{3.i}, a, u_2 \wedge \beta_{3.i})$ arrows are retained. Furthermore, the $(\Diamond_U(\beta_{3.i} \wedge \neg\Diamond_U \beta_{3.i}) \wedge \Diamond_U(\beta_{3.i} \wedge \Diamond_U \beta_{3.i}), \mathcal{A}, \beta_{3.i} \wedge \neg\Diamond_U \beta_{3.0})$ clause guarantees that no ϑ worlds are created in these $\beta_{3.i}$ worlds. We therefore have $[U^{\psi}_{3.i.3}] C_{\mathcal{A}}\neg\vartheta$, independent of ψ, so $\mathcal{M}, w \models \alpha_{3.i.3}$. $\square$

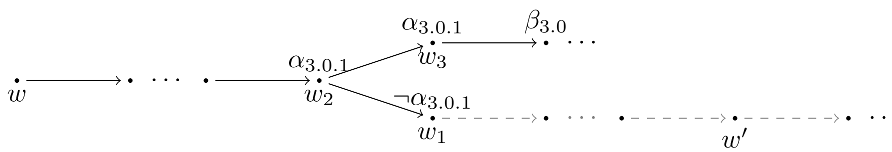

(a) Possibility 1: $\mathcal{M}, w_3 \models \beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U \beta_{3.0}$. We have $\mathcal{M}, w \not\models [U^{+.1}_{3.0.2}]C_{\mathcal{A}}\neg\vartheta$.

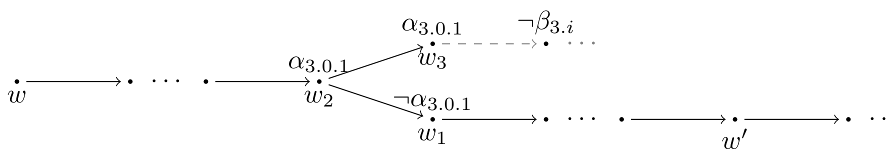

(b) Possibility 2: $\mathcal{M}, w_3 \models \beta_{3.0} \wedge \alpha_{3.0.1} \wedge \neg\Diamond_U \beta_{3.0}$. We have $\mathcal{M}, w \not\models [U^{+.2}_{3.0.2}]C_{\mathcal{A}}\neg\vartheta$.

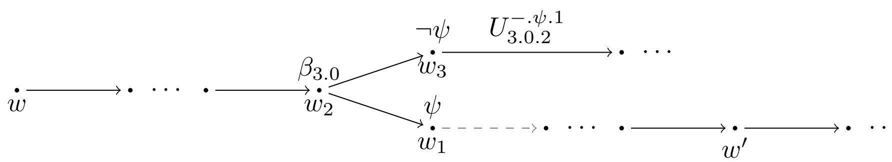

(c) Possibility 3: $\mathcal{M}, w_2 \models \Diamond\psi \wedge \Diamond\neg\psi$ and $\mathcal{M}, w_3 \models \Diamond_{U^{-.\psi.1}_{3.0.2}}\top$. We have $\mathcal{M}, w \not\models [U^{-.\psi.1}_{3.0.2}]C_{\mathcal{A}}\neg\vartheta$.

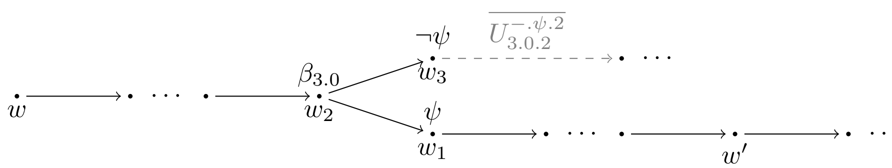

(d) Possibility 4: $\mathcal{M}, w_2 \models \Diamond\psi \wedge \Diamond\neg\psi$ and $\mathcal{M}, w_3 \models \Box_{U^{-.\psi.2}_{3.0.2}}\bot$. We have $\mathcal{M}, w \not\models [U^{-.\psi.2}_{3.0.2}]C_{\mathcal{A}}\neg\vartheta$.

Figure 14: The four possibilities for the conditions of Lemma 13 to hold.

Now for the other side, with another harmless extra condition.

Lemma 15. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{3.i}$ we have that if*

- *there is a U-reachable world w' with $\mathcal{M}, w' \models \neg\beta_{3.i} \wedge \neg\varphi$ and*
- *$\mathcal{M}, w \models \alpha_{3.i.1} \wedge \alpha_{3.i.2}$*

then $\mathcal{M}, w \models \neg\alpha_{3.i.3}$.

Proof. Fix any world w' satisfying the condition of the Lemma. First note that from $\mathcal{M}, w \models \alpha_{3.i.1} \wedge \alpha_{3.i.2}$ it follows that there is no U-reachable world satisfying $\beta_{3.i} \wedge \neg\varphi'$. The w' satisfying the condition of the Lemma must therefore be the successor of another U-reachable $\neg\beta_{3.i}$ world.

By the assumption that w' is U-reachable there is a U-path from w to w'. Let w_1 be the first $\neg\beta_{3.i}$ world on this path. In particular this implies that $w_1 \neq w'$. Let w_2 be the predecessor of w_1 on the path. There are four possibilities for the situation around w_2, the third of which can only occur if $i > 0$ and the fourth of which can only occur if $i = 0$.

The first possibility is that there is a successor w_3 of w_2 such that $\mathcal{M}, w_3 \models \beta_{3.i} \wedge \Diamond_{U^{\psi}_{3.i.3}} \top$ for some ψ. Note that the only clause that depends on ψ has $\neg\beta_{3.i}$ in the start conditions so it follows that w_3 satisfies this condition for all ψ. The second possibility is that w_2 has no successor as in case 1, but it does have a successor w_3 such that $\mathcal{M}, w_3 \models \beta_{3.i} \wedge \Box_{U^{\psi}_{3.i.3}} \bot$. The third possibility is if $i > 0$ and w_2 has no $\beta_{3.i}$ successors, in which case it must have a $\Diamond^{i-1}(\Diamond\psi' \wedge \Diamond\neg\psi')$ successor w_3. The fourth possibility is if $i = 0$ and w_2 has no $\beta_{3.i}$ successors, in which case it must have ψ' and $\neg\psi'$ successors for some ψ'.

The four different cases are shown in Figure 14. There may be more arrows at some points than the ones shown in the figure but such arrows do not matter as long as the arrows that are shown exist. Arrows that are not retained are drawn in gray and dashed. In all four cases the relevant update is $U^{\psi}_{3.i.3}$ for some ψ, but the ψ in question may differ.

In the first case take ψ such that $\mathcal{M}, w_1 \not\models \psi$, see Figure 15a. The arrows from w_2 to w_3 and w_1 are retained because they are $(\beta_{3.i} \wedge \Diamond_U(\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1}), \mathcal{A}, \top)$ arrows (and neither w_1 not w_3 satisfies $\beta_{3.i} \wedge \neg\Diamond_U(\beta_{3.i} \vee \zeta_{3.i+1})$ so the overlined clause does not apply). The arrow from w_3 to at least one of its successors is retained because of the assumption that $\mathcal{M}, w_3 \models \beta_{3.i} \wedge \Diamond_{U^{\psi}_{3.i.3}} \top$. The arrows from w_1 to its successors are not retained because the only arrows from $\neg\beta_{3.i}$ worlds that are retained have a conjunct ψ in the start condition. We therefore have $\mathcal{M}_{[U^{\psi}_{3.i.3}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.3}$.

In the second case take ψ such that $\mathcal{M}, w_1 \models \psi$, see Figure 15b. The arrows from w_2 to w_3 and w_1 are retained because they are $(\beta_{3.i} \wedge \Diamond_U(\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1}), \mathcal{A}, \top)$ arrows (and neither w_1 not w_3 satisfies $\beta_{3.i} \wedge \neg\Diamond_U(\beta_{3.i} \vee \zeta_{3.i+1})$ so the overlined clause does not apply). The arrows from w_3 to its successors to its successors are not retained because of the assumption that $\mathcal{M}, w_3 \models \beta_{3.i} \wedge \Box_{U^{\psi}_{3.i.3}} \bot$. The arrows from w_1 to its successors are however retained because they are $(\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1} \wedge \psi, \mathcal{A}, \top)$ arrows. We therefore have $\mathcal{M}_{[U^{\psi}_{3.i.3}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.3}$.

In the third case we have $\mathcal{M}, w_3 \models \Diamond^{i-1}(\Diamond\psi' \wedge \Diamond\neg\psi')$. This implies that the arrow from w_2 to w_3 is not a U-arrow, as we couldn't other wise have $\mathcal{M}, w_2 \models \beta_{3.i}$. Since the arrow from w_2 to w_1 is a U-arrow this implies that there are formulas in $\Phi^{kn}_{\mathrm{Pvar}(\chi)}$ that distinguish between w_1 and w_3. Let ψ be such a distinguishing formula with the additional property that $\mathcal{M}, w_3 \models \neg\psi$, see Figure 15c. The arrows from w_2 to w_3 and w_1 are retained because they are $(\beta_{3.i} \wedge \Diamond_U(\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1}), \mathcal{A}, \top)$ arrows. Arrows from w_3 to its successors are not retained because the only arrows from $\neg\beta_{3.i}$ worlds that are retained have a conjunct ψ in the start condition. Arrows from w_1 to its successors are however retained because they are $(\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1} \wedge \psi, \mathcal{A}, \top)$ arrows. We therefore have $\mathcal{M}_{[U^{\psi}_{3.i.3}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.3}$.

In the fourth case we have $\mathcal{M}, w_2 \models \Diamond\psi' \wedge \Diamond\neg\psi'$. Choose ψ such that $\psi \equiv \psi'$ or $\psi \equiv \neg\psi'$ and furthermore $\mathcal{M}, w_1 \models \psi$. Let w_3 be a successor of w_2 with $\mathcal{M}, w_2 \models \neg\psi$. Then, like in the third case the arrows from w_2 to w_3 and w_1 are retained because they are $(\beta_{3.i} \wedge \Diamond_U(\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1}), \mathcal{A}, \top)$ arrows. Arrows from w_3 to its successors are not retained because the only arrows from $\neg\beta_{3.i}$ worlds that are retained have a conjunct ψ in the start condition. Arrows from w_1 to its successors are however retained because they are $(\neg\beta_{3.i} \wedge \neg\zeta_{3.i+1} \wedge \psi, \mathcal{A}, \top)$ arrows. We therefore have $\mathcal{M}_{[U^{\psi}_{3.i.3}]}, w \models \neg C_{\mathcal{A}}\neg\vartheta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.3}$.

These four cases are exhaustive so this proves the Lemma. □

B.3.5 Constructing $\alpha_{3.i}$

Let

$$\alpha_{3.i} := \alpha_{3.i.1} \wedge \alpha_{3.i.2} \wedge \alpha_{3.i.3}$$

and

$$\zeta_{3.i} := \zeta_{3.i+1} \wedge (\beta_{3.i} \to \alpha_{3.i}).$$

For any $\mathcal{M}, w$ such that $\mathcal{M}, w \models \beta_{3.i}$ we now have the following results:

- $\mathcal{M}, w \models \chi \Rightarrow \mathcal{M}, w \models \alpha_{3.i.1} \wedge \alpha_{3.i.2} \wedge \alpha_{3.i.3}$
- if there is a U-reachable $\neg\varphi \wedge \beta_{3.i}$ world and $\mathcal{M}, w \models \alpha_{3.i.1}$ then $\mathcal{M}, w \models \neg\alpha_{3.i.2}$
- if there is a U-reachable $\neg\varphi \wedge \neg\beta_{3.i}$ world and $\mathcal{M}, w \models \alpha_{3.i.1} \wedge \alpha_{3.i.2}$ then $\mathcal{M}, w \models \neg\alpha_{3.i.3}$.

Since any U-reachable $\neg\varphi$ world must satisfy either $\neg\varphi \wedge \beta_{3.i}$ or $\neg\varphi \wedge \neg\beta_{3.i}$ this is sufficient to show that $\mathcal{M}, w \models \chi \Leftrightarrow \mathcal{M}, w \models \alpha_{3.i}$.

We can then take

$$\alpha_3 := \bigwedge_{0 \leq i \leq n} (\beta_{3.i} \to \alpha_{3.i}).$$

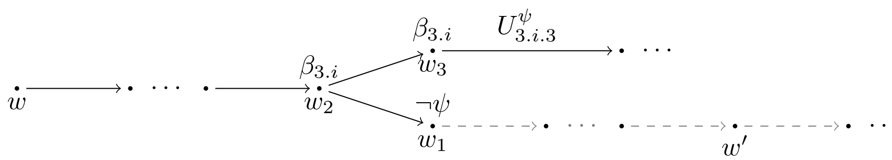

(a) Possibility 1: $\mathcal{M}, w_3 \models \beta_{3.i} \wedge \Diamond_{U^{\psi}_{3.i.3}} \top$.

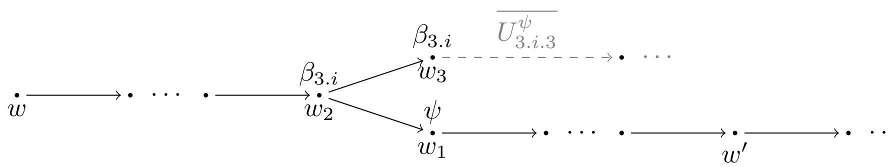

(b) Possibility 2: $\mathcal{M}, w_3 \models \beta_{3.i} \wedge \Box_{U^{\psi}_{3.i.3}} \bot$.

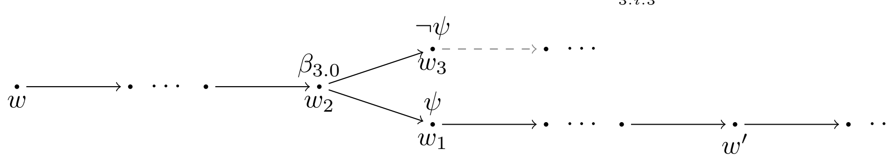

(c) Possibility 3: $\mathcal{M}, w_3 \models \wedge\Diamond^{i-1}(\Diamond\psi' \wedge \Diamond\neg\psi')$.

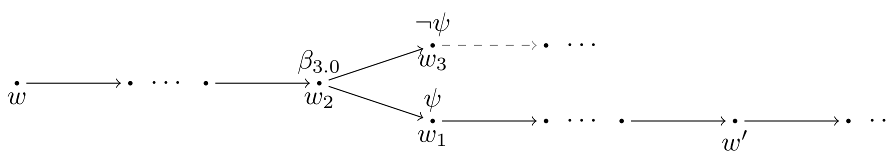

(d) Possibility 4: $\mathcal{M}, w_2 \models \Diamond\psi \wedge \Diamond\neg\psi$.

Figure 15: The four possibilities for the conditions of Lemma 15 to hold. In each case we have $\mathcal{M}, w \models \neg[U^{\psi}_{3.i.3}]C_{\mathcal{A}}\neg\vartheta$ for an appropriate choice of ψ.

B.4 Constructing α_2

Since α_2 has two extra indices for subcases the formula $\zeta_{2.i.j}$ is slightly harder to define than $\zeta_{3.i}$. The base case is $\zeta_{2.|\mathrm{Pvar}(\chi)|+1.0}$ given here, the other cases are defined by induction at the end of this section.

$$\zeta_{2.|\mathrm{Pvar}(\chi)|+1.0} := (\beta_6 \to \alpha_6) \wedge (\beta_5 \to \alpha_5) \wedge (\beta_4 \to \alpha_4) \wedge (\beta_3 \to \alpha_3).$$

B.4.1 Constructing $\alpha_{2.i.0}$

If we are in case $2.i.0$ we know that there are both a U-reachable p_i world and a U-reachable $\neg p_i$ world. The solution $\alpha_{2.i.0}$ works by creating two updates. The $+$ update will guarantee that the only $\Box\bot$ worlds are p_i worlds—except if there is a U-reachable $\neg p_i \wedge \neg\chi$ world. The $-$ update will likewise guarantee that the $\Box\bot$ worlds are $\neg p_i$ worlds unless there is a U-reachable $p_i \wedge \neg\chi$ world.

For $1 \leq i \leq |\mathrm{Pvar}(\chi)|$ let

$$\begin{aligned} U^+_{2.i.0} &:= \{(u_1 \wedge \beta_{2.i.0}, a, u_2 \wedge (\beta_{2.i.0} \vee p_i \vee \neg\zeta_{2.i.1})) \mid (u_1, a, u_2) \in U\} \\ &\quad \cup \{\overline{(\neg p_i \wedge \neg\varphi, \mathcal{A}, \top)}\}, \\ U^-_{2.i.0} &:= \{(u_1 \wedge \beta_{2.i.0}, a, u_2 \wedge (\beta_{2.i.0} \vee \neg p_i \vee \neg\zeta_{2.i.1})) \mid (u_1, a, u_2) \in U\} \\ &\quad \cup \{\overline{(p_i \wedge \neg\varphi, \mathcal{A}, \top)}\} \end{aligned}$$

and

$$\alpha_{2.i.0} := [U^+_{2.i.0}]C_{\mathcal{A}}(\Box\bot \to p_i) \wedge [U^-_{2.i.0}]C_{\mathcal{A}}(\Box\bot \to \neg p_i).$$

Lemma 16. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{2.i.0}$ we have*

$$\mathcal{M}, w \models \chi \Leftrightarrow \mathcal{M}, w \models \alpha_{2.i.0}.$$

Proof. I show the results for $U^+_{2.i.0}$, the $-$ case is the same except for an interchanging of p_i and $\neg p_i$.

First suppose $\mathcal{M}, w \models \chi$. Then the overlined clause in $U^+_{2.i.0}$ cannot apply. Likewise, the $\neg\zeta_{2.i.1}$ term cannot occur, so $U^+_{2.i.0}$ simplifies to

$$U^{+\prime}_{2.i.0} := \{(u_1 \wedge \beta_{2.i.0}, a, u_2 \wedge (\beta_{2.i.0} \vee p_i)) \mid (u_1, a, u_2) \in U\}.$$

Let w' be any U-reachable world that has no successors after the update. Then it is either a $\neg\beta_{2.i.0}$ world or a $\beta_{2.i.0} \wedge \neg\Diamond_U(\beta_{2.i.0} \vee p_i)$ world.

If w' is a $\neg\beta_{2.i.0}$ world then it must be a p_i world in order to satisfy the end condition of an arrow. If it is a $\beta_{2.i.0} \wedge \neg\Diamond_U(\beta_{2.i.0} \vee p_i)$ world then there are both a p_i world and a $\neg p_i$ world U-reachable from w' but this is not the case for any of its successors. But all its successors satisfy $\neg p_i$ so w' must itself be a p_i world. This shows that $\mathcal{M}, w \models [U^+_{2.i.0}]C_{\mathcal{A}}(\Box\bot \to p_i)$.

Now suppose that $\mathcal{M}, w \models \neg\chi$. Then there is a U-reachable $\neg\varphi$ world w'. Assume without loss of generality that w' is the first $\neg\varphi$ world on the U-path from w to w'.

If $\mathcal{M}, w_1 \models \beta_{2.i.0} \wedge \neg p_i$ then the path to w_1 is retained by the update because all arrows in it are $(u_1 \wedge \beta_{2.i.0}, a, u_2 \wedge \beta_{2.1.0})$ arrows. Arrows from w' are not retained however, because they are $\{\overline{(\neg p_i \wedge \neg\varphi, \mathcal{A}, \top)}$ arrows. We therefore have $\mathcal{M}, w \models \neg[U^+_{2.i.0}]C_\mathcal{A}(\Box\bot \to p_i)$.

If $\mathcal{M}, w_1 \models \neg\beta_{2.i.0} \wedge \neg p_i$ let w_1 be the last $\beta_{2.i.0}$ world on the U-path from w to w' and w_2 the successor of w_1 along this path. There is a $\neg p_i$ world U-reachable from w_2, but not both a $\neg p_i$ and a p_i world. This implies that in particular $\mathcal{M}, w_2 \models \neg p_i$. Furthermore, w_2 is in a later case and has a U-reachable $\neg\varphi$ world so $\mathcal{M}, w_2 \models \neg\zeta_{2.i.1}$. The arrow from w_1 to w_2 is therefore a $(u_1 \wedge \beta_{2.i.0}, a, u_2 \wedge \neg\zeta_{2.i.1})$ arrow and therefore retained by the update. No arrow from w_2 is retained because every start condition includes a $\beta_{2.i.1}$ conjunct. We therefore have $\mathcal{M}, w \models \neg[U^+_{2.i.0}]C_\mathcal{A}(\Box\bot \to p_i)$.

Mutatis mutandis this also shows that if $\mathcal{M}, w \models \chi$ then also $\mathcal{M}, w \models [U^-_{2.i.0}]C_\mathcal{A}(\Box\bot \to p_i)$ and that if the first U-reachable $\neg\varphi$ world is a p_i world then $\mathcal{M}, w \models \neg[U^-_{2.i.0}]C_\mathcal{A}(\Box\bot \to p_i)$, which completes the proof. $\Box$

B.4.2 Constructing $\alpha_{2.i.j}$ with $j > 0$

Let

$$\begin{aligned} U^+_{2.i.j} := \{&(u_1 \wedge \beta_{2.i.j}, a, u_2 \wedge (\beta_{2.i.j} \vee \neg\zeta_{2.i.j+1}) \mid (u_1, a, u_2) \in U)\} \cup \\ \{&(\Diamond^j \neg p_i, \mathcal{A}, \Diamond^{j-1} \neg p_i), \overline{(\beta_{2.i.j} \wedge \neg\varphi, \mathcal{A}, \top)}, \overline{(\neg p_i, \mathcal{A}, \top)}\}, \end{aligned}$$

$$\begin{aligned} U^-_{2.i.j} := \{&(u_1 \wedge \beta_{2.i.j}, a, u_2 \wedge (\beta_{2.i.j} \vee \neg\zeta_{2.i.j+1}) \mid (u_1, a, u_2) \in U)\} \cup \\ \{&(\Diamond^j p_i, \mathcal{A}, \Diamond^{j-1} p_i), \overline{(\beta_{2.i.j} \wedge \neg\varphi, \mathcal{A}, \top)}, \overline{(p_i, \mathcal{A}, \top)}\} \end{aligned}$$

and

$$\alpha_{2.i.j} := (p_i \to [U^+_{2.i.j}]C_\mathcal{A}(\Box\bot \to \neg p_i)) \wedge (\neg p_i \to [U^-_{2.i.j}]C_\mathcal{A}(\Box\bot \to p_i)).$$

Lemma 17. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{2.i.j}$ we have*

$$\mathcal{M}, w \models \chi \Leftrightarrow \mathcal{M}, w \models \alpha_{2.i.j}.$$

Proof. I show the results for $U^+_{2.i.0}$, the $-$ case is the same except for an interchanging of p_i and $\neg p_i$.

Unless a $\beta_{2.i.j} \wedge \neg\varphi$ or $\neg\zeta_{2.i.j+1}$ world is encountered the update $U^+_{2.i.0}$ retains all arrows to U-reachable $\beta_{2.i.j}$ worlds by the $(u_1 \wedge \beta_{2.i.j}, a, u_2 \wedge (\beta_{2.i.j} \vee \neg\zeta_{2.i.j+1})$ clauses and all paths to the nearby $\neg p_i$ worlds by the $(\Diamond^j \neg p_i, \mathcal{A}, \Diamond^{j-1} \neg p_i)$ clause. The worlds on the way to the $\neg p_i$ world are p_i worlds with a $\neg p_i$ world reachable in less than j steps so they are not $\beta_{2.i.j}$ worlds. The $\neg p_i$ world itself may be a $\beta_{2.i.j}$ world but its outgoing arrows are not retained because of the $\overline{(\neg p_i, \mathcal{A}, \top)}$ clause.

This implies that the $\neg\zeta_{2.i.j+1}$ possibility in the end condition of $(u_1 \wedge \beta_{2.i.j}, a, u_2 \wedge (\beta_{2.i.j} \vee \neg\zeta_{2.i.j+1})$ and the $\overline{(\beta_{2.i.j} \wedge \neg\varphi, \mathcal{A}, \top)}$ can only apply in U-reachable worlds.

Suppose $\mathcal{M}, w \models \chi$ and $\mathcal{M}, w \models p_i$. Then the $\neg\zeta_{2.i.j+1}$ term and the $\overline{(\beta_{2.i.j} \land \neg\varphi, \mathcal{A}, \top)}$ clause cannot apply in U-reachable worlds so the update simplifies to

$$\begin{aligned} U^{+\prime}_{2.i.j} :=& \{(u_1 \land \beta_{2.i.j}, a, u_2 \land \beta_{2.i.j}) \mid (u_1, a, u_2) \in U)\} \cup \\ & \{(\Diamond^j \neg p_i, \mathcal{A}, \Diamond^{j-1} \neg p_i), \overline{(\neg p_i, \mathcal{A}, \top)}\}. \end{aligned}$$

Any $p_i \land \beta_{2.i.j}$ world has a successor after this update, since it has either an arrow to a $\beta_{2.i.j}$ world that is retained or an arrow to a $\Diamond^{j-1}\neg p_i$ world that is retained. The p_i worlds on the way from a $\beta_{2.i.j}$ world to a $\neg p_i$ world also have a successor after the update because they have an arrow to a $\Diamond^{j-1}\neg p_i$ world that is retained. These are the only $U^{+\prime}_{2.i.j}$-reachable p_i worlds, so $\mathcal{M}, w \models [U^{+\prime}_{2.i.j}]C_{\mathcal{A}}(\Box\bot \to \neg p_i)$ and therefore also $\mathcal{M}, w \models [U^{+}_{2.i.j}]C_{\mathcal{A}}(\Box\bot \to \neg p_i)$.

Suppose on the other hand that $\mathcal{M}, w \models \neg\chi$ and $\mathcal{M}, w \models p_i$. Then there is a U-reachable world w' with $\mathcal{M}, w' \models \neg\varphi$. Suppose without loss of generality that w' is the first $\neg\varphi$ world on the U-path from w to w'.

If $\mathcal{M}, w' \models \beta_{2.i.j}$ it is reachable after the update $U^{+}_{2.i.j}$ but has no successor after that update because of the $\overline{(\beta_{2.i.j} \land \neg\varphi, \mathcal{A}, \top)}$ clause. From $\mathcal{M}, w \models p_i$ and $\mathcal{M}, w \models \beta_{2.i.j}$ with $j > 0$ it also follows that $\mathcal{M}, w' \models p_i$, so $\mathcal{M}, w \models \neg[U^{+}_{2.i.j}]C_{\mathcal{A}}(\Box\bot \to \neg p_i)$.

If $\mathcal{M}, w' \models \neg\beta_{2.i.j}$ let w_1 be the last $\beta_{2.i.j}$ world on the U-path from w to w' and w_2 the successor of w_1 along this path. Then w_1 is reachable after the update $U^{+}_{2.i.j}$. The arrow from w_1 to w_2 is also retained by the update because it is an $(u_1 \land \beta_{2.i.j}, a, u_2 \land \neg\zeta_{2.i.j+1})$ arrow. Arrows from w_2 are not retained however, because $\mathcal{M}, w_2 \models \neg\beta_{2.i.j} \land \neg\Diamond^j\neg p_i$. From $\mathcal{M}, w \models p_i$ and $\mathcal{M}, w \models \beta_{2.i.j}$ with $j > 0$ it also follows that $\mathcal{M}, w' \models p_i$, so $\mathcal{M}, w \models \neg[U^{+}_{2.i.j}]C_{\mathcal{A}}(\Box\bot \to \neg p_i)$.

This shows that $[U^{+}_{2.i.j}]C_{\mathcal{A}}(\Box\bot \to \neg p_i)$ is equivalent to χ under the conditions $\beta_{2.i.j}$ and p_i. Mutatis mutandis it also shows that $[U^{-}_{2.i.j}]C_{\mathcal{A}}(\Box\bot \to p_i)$ is equivalent to χ under the conditions $\beta_{2.i.j}$ and $\neg p_i$. This proves the Lemma. $\Box$

B.4.3 Constructing α_2

We can now give the definition of $\zeta_{2.i.j}$:

$$\zeta_{2.i.3n+1} := \zeta_{2.i+1.0},$$

$$\zeta_{2.i.j} := \zeta_{2.i.j+1} \land (\beta_{2.i.j} \to \alpha_{2.i.j}).$$

We can also define α_2:

$$\alpha_2 := \bigwedge_{1 \le i \le |\mathrm{Pvar}(\chi)|} \bigwedge_{0 \le j \le 3n} (\beta_{2.i.j} \to \alpha_{2.i.j})$$

B.5 Constructing α_1

B.5.1 Constructing $\alpha_{1.i.j.k}$

Let

$$U_{1.i.j.0} := \{(u_1 \wedge \beta_{1.i.j.0}, a, u_2 \wedge (\beta_{1.i.j.0} \vee \neg\zeta_{1.i.j.0+1})) \mid (u_1, a, u_2) \in U\} \cup \{(\top, a_j, \top), \overline{(\beta_{1.i.j.0} \wedge \neg\varphi, \mathcal{A}, \top)}\},$$

$$U_{1.i.j.k} := \{(u_1 \wedge \beta_{1.i.j.k}, a, u_2 \wedge (\beta_{1.i.j.k} \vee \neg\zeta_{1.i.j.k+1})) \mid (u_1, a, u_2) \in U\} \cup \{(\Diamond\!\!\!\cdot^{k} \Diamond_{a_j} \top, \mathcal{A}, \Diamond\!\!\!\cdot^{k-1} \Diamond_{a_j} \top), (\top, a_j, \top), \overline{(\beta_{1.i.j.k} \wedge \neg\varphi, \mathcal{A}, \top)}\}$$

and

$$\alpha_{1.i.j.k} := [U_{1.i.j.k}] C_{\mathcal{A}\setminus\{a_j\}} \neg\Box\bot.$$

Lemma 18. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{1.i.j.k}$ we have*

$$\mathcal{M}, w \models \chi \Leftrightarrow \mathcal{M}, w \models \alpha_{1.i.j.k}.$$

Proof. First suppose that $\mathcal{M}, w \models \chi$. Then the $\neg\zeta_{1.i.j.k+1}$ possibility and the $\overline{(\beta_{1.i.j.k} \wedge \neg\varphi, \mathcal{A}, \top)}$ cannot apply. The remaining possibilities leave exactly those paths intact that lead to U-reachable $\beta_{1.i.j.k}$ worlds or from a U-reachable $\beta_{1.i.j.k}$ world a nearby $\Diamond_{a_j}\top$ world. The only candidates for being $\Box\bot$ worlds are the worlds that are reached by the $(\top, a_j, \top)$ clause and worlds that can only be reached by passing through such a world. This implies that $\mathcal{M}_{[U_{1.i.j.k}]}, w \models C_{\mathcal{A}\setminus\{a_j\}} \neg\Box\bot$, so $\mathcal{M}, w \models \alpha_{1.i.j.k}$.

Now suppose that $\mathcal{M}, w \models \neg\chi$. Then there is a U-reachable world w' with $\mathcal{M}, w' \models \neg\varphi$. Suppose without loss of generality that w' is the first $\neg\varphi$ world on the U-path from w to w'.

If $\mathcal{M}, w' \models \beta_{1.i.j.k}$ it is reachable after the update $U_{1.i.j.k}$ but has no successor after that update because of the $\overline{(\beta_{1.i.j.k} \wedge \neg\varphi, \mathcal{A}, \top)}$ clause so $\mathcal{M}_{[U_{1.i.j.k}]}, w' \models \Box\bot$. We therefore have $\mathcal{M}, w \models \neg\alpha_{1.i.j.k}$.

If $\mathcal{M}, w' \models \neg\beta_{1.i.j.k}$ let w_1 be the last $\beta_{1.i.j.k}$ world on the U-path from w to w' and w_2 the successor of w_1 along this path. Then w_1 is reachable after the update $U_{1.i.j.k}$. The arrow from w_1 to w_2 is also retained by the update because it is an $(u_1 \wedge \beta_{1.i.j.k}, a, u_2 \wedge \neg\zeta_{1.i.j.k+1})$ arrow. Arrows from w_2 are not retained however, because $\mathcal{M}, w_2 \models \neg\beta_{1.i.j.k} \wedge \neg\Diamond\!\!\!\cdot^{j} \Diamond_{a_j}\top$. This implies that w_2 is a reachable $\Box\bot$ world after the update, so $\mathcal{M}, w \models \neg\alpha_{1.i.j.k}$. $\Box$

B.5.2 Constructing $\alpha_{1.i.-1}$

Recall that $B_1, \cdots, B_{2^{|\mathcal{A}|}-|\mathcal{A}|-1}$ are all the subsets of $\mathcal{A}$ with at least two elements, ordered in such a way that if $B_i \subset B_j$ then $i > j$ and $\beta_{1.i.-1}$ is the case where the agents in B_i are exactly the agents for which there is a U-arrow departing from a U-reachable world.

For any $1 \leq i \leq 2^{|\mathcal{A}|} - |\mathcal{A}| - 1$ and any j such that $a_j \in B_i$ let

$$\begin{aligned}U^j_{1.i.-1} := \{&(u_1 \wedge (\beta_{1.i.-1} \vee \neg\zeta_{1.i+1.-1}), a, u_2 \wedge \\ &(\beta_{1.i.-1} \vee \neg\zeta_{1.i+1.-1})) \mid (u_1, a, u_2) \in U\} \cup \\ &\{(u_1 \wedge [U]\neg C_{\mathcal{A}}\Box_{a_j}\bot, a, u_2 \wedge [U]\neg C_{\mathcal{A}}\Box_{a_j}\bot) \mid (u_1, a, u_2) \in U\} \cup \\ &\{(u_1, a_j, u_2) \mid (u_1, a_j, u_2) \in U\} \cup \{\overline{(\neg\varphi, \mathcal{A}, \top)}\}\end{aligned}$$

and

$$\alpha_{1.i.-1} := \bigwedge_{j: a_j \in B_i} [U^j_{1.i.-1}] C_{\mathcal{A}} \Box_{\mathcal{A}\setminus a_j} \Diamond \top$$

Lemma 19. *For any $\mathcal{M}$ and w such that $\mathcal{M}, w \models \beta_{1.i.-1}$ we have*

$$\mathcal{M}, w \models \chi \Leftrightarrow \mathcal{M}, w \models \alpha_{1.i.-1}.$$

Proof. First, suppose $\mathcal{M}, w \models \chi$. The update $U^j_{1.i.-1}$ only retains U-arrows so then the $\zeta_{1.i+1.-1}$ term and the $\overline{(\neg\varphi, \mathcal{A}, \top)}$ clause cannot apply. The clauses $(u_1 \wedge \beta_{1.i.-1}, a, u_2 \wedge \beta_{1.i.-1})$ and $(u_1 \wedge [U]\neg C_{\mathcal{A}}\Box_{a_j}\bot, a, u_2 \wedge [U]\neg C_{\mathcal{A}}\Box_{a_j}\bot)$ retain exactly the paths that go to worlds from which there is a U-reachable world with a departing (u_1, a_j, u_2) arrow. Let w' be any $U^j_{1.i.-1}$ reachable $\Box\bot$ world. Then it cannot have been reached by a $(u_1 \wedge \beta_{1.i.-1}, a, u_2 \wedge \beta_{1.i.-1})$ or $(u_1 \wedge [U]\neg C_{\mathcal{A}}\Box_{a_j}\bot, a, u_2 \wedge [U]\neg C_{\mathcal{A}}\Box_{a_j}\bot)$ arrows, as in those cases there is always either an $(u_1 \wedge [U]\neg C_{\mathcal{A}}\Box_{a_j}\bot, a, u_2 \wedge [U]\neg C_{\mathcal{A}}\Box_{a_j}\bot)$ arrow or an (u_1, a_j, u_2) arrow departing from the target world. The world w' can therefore only be reached by a (u_1, a_j, u_2) arrow. This implies that $\mathcal{M}, w \models [U^j_{1.i.-1}] C_{\mathcal{A}} \Box_{\mathcal{A}\setminus a_j} \Diamond \top$.

Now suppose that $\mathcal{M}, w \models \neg\chi$. Then there is a U-reachable world w' with $\mathcal{M}, w' \models \neg\varphi$. Suppose without loss of generality that w' is the first $\neg\varphi$ world on the U-path from w to w'. Then for any $j \in B_i$ we have $\mathcal{M}, w' \models [U^j_{1.i.-1}]\Box\bot$ because of the $\overline{(\neg\varphi, \mathcal{A}, \top)}$ clause.

If $\mathcal{M}, w' \models \beta_{1.i.-1}$ then w' is reachable in $\mathcal{M}_{U^j_{1.i.-1}}$ for any $j \in B_i$ because of the $(u_1 \wedge (\beta_{1.i.-1} \vee \neg\zeta_{1.i+1.-1}), a, u_2 \wedge (\beta_{1.i.-1} \vee \neg\zeta_{1.i+1.-1}))$ clauses. Take any j such that there at least one of the retained arrows from the predecessor of w' to w' is a non-j arrow. Then $\mathcal{M}, w \models [U^j_{1.i.-1}]\neg C_{\mathcal{A}} \Box_{\mathcal{A}\setminus a_j} \Diamond \top$.

If $\mathcal{M}, w' \models \neg\beta_{1.i.-1}$ let w_1 be the last $\beta_{1.i.-1}$ world on the U-path from w to w' and w_2 the successor of w_1 along this path. The world w' is reachable in $\mathcal{M}_{U^j_{1.i.-1}}$ for any $j \in B_i$ because of the $(u_1 \wedge (\beta_{1.i.-1} \vee \neg\zeta_{1.i+1.-1}), a, u_2 \wedge (\beta_{1.i.-1} \vee \neg\zeta_{1.i+1.-1}))$ clauses; the path up to w_1 consists of $(u_1 \wedge \beta_{1.i.-1}, a, u_2 \wedge \beta_{1.i.-1})$ arrows, the arrow from w_1 to w_2 is an $(u_1 \wedge \beta_{1.i.-1}, a, u_2 \wedge \neg\zeta_{1.i+1.-1})$ arrow and the path from w_2 to w' consists of $(u_1 \wedge \neg\zeta_{1.i+1.-1}, a, u_2 \wedge \neg\zeta_{1.i+1.-1})$ arrows. Take any j such that there at least one of the retained arrows from the predecessor of w' to w' is a non-j arrow. Then $\mathcal{M}, w \models [U^j_{1.i.-1}]\neg C_{\mathcal{A}} \Box_{\mathcal{A}\setminus a_j} \Diamond \top$. $\square$

B.5.3 Constructing α_1

We can now define α_1 and ζ for the appropriate indices.

$$\alpha_1 := \bigwedge_{1\leq i\leq 2^{|\mathcal{A}|}-|\mathcal{A}|-1} (\beta_{1.i.-1} \to \alpha_{1.i.-1}) \wedge \bigwedge_{1\leq i\leq|\mathcal{A}|} \bigwedge_{1\leq j\leq|\mathcal{A}|, j\neq i} \bigwedge_{1\leq k\leq 3n} (\beta_{1.i.j.k} \to \alpha_{1.i.j.k}).$$

The definition of ζ is a bit more complicated in this case than it is in the other cases due to the more complex indexing. First let us define $\zeta_{1.i.j.k}$.

$$\zeta_{1.|\mathcal{A}|.|\mathcal{A}|-1.3n+1} := (\beta_6 \to \alpha_6) \wedge (\beta_5 \to \alpha_5) \wedge (\beta_4 \to \alpha_4) \wedge (\beta_3 \to \alpha_3) \wedge (\beta_2 \to \alpha_2)$$

$$\zeta_{1.i.j.3n+1} := \begin{cases} \zeta_{1.i.j+1.0} & \text{if } j+1 \neq i \\ \zeta_{1.i.j+2.0} & \text{if } j+1 = i \text{ and } i \neq |\mathcal{A}| \end{cases}$$

$$\zeta_{1.i.|\mathcal{A}|+1.0} := \zeta_{1.i+1.1.0}$$

$$\zeta_{1.i.j.k} := \zeta_{1.i.j.k+1} \wedge (\beta_{1.i.j.k} \to \alpha_{1.i.j.k}).$$

Now we can define $\zeta_{1.i.-1}$ by

$$\zeta_{1.|\mathcal{A}|+1.-1} := \zeta_{1.1.2.0}$$

and

$$\zeta_{1.i.-1} := \zeta_{1.i+1.-1} \wedge (\beta_{1.i.-1} \to \alpha_{1.i.-1}).$$

C Proofs of auxiliary lemmas

Lemma 20. *Let $\mathcal{M}$ be a model, w_1 and w_2 worlds in the model and ψ an AUC formula such that w_2 is a successor of w_1, $\mathcal{M}, w_1 \models \psi$ and $\mathcal{M}, w_2 \models \neg\psi$. Then one of the following holds:*

1. *there are two agents b_1, b_2 and two worlds w_3, w_4 such that $b_1 \neq b_2$, $\mathcal{M}, w_3 \models \Diamond_{b_1}\top$, $\mathcal{M}, w_4 \models \Diamond_{b_2}\top$ and both w_3 and w_4 are reachable from w_1 in at most $d(\psi)$ steps.*

2. *there are a propositional variable $p \in \mathrm{Pvar}(\psi)$ and two worlds w_3, w_4 such that $\mathcal{M}, w_3 \models p$, $\mathcal{M}, w_4 \models \neg p$ and both w_3 and w_4 are reachable from w_1 in at most $d(\psi)+1$ steps.*

3. *there are a formula $\psi' \in \Phi^{d(\psi)}_{\mathrm{Pvar}(\psi)}$ and a world w_3 such that $\mathcal{M}, w_3 \models \Diamond\psi' \wedge \Diamond\neg\psi'$ and there is a $k \in \mathbb{N}$ such that w_3 is reachable from w_1 in at most k steps and $k + d(\psi') \leq d(\psi)$.*

4. *there is a $k \leq d(\psi)$ such that $\mathcal{M}, w_1 \models \Diamond^k\Box\bot \wedge \Box^{k+1}\bot$.*

Proof. The proof is by showing that if the conditions hold and we are not in one of the first two cases then we are in one of the last two cases, and by induction on the depth of ψ.

The base case is trivial; if ψ is of depth 0 then it is a boolean combination of propositional variables so there is at least one propositional variable that takes different values in the two worlds so we are in case 2.

Suppose therefore that $\mathcal{M}$, w_1, w_2 and ψ are as in the Lemma, that $d(\psi) > 0$, that the first two possibilities do not hold and that the Lemma holds for all ψ' with $d(\psi') < d(\psi)$. If a boolean combination of formulas distinguishes between two worlds then so does at least one of the combined formulas so we can assume without loss of generality that ψ is either of pure depth or the negation of a formula of pure depth.[11]

A formula of pure depth > 0 or the negation thereof must have $\Diamond_a, \Box_a, [U], \neg[U], C_B$ or $\neg C_B$ as main connective. Now consider the following validities:

$$\begin{array}{rcl} [U]p & \leftrightarrow & p \\ [U]\neg\xi & \leftrightarrow & \neg[U]\xi \\ [U](\xi_1 \vee \xi_2) & \leftrightarrow & [U]\xi_1 \vee [U]\xi_2 \\ [U]\Box_a\xi_1 & \leftrightarrow & \bigwedge_{(\xi_2,a,\xi_3)\in U}(\xi_2 \rightarrow \Box_a(\xi_3 \rightarrow [U]\xi_1)) \\ \xi & \leftrightarrow & [(\top, \mathcal{A}, \top)]\xi. \end{array}$$

They allow us to find a formula equivalent to ψ that is of the form p, of the form $\Diamond_a\psi'$, of the form $\Box_a\psi'$, of the form $[U']C_B\psi'$ or of the form $\neg[U']C_B\psi'$ for some $p \in \mathcal{P}$, $a \in \mathcal{A}, B \subseteq \mathcal{A}$, $\psi' \in \Phi^{d(\psi)-1}_{\mathrm{Pvar}(\psi)}$ and update U' with $d(U') < d(\psi)$. We already discussed formulas of the form p, so we can restrict to formulas of the forms $\Diamond_a\psi'$, $\Box_a\psi'$, $[U']C_B\psi'$ or $\neg[U']C_B\psi'$ without loss of generality.

- Suppose ψ is of the form $\Diamond_a\psi'$. Then there are three possibilities.
  - Suppose $\mathcal{M}, w_2 \not\models \psi'$. Then $\mathcal{M}, w_1 \models \Diamond\psi' \wedge \Diamond\neg\psi'$ so we are in case 3.
  - Suppose $\mathcal{M}, w_2 \models \psi' \wedge \Box_a\bot$. Then either $\mathcal{M}, w_1 \models \Diamond\Box_a\bot \wedge \Diamond\neg\Box_a\bot$ so we are in case 3 or $\mathcal{M}, w_1 \models \Diamond\Box_a\bot \wedge \Box\Box_a\bot$. In the latter case it follows from the fact that we are not in case 1 and that there is an a-arrow departing from w_1 that $\mathcal{M}, w_1 \models \Diamond\Box\bot \wedge \Box\Box\bot$ so we are in case 4.
  - Suppose $\mathcal{M}, w_2 \models \psi' \wedge \Diamond_a\top$. Then w_2 has a successor w_3 with $\mathcal{M}, w_3 \models \neg\psi'$. We can then apply the Lemma to ψ', w_2 and w_3 by the induction hypothesis. If one of the first three cases holds for ψ', w_2 and w_3 it immediately follows that the same case holds for ψ, w_1 and w_2 as these cases allow the relevant worlds to be a certain distance away. Suppose then that the fourth case holds for ψ', w_2 and w_3 so $\mathcal{M}, w_2 \models \Diamond^{k'}\Box\bot \wedge \Box^{k'+1}\bot$ for some $k' \leq d(\psi')$. Then either all successors of w_1 satisfy the same formula in which case $\mathcal{M}, w_1 \models \Diamond^k\Box\bot \wedge \Box^{k+1}\bot$ for $k = k' + 1 \leq d(\psi)$ so we are in case 4 or at least one successor of w_1 does not satisfy the formula in which case $\mathcal{M}, w_1 \models \Diamond(\Diamond^{k'}\Box\bot \wedge \Box^{k'+1}\bot) \wedge \Diamond\neg(\Diamond^{k'}\Box\bot \wedge \Box^{k'+1}\bot)$ so we are in case 3.

[11] We could of course require ψ to be of pure depth and still have it distinguish w_1 and w_2. Allowing ψ to be a negation of a formula of pure depth allows us to guarantee that ψ holds in w_1 and not in w_2.

- Suppose ψ is of the form $\Box_a\psi'$. We are not in case 1 so there is only one agent nearby. We therefore also have $\mathcal{M}, w_1 \models \Box\psi'$ and $\mathcal{M}, w_2 \models \neg\Box\psi'$. Since w_2 is a successor of w_1 we have $\mathcal{M}, w_2 \models \psi'$. But w_2 has a successor w_3 with $\mathcal{M}, w_3 \models \neg\psi'$. We can then apply the Lemma to ψ', w_2 and w_3 by the induction hypothesis. By the same reasoning as in the last subcase of the previous possibility it then follows that the Lemma holds for ψ, w_1 and w_2.
- Suppose ψ is of the form $[U']C_B\psi'$. Then there are three possibilities.
  - Suppose there is no B-arrow from w_1 to w_2. Then the arrow from w_1 to w_2 must be of an agents $a \notin B$ and from the fact that we are not in case 1 it follows that there are only a arrows from w_2. But then we must have $\mathcal{M}, w_2 \models \neg\psi'$ and $\mathcal{M}, w \models \psi'$ so we can apply the Lemma to ψ', w_1 and w_2, from which it immediately follows that the Lemma holds for ψ, w_1 and w_2.
  - Suppose there is a B-arrow from w_1 to w_2 and $\mathcal{M}, w_2 \models \neg\psi'$. Then we can apply the Lemma to ψ', w_1 and w_2, from which it immediately follows that the Lemma holds for ψ, w_1 and w_2.
  - Suppose there is a B-arrow from w_1 to w_2 and $\mathcal{M}, w_2 \models \psi'$. Then the arrow from w_1 to w_2 must not be a U-arrow and there must be a U-arrow from w_2 to a successor w_3 of w_2. From the fact that we are not in case 1 it follows that the arrow from w_1 to w_2 and the arrow from w_2 to w_3 must belong to the same agent. Let (u_1, a, u_2) be the U' clause for which there is an arrow from w_2 to w_3. Then we must have either $\mathcal{M}, w_1 \models \neg u_1$ and $\mathcal{M}, w_2 \models u_1$ or $\mathcal{M}, w_2 \models \neg u_2$ and $\mathcal{M}, w_3 \models u_2$. In the first case we can apply the Lemma to $\neg u_1$, w_1 and w_2 and it follows immediately that the Lemma holds for ψ', w_1 and w_2. In the second case we can apply the lemma to u_2, w_1 and w_2 from which it follows that the Lemma holds for ψ', w_1 and w_2 by the same reasoning as in the last subcase of the first possibility.
- Suppose ψ is of the form $\neg[U']C_B\psi'$. Then there are two possibilities.
  - Suppose $\mathcal{M}, w_1 \models \neg\psi'$. From $\mathcal{M}, w_2 \models \neg\neg[U']C_B\psi'$ it follows that $\mathcal{M}, w_2 \models \psi'$. We can then apply the Lemma for ψ', w_1 and w_2 and it follows immediately that the Lemma holds for ψ, w_1 and w_2.
  - Suppose $\mathcal{M}, w_2 \models \psi'$. Then there must be a successor w_3 of w_1 with $\mathcal{M}, w_3 \models \neg[U']C_B\psi'$, so we have $\mathcal{M}, w_1 \models \Diamond\psi \wedge \Diamond\neg\psi$. We therefore are in case 3.

This completes the induction step and thereby the proof. $\Box$

Lemma 21. *Let $\mathcal{M}$ be a model, w and w_1 worlds in $\mathcal{M}$, ψ a formula in $\Phi^k_{\mathrm{Pvar}(\chi)}$ with $k \leq 3n$ and $\pi = ((w, b, w'), \cdots, (w'', b', w_1))$ a path from w to w' such that*

- *all arrows in π except possibly some or all of the last $3n - k$ ones are U-arrows,*

- *$\mathcal{M}, w \models \neg\beta_1 \wedge \neg\beta_2$,*
- *$\mathcal{M}, w_1 \models \Diamond\psi \wedge \Diamond\neg\psi$,*
- *there is no $\psi' \in \Phi^k_{\mathrm{Pvar}(\chi)}$ with $d(\psi') < d(\psi)$ and $\mathcal{M}, w_1 \models \Diamond\psi' \wedge \Diamond\psi'$ and*
- *there are no $\psi' \in \Phi^k_{\mathrm{Pvar}(\chi)}$ and w_2 on π with $d(\psi') \leq d(\psi)$, $w_2 \neq w_1$ and $\mathcal{M}, w_2 \models \Diamond\psi' \wedge \Diamond\psi'$.*

Then there is a formula $\xi \in \Phi^k_{\mathrm{Pvar}(\chi)}$ such that $\mathcal{M}, w_1 \models \Diamond\xi \wedge \Diamond\neg\xi$, there is no w_2 on π with $\mathcal{M}, w_2 \models \xi$ and for any successor w_3 of w_1 with $\mathcal{M}, w_3 \models \neg\xi$ we have $\mathcal{M}, w_3 \models \Diamond\top$.

Proof. First, suppose that there is a successor w_3 of w_1 with $\mathcal{M}, w_3 \models \Box\bot$. From $\mathcal{M}, w \models \neg\beta_2$ and the fact that all but the last n arrows in π are U arrows it follows that for each propositional variable p all successors of w_1 have the same value for p. Since ψ distinguishes two successors of w_1 this implies that w_1 must also have a successor satisfying $\neg\Box\bot$. We can then take $\xi = \Box\bot$.

Suppose then that every successor of w_1 satisfies $\Diamond\top$. If a boolean combination of formulas distinguishes between two worlds then at least one of the combined formulas distinguishes them as well, so we can assume without loss of generality that ψ is either of pure depth or the negation of a formula of pure depth. Since this still allows for the negating of a formula we can furthermore assume that $\mathcal{M}, w \models \neg\psi$.

If there is no ψ world on π we can take $\xi = \psi$. Assume therefore that there is a ψ world on π and let w_2 be the first ψ world on the path. We have taken ψ such that $\mathcal{M}, w \models \neg\psi$, so $w_2 \neq w$, so there is a predecessor w_3 of w_2 on π with $\mathcal{M}, w_3 \models \neg\psi$.

We can then apply Lemma 20 to $\neg\psi$, w_3 and w_2. The first two possibilities of Lemma 20 cannot be the case, as this would require either two agents to have arrows within $3n$ steps of a U-reachable world or a propositional variable p such that both p and $\neg p$ hold in some world within $3n$ steps of a U-reachable world.

The fourth possibility cannot occur either, as no $\Diamond\psi \wedge \Diamond\neg\psi$ world can occur after a $\Diamond^j\Box\bot \wedge \Box^{j+1}\bot$ world unless there are either multiple agents or a propositional difference nearby, which there aren't.

We must therefore be in the third possibility: there are a formula $\psi' \in \Phi^{d(\psi)}_{\mathrm{Pvar}(\psi)}$ and a world w_4 such that $\mathcal{M}, w_4 \models \Diamond\psi' \wedge \Diamond\neg\psi'$, w_4 is reachable from w_3 in at most l steps and $l + d(\psi') \leq d(\psi)$.

There may be multiple worlds on π that satisfy $\Diamond^{l'}(\Diamond\psi' \wedge \Diamond\neg\psi')$ for some l' with $l' + d(\psi') \leq d(\psi)$. Let w_5 be such a world on π that minimizes $l' - m$ where m is the distance between the world and w_1. Note that since there is no $\Diamond\psi' \wedge \Diamond\neg\psi'$ world on π before w_1 we must have $l \geq 1$, so $d(\psi') < d(\psi)$.

Every successor of w_5 must satisfy $\Diamond^{l'-1}(\Diamond\psi' \wedge \Diamond\neg\psi')$ since otherwise we would have $\mathcal{M}, w_5 \models \Diamond(\Diamond^{l'-1}(\Diamond\psi' \wedge \Diamond\neg\psi')) \wedge \Diamond\neg(\Diamond^{l'-1}(\Diamond\psi' \wedge \Diamond\neg\psi'))$. In particular the successor w_6 of w_5 along π satisfies $\Diamond^{l'-1}(\Diamond\psi' \wedge \Diamond\neg\psi')$. But for the same reason every successor of w_6 satisfies $\Diamond^{l'-2}(\Diamond\psi' \wedge \Diamond\neg\psi')$.

This can be repeated until we either reach the l'-th successor of w_5 or until we reach w_1, whichever comes first. If the distance m between w_5 and w_1 is at least l' we will reach the l'-th successor w_7 of w_5 on π which satisfies $\Diamond\psi' \wedge \Diamond\neg\psi'$. But $d(\psi') < d(\psi)$ so there can be no such world on π. This is a contradiction, so m must be less than l' and we have $\mathcal{M}, w_1 \models \Diamond^{l'-m}(\Diamond\psi' \wedge \Diamond\neg\psi')$. Since the depth of ψ is minimal for distinguishing the successors of w_1 we also have that all successors of w_1 satisfy $\Diamond^{l'-m}(\Diamond\psi' \wedge \Diamond\neg\psi')$.

Let $\xi = \psi \wedge \Diamond^{l'-m-1}(\Diamond\psi' \wedge \Diamond\neg\psi')$. Then $d(\xi) = d(\psi)$ and $\mathrm{Pvar}(\xi) = \mathrm{Pvar}(\xi)$ so $\xi \in \Phi^k_{\mathrm{Pvar}(\chi)}$. Furthermore, we have $\mathcal{M}, w_1 \models \Diamond\xi \wedge \Diamond\neg\xi$. And, because w_5 was chosen to minimize $l' - m$ we have $\mathcal{M}, w' \models \neg\Diamond^{l'-m-1}(\Diamond\psi' \wedge \Diamond\neg\psi')$ and therefore $\mathcal{M}, w' \models \neg\xi$ for all w' on π. $\square$