Expressivity of Logics of Knowledge and Action

Kuijer, Bouke

Appendix A

Appendix: Detailed proof that U* $\preceq$UC

A.1 Constructing $\beta_1, \cdots, \beta_6$

In order to prove Lemma 3.4 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.1.1 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

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

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 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.

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.

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

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 $1.i.j.k$.

Constructing β_1

We can then simply take

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

A.1.2 Constructing β_2

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

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

Subcases of case 2

Let $\text{Pvar}(\chi) = \{p_1, \cdots, p_{|\text{Pvar}(\chi)|}\}$. There are also several subcases of case 2, cases $2.i.j$ for $1 \leq i \leq |\text{Pvar}(\chi)|$ 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.

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 $1 \leq i \leq |\text{Pvar}(\chi)|$ 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

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

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.

Constructing β_2

We can then simply take

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

A.1.3 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 $\mathcal{M}, w_1 \models \lozenge\psi \wedge \lozenge\neg\psi$. We have

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

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 $\Phi^{2n}_{\text{Pvar}(\chi)}$ is at distance i from a U-reachable world.

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 $\lozenge\psi \wedge \lozenge\neg\psi$ 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 $\lozenge\psi \wedge \lozenge\neg\psi$ in $\mathcal{M}$.

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

$$\theta := \lozenge\Box\bot \wedge \lozenge\lozenge\top.$$

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 $\mathcal{M}, w_1 \models \lozenge\psi \wedge \lozenge\neg\psi$ then $\mathcal{M}_{[U^\psi_{3.i}]}, w_1 \models \theta$. In order guarantee that $\neg\psi$ worlds have successors we need another case distinction. Let

$$\tau^\psi_i := \lozenge_U\top \vee \Diamond^i(\lozenge\psi \wedge \lozenge\neg\psi).$$

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.

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 $\mathcal{M}_{[U^{\psi.+}_{3.i}]}, w_1 \models \theta$ 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

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

and for $0 < i \leq n$ let

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

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.

Constructing $\beta_{3.i}$

For $1 \leq i \leq n$ let

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

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

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 A.1. *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

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

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

Suppose we are in the situation where this happens. Then there is no U-reachable $\Diamond\!\!\!\cdot^i(\Diamond\psi \wedge \Diamond\neg\psi)$ 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 $\mathcal{M}_{[U^{\psi,-}_{3.i}]}, w_1 \models \theta$. 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 $\Diamond\!\!\!\cdot^i(\Diamond\psi \wedge \Diamond\neg\psi)$ 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 length 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 $\mathcal{M}, w \models \neg\{U\}^*\neg\Diamond\!\!\!\cdot^i(\Diamond\psi \wedge \Diamond\neg\psi)$ 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 $\mathcal{M}, w \models \neg\{U\}^*\neg\Diamond\!\!\!\cdot^i(\Diamond\psi \wedge \Diamond\neg\psi)$ without $\mathcal{M}, w \models \neg[U^{\psi,+}_{3.0}]C_{\mathcal{A}}\neg\theta \vee \neg[U^{\psi,-}_{3.0}]C_{\mathcal{A}}\neg\theta$ 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 $\mathcal{M}, w \models \neg[U^{\psi,+}_{3.0}]C_{\mathcal{A}}\neg\theta \vee \neg[U^{\psi,-}_{3.0}]C_{\mathcal{A}}\neg\theta$ we do have

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

The proof of the existence of such a ψ' is rather long and technical and is therefore included in Section A.3 of the Appendix as Lemma A.15.

Constructing β_3

We can then simply take

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

A.1.4 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 3.6.4 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 $p \in \text{Pvar}(\psi)$ 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 $\psi' \in \Phi^{d(\psi)}_{\text{Pvar}(\psi)}$ 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 $k + d(\psi') \le d(\psi)$.

4. there is a $k \le d(\psi)$ 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

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

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 $1 \le i \le n$ 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.1.5 Constructing β_5 and β_6

Finding a CU 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 a CU formula so we can simply take

$$\beta_6 := \delta_6.$$

A.2 Constructing $\alpha_6, \cdots, \alpha_1$

In order to prove Lemma 3.5 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 3.6.3: 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.

A.2.1 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

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

A.2.2 Constructing α_4

In case 4 the only possible cause of two worlds being distinguishable (by a formula of length 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 length 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

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

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

A.2.3 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' in stead 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.

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

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

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 A.2. *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 A.1.

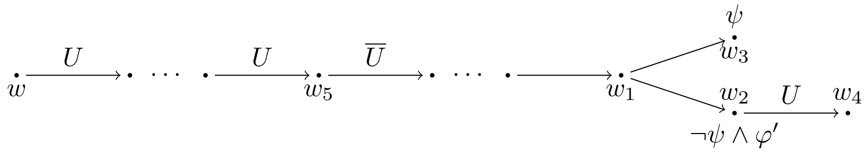

(a) Possibility 1, solved by $U^{\psi.+}_{3.i.1.1}$

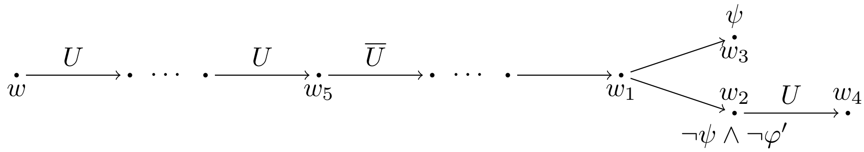

(b) Possibility 2, solved by $U^{\psi.+}_{3.i.1.2}$

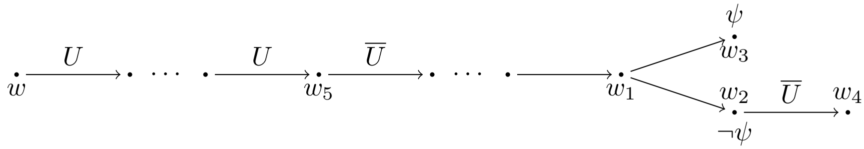

(c) Possibility 3, solved by $U^{\psi.-}_{3.i.1}$

Figure A.1: Three possibilities for a $\Diamond\psi \wedge \Diamond\neg\psi$ world.

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 A.1a, 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 A.1b, 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 A.1c, 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. □

Lemma A.3. *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.[1] 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 length 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 $i = 0$. □

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

For $1 \leq i \leq n$ let

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

and

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

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 A.4. *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

[1] 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.

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 A.2 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

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

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

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. □

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 A.5. *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 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

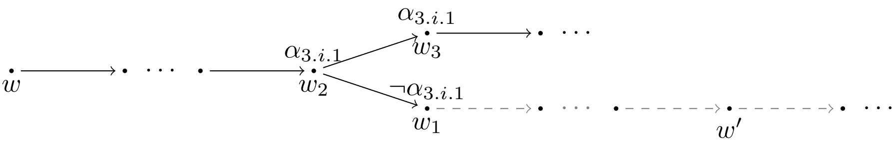

(a) Possibility 1: w_1 has a successor w_3 with $\mathcal{M}, w_3 \models \alpha_{3.i.1}$. We have $\mathcal{M}, w \not\models [U^+_{3.i.2}]C_{\mathcal{A}}\neg\theta$.

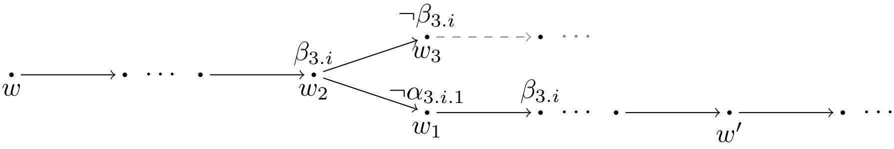

(b) Possibility 2: $\mathcal{M}, w_1 \models \Diamond^i(\Diamond\psi \wedge \Diamond\neg\psi)$ for some $\psi \in \Phi^{2n}_{\text{Pvar}(\chi)}$. We have $\mathcal{M}, w \not\models [U^-_{3.i.2}]C_{\mathcal{A}}\neg\theta$.

Figure A.2: The two possibilities for the conditions of Lemma A.5 to hold.

shown in Figure A.2. 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 A.2a. 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 $\mathcal{M}_{[U^+_{3.i.2}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.2}$.

In the second case, shown in Figure A.2b, 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 $(\beta_{3.i} \wedge \Box_U\neg(\beta_{3.i} \wedge \alpha_{3.i.1}) \wedge \Diamond^i \bigvee_{\psi\in\Phi^{2n}_{\text{Pvar}(\chi)}}(\Diamond\psi \wedge \Diamond\neg\psi), \mathcal{A}, \Diamond^{i-1}\bigvee_{\psi\in\Phi^{2n}_{\text{Pvar}(\chi)}}(\Diamond\psi \wedge \Diamond\neg\psi)$ arrow. We therefore have $\mathcal{M}_{[U^+_{3.i.2}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.2}$.

□

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

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

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

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

and

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

Lemma A.6. *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 $\mathcal{M}_{[U'_{3.0.2}]}, w_1 \models \theta$. 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 $\mathcal{M}, w_3 \models \beta_{3.0} \wedge \neg\Diamond_U \beta_{3.0}$, which implies $\mathcal{M}, w_1 \models \Diamond_U(\beta_{3.0} \wedge \neg\Diamond_U \beta_{3.0})$. But then 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 A.7. *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 $\beta_{3.0} \land \alpha_{3.0.1} \land \Diamond_U \beta_{3.0}$. The second possibility is that w_2 has no successor of the kind in case 1, but does have a successor w_3 satisfying $\beta_{3.0} \land \alpha_{3.0.1} \land \neg\Diamond_U \beta_{3.0}$.

In the third and fourth possibilities w_2 has no successor satisfying $\beta_{3.0} \land \alpha_{3.0.1}$. From $\mathcal{M}, w_2 \models \alpha_{3.0.1}$ it follows that there is some $\Diamond\psi \land \Diamond\neg\psi$ 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 $\Diamond\psi \land \Diamond\neg\psi$ 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 $U'_{3.0.2} := U^{-.\psi.1}_{3.0.2} \setminus \{\overline{(\beta_{3.0} \land \neg\alpha_{3.0.1} \land \psi, \mathcal{A}, \top)}\} = U^{-.\psi.2}_{3.0.2} \setminus \{(\beta_{3.0} \land \neg\alpha_{3.0.1} \land \psi, \mathcal{A}, \top)\}$. 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 $\overline{(\beta_{3.0} \land \neg\alpha_{3.0.1} \land \psi, \mathcal{A}, \top)}$ and $(\beta_{3.0} \land \neg\alpha_{3.0.1} \land \psi, \mathcal{A}, \top)$ 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 A.3. 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 A.3a. 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 $(u_1 \land \beta_{3.0}, a, u_2 \land \beta_{3.0})$ arrows and, because $\mathcal{M}, w_2 \models \alpha_{3.0.1}$ and $\mathcal{M}, w_3 \models \alpha_{3.0.1}$, not $\overline{(\beta_{3.0} \land \neg\alpha_{3.0.1}, \mathcal{A}, \top)}$ arrows. The arrows from w_1 to its successors are not retained because they are $\overline{(\beta_{3.0} \land \neg\alpha_{3.0.1}, \mathcal{A}, \top)}$ arrows. We therefore have $\mathcal{M}_{[U^{+.1}_{3.0.2}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ 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 A.3b. The arrows from w_2 to w_3 and w_1 are retained because they are $(u_1 \land \beta_{3.0}, a, u_2 \land \beta_{3.0})$ 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 $(\beta_{3.0} \land \neg\alpha_{3.0.1}, \mathcal{A}, \top)$ arrows. Arrows from w_3 to its successors are not retained; they are not $(u_1 \land \beta_{3.0}, a, u_2 \land \beta_{3.0})$ arrows because the successors of w_3 are not $\beta_{3.0}$ worlds and they are not $(\beta_{3.0} \land \neg\alpha_{3.0.1}, \mathcal{A}, \top)$ arrows because $\mathcal{M}, w_3 \models \alpha_{3.0.1}$. We therefore have $\mathcal{M}_{[U^{+.2}_{3.0.2}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ 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 A.3c. The arrow from w_2 to w_1 is retained because it is an $(u_1 \land \beta_{3.0}, a, u_2 \land \beta_{3.0})$ arrow. The arrow from w_2 to w_3 is retained because it is an $(\beta_{3.0} \land \alpha_{3.0.1} \land \Diamond_U \neg\alpha_{3.0.1}, \mathcal{A}, \neg\psi)$ 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 $\overline{(\beta_{3.0} \land \neg\alpha_{3.0.1} \land \psi, \mathcal{A}, \top)}$ arrows. We therefore have $\mathcal{M}_{[U^{-.\psi.1}_{3.0.2}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ 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 A.3d. The arrow from w_2 to w_1 is retained because it is an $(u_1 \land \beta_{3.0}, a, u_2 \land \beta_{3.0})$ arrow. The arrow from w_1 to w_3 is retained because it is an $(\beta_{3.0} \land \alpha_{3.0.1} \land \Diamond_U \neg\alpha_{3.0.1}, \mathcal{A}, \neg\psi)$ 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 $(\beta_{3.0} \land \neg\alpha_{3.0.1} \land \psi, \mathcal{A}, \top)$ arrows. We therefore have $\mathcal{M}_{[U^{-.\psi.2}_{3.0.2}]}, w \models \neg C_{\mathcal{A}}\neg\theta$

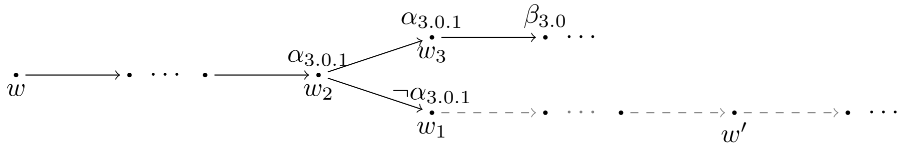

(a) Possibility 1: $\mathcal{M}, w_3 \models \beta_{3.0} \wedge \alpha_{3.0.1} \wedge \Diamond_U \beta_{3.0}$. We have $\mathcal{M}, w \not\models [U^{+.1}_{3.0.2}]C_{\mathcal{A}}\neg\theta$.

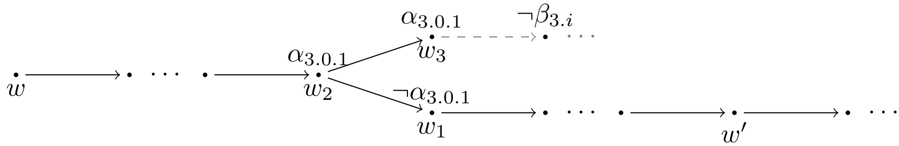

(b) Possibility 2: $\mathcal{M}, w_3 \models \beta_{3.0} \wedge \alpha_{3.0.1} \wedge \neg\Diamond_U \beta_{3.0}$. We have $\mathcal{M}, w \not\models [U^{+.2}_{3.0.2}]C_{\mathcal{A}}\neg\theta$.

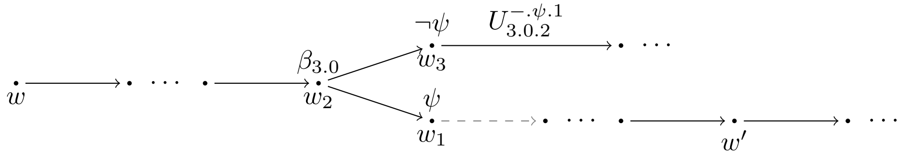

(c) Possibility 3: $\mathcal{M}, w_2 \models \Diamond\psi \wedge \Diamond\neg\psi$ and $\mathcal{M}, w_3 \models \Diamond_{U^{-.\psi.1}_{3.0.2}}\top$. We have $\mathcal{M}, w \not\models [U^{-.\psi.1}_{3.0.2}]C_{\mathcal{A}}\neg\theta$.

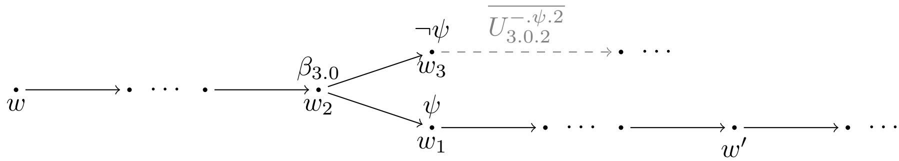

(d) Possibility 4: $\mathcal{M}, w_2 \models \Diamond\psi \wedge \Diamond\neg\psi$ and $\mathcal{M}, w_3 \models \Box_{U^{-.\psi.2}_{3.0.2}}\bot$. We have $\mathcal{M}, w \not\models [U^{-.\psi.2}_{3.0.2}]C_{\mathcal{A}}\neg\theta$.

Figure A.3: The four possibilities for the conditions of Lemma A.7 to hold.

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

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

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

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

and

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

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

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

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

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

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

Lemma A.9. *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 A.3. 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 A.4a. 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 $\mathcal{M}_{[U^{\psi}_{3.i.3}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ 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 A.4b. 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 $\mathcal{M}_{[U^{\psi}_{3.i.3}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ 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 A.4c. 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 $\mathcal{M}_{[U^{\psi}_{3.i.3}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ 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 $\mathcal{M}_{[U^{\psi}_{3.i.3}]}, w \models \neg C_{\mathcal{A}}\neg\theta$ so also $\mathcal{M}, w \models \neg\alpha_{3.i.3}$.

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

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

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

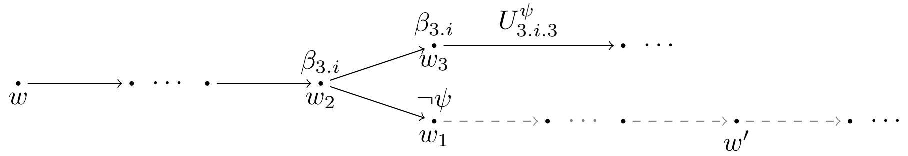

(a) Possibility 1: $\mathcal{M}, w_3 \models \beta_{3.i} \wedge \Diamond_{U^{\psi}_{3.i.3}} \top$.

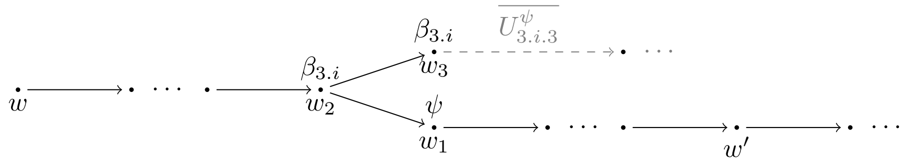

(b) Possibility 2: $\mathcal{M}, w_3 \models \beta_{3.i} \wedge \Box_{U^{\psi}_{3.i.3}} \bot$.

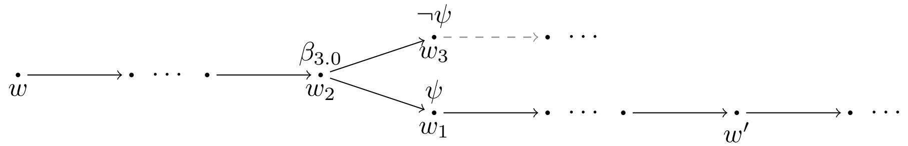

(c) Possibility 3: $\mathcal{M}, w_3 \models \wedge\Diamond^{i-1}(\Diamond\psi' \wedge \Diamond\neg\psi')$.

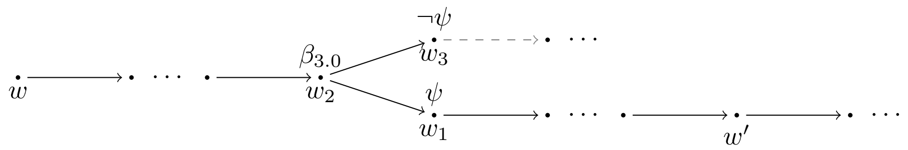

(d) Possibility 4: $\mathcal{M}, w_2 \models \Diamond\psi \wedge \Diamond\neg\psi$.

Figure A.4: The four possibilities for the conditions of Lemma A.9 to hold. In each case we have $\mathcal{M}, w \models \neg[U^{\psi}_{3.i.3}]C_{\mathcal{A}}\neg\theta$ for an appropriate choice of ψ.

A.2.4 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.

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

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

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

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 A.10. *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 $\mathcal{M}, w \models \neg[U^+_{2.i.0}]C_{\mathcal{A}}(\Box\bot \rightarrow p_i)$.

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

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

Let

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

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

and

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

Lemma A.11. *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 clause $\overline{(\beta_{2.i.j} \wedge \neg\varphi, \mathcal{A}, \top)}$ cannot apply in U-reachable worlds so the update simplifies to

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

Any $p_i \wedge \beta_{2.i.j}$ 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 $\mathcal{M}, w \models [U^{+\prime}_{2.i.j}]C_{\mathcal{A}}(\Box\bot \rightarrow \neg p_i)$ and therefore also $\mathcal{M}, w \models [U^+_{2.i.j}]C_{\mathcal{A}}(\Box\bot \rightarrow \neg p_i)$.

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 $\overline{(\beta_{2.i.j} \wedge \neg\varphi, \mathcal{A}, \top)}$ 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 $\mathcal{M}, w \models \neg[U^+_{2.i.j}]C_{\mathcal{A}}(\Box\bot \rightarrow \neg p_i)$.

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 $(u_1 \wedge \beta_{2.i.j}, a, u_2 \wedge \neg\zeta_{2.i.j+1})$ arrow. Arrows from w_2 are not retained however, because $\mathcal{M}, w_2 \models \neg\beta_{2.i.j} \wedge \neg\Diamond^j\neg p_i$. 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 $\mathcal{M}, w \models \neg[U^+_{2.i.j}]C_{\mathcal{A}}(\Box\bot \rightarrow \neg p_i)$.

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

Constructing α_2

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

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

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

We can also define α_2:

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

A.2.5 Constructing α_1

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

Let

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

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

and

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

Lemma A.12. *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 clause $\overline{(\beta_{1.i.j.k} \wedge \neg\varphi, \mathcal{A}, \top)}$ 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 $\mathcal{M}, w_2 \models \neg\beta_{1.i.j.k} \wedge \neg\Diamond^j\Diamond_{a_j}\top$. This implies that w_2 is a reachable $\Box\bot$ world after the update, so $\mathcal{M}, w \models \neg\alpha_{1.i.j.k}$. □

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 A.13. *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 $\mathcal{M}, w \models [U^j_{1.i.-1}]\neg C_{\mathcal{A}} \Box_{\mathcal{A}\backslash a_j} \Diamond \top$.

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 $\mathcal{M}, w \models [U^j_{1.i.-1}]\neg C_{\mathcal{A}} \Box_{\mathcal{A}\backslash a_j} \Diamond \top$. □

Constructing α_1

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

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

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 $\zeta_{1.i.j.k}$.

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

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

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

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

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

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

and

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

A.3 Proofs of auxiliary lemmas

Lemma A.14. *Let $\mathcal{M}$ be a model, w_1 and w_2 worlds in the model and ψ a CU 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 $b_1 \ne b_2$, $\mathcal{M}, w_3 \models \Diamond_{b_1} \top$, $\mathcal{M}, w_4 \models \Diamond_{b_2} \top$ and both w_3 and w_4 are reachable from w_1 in at most $d(\psi)$ steps.*

2. *there are a propositional variable $p \in$ Pvar(ψ) 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 length of ψ.

The base case is trivial; if ψ is of length 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 length or the negation of a formula of pure length.[2]

We can restrict ourselves to four possibilities for ψ that can together generate all CU formulas that are of pure length or the negation of a formula of pure length. It can be 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 $a \in \mathcal{A}, B \subseteq \mathcal{A}$, $\psi' \in \Phi^{d(\psi)-1}_{\mathrm{Pvar}(\psi)}$ and update U' with $d(U') < d(\psi)$.

- 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.
- 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.

[2]We could of course require ψ to be of pure length and still have it distinguish w_1 and w_2. Allowing ψ to be a negation of a formula of pure length allows us to guarantee that ψ holds in w_1 and not in 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 $\mathcal{M}, w_1 \models \neg\psi'$. Because $\mathcal{M}, w_2 \models \neg\neg[U']C_B\psi'$ we have $\mathcal{M}, w_2 \models \psi'$. 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. □

Lemma A.15. *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 length or the negation of a formula of pure length. 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 A.14 to $\neg\psi$, w_3 and w_2. The first two possibilities of Lemma A.14 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 length 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. □