Exponential Succinctness of Two-Agent Public Announcement Logic over S5

A direct proof using shared incidence models
(expanded version with detailed explanations)

Prompted by Yanjing Wang, Proof found by GPT 6 Astra,
report polished by Claude Opus 5.5 and GPT 6 Astra

14 October 2026; revised 15 October 2026

Abstract

We give a family of public-announcement formulas Φ_n, for $n \geq 1$, of size $O(n)$, over two agents and four atoms, such that every announcement-free modal formula equivalent to Φ_n over two-agent S5 models contains at least 2^n occurrences of the atom q, and hence has size at least 2^n. The proof uses one explicit finite model $\mathcal{M}_n$ with $19n-4$ worlds, together with a numerical invariant: a quantity $N(\theta)$ that every modal formula θ "pays for" with occurrences of q. This version spells out every step and is meant for readers who have seen basic modal logic, but not necessarily succinctness proofs.

Contents

1 Introduction and roadmap **2**

2 Preliminaries **3**
2.1 Two-agent S5 models 3
2.2 The modal language 4
2.3 Public announcements 4
2.4 Size and occurrences of q 4

3 The shared incidence model $\mathcal{M}_n$ **5**
3.1 Definition 5
3.2 Basic properties 6

4 The short PAL formula **7**
4.1 The announced formula χ 7
4.2 Round-by-round behaviour 7

5 Why we expect an exponential lower bound **9**

6 The occurrence-count invariant **9**

7 Tools: weighted cuts **10**
7.1 Cuts . . . 10
7.2 Boolean behaviour . . . 11
7.3 Boolean skeletons and the transfer principle . . . 11
7.4 Projection: reducing modal leaves to junctions . . . 12

8 Proof of the Invariant Lemma **12**
8.1 Atoms, constants and Boolean connectives . . . 12
8.2 The modal case: general recipe . . . 12
8.3 The a-diamond at rank $h \geq 1$ . . . 13
8.4 The a-diamond at rank 0 . . . 14
8.5 The b-diamond at rank $h \geq 1$ . . . 14
8.6 The b-diamond at rank 0 . . . 15
8.7 Conclusion of the induction . . . 16

9 Completing the proof of the Main Theorem **16**

10 Discussion **16**

References **18**

1 Introduction and roadmap

Public announcements. Epistemic logic talks about what agents know. The formula $K_a\alpha$ ("agent a knows α") is true at a world when α holds at every world that a cannot tell apart from the current one. *Public announcement logic* (PAL) adds formulas $\langle\alpha\rangle\beta$, read "$\alpha$ is true, and after α is publicly announced, β is true". Semantically, an announcement of α deletes every world where α is false.

For PAL without common knowledge, announcements add no *expressive power*: Plaza's reduction result, originally published in 1989 and reprinted in 2007, translates every PAL formula into an equivalent announcement-free epistemic formula [1, 2]. The rewriting can make formulas much longer, though. The question of *succinctness* is whether this blow-up can be avoided. The theorem below says it cannot. For some formulas, *every* announcement-free equivalent is exponentially longer.

Theorem 1 (Main Theorem)**.** *There is a sequence of public-announcement formulas Φ_n, indexed by $n \geq 1$, over exactly two agents a, b and the fixed atoms t, p_0, p_1, q, with $|\Phi_n| = O(n)$, such that every ordinary (announcement-free) modal formula θ with*

$$\models_{\mathbf{S5}_2} \ \Phi_n \leftrightarrow \theta$$

contains at least 2^n occurrences of q. In particular $|\theta| \geq 2^n$. Here formula size is the size of the syntax tree, with repeated occurrences counted separately. The lower bound already holds on finite two-agent S5 models; in fact on a single model with $19n - 4$ worlds.

Remark 2 (Previous succinctness results). Lutz [3, Theorem 2] proved that PAL is exponentially more succinct than basic modal logic over *unrestricted Kripke models*, using two agents. Van Ditmarsch, Fan, van der Hoek and Iliev [6] strengthened the unrestricted-model result to one modality and one propositional atom.

For *S5 models*, French, van der Hoek, Iliev and Kooi [4, 5] established exponential succinctness with *at least four agents*, using formula-size games. These agent and frame restrictions matter: the cited results do not themselves establish the two-agent S5 case. Theorem 1 treats exactly two agents and four fixed atoms, by a direct occurrence-counting argument on the explicit finite models below.

Roadmap. The proof has three ingredients.

(I) **A model** $\mathcal{M}_n$ (Section 3). It is a layered "graph" with $n+1$ ranks. At each rank $h \geq 1$ there is one *good* node A_h and two *faulty* nodes B_h^0, B_h^1. The outgoing-port pattern of B_h^c is that of A_h with one port removed: the colour-c port to the good node below. Their incoming classes may also differ; the invariant accounts for this.

(II) **A short PAL formula** Φ_n (Section 4). It announces one fixed, constant-size formula χ exactly $2n+1$ times. Each announcement "peels off" a layer of faultiness. After $2n+1$ rounds, A_n survives but B_n^0 and B_n^1 have been deleted. So Φ_n is true at A_n and false at B_n^0, B_n^1.

(III) **A counting invariant** $N(\theta)$ (Sections 6–8). For every announcement-free formula θ we define a number $N(\theta)$ that measures how well θ separates good nodes from faulty nodes, with separations at positive rank h weighted by 2^{h-1} and the rank-0 separation weighted by 1. We prove $N(\theta) \leq \mathrm{occ}_q(\theta)$, the number of occurrences of q in θ. Any formula that separates A_n from both B_n^0 and B_n^1 has $N(\theta) \geq 2^n$.

Putting these together: an equivalent θ of Φ_n must separate A_n from B_n^0, B_n^1, so $2^n \leq N(\theta) \leq \mathrm{occ}_q(\theta) \leq |\theta|$.

Proof idea. Repeated announcements apply the same constant-size test simultaneously throughout the current model. The lower-bound proof shows that every equivalent modal formula must contain exponentially many occurrences of q, by an invariant that accounts for both downward and upward movement. The size of a particular unrolling alone would not establish this claim.

2 Preliminaries

2.1 Two-agent S5 models

Definition 3. A *two-agent S5 model* is a tuple $\mathcal{M} = (W, \sim_a, \sim_b, V)$ where

- W is a non-empty set of *worlds*;
- $\sim_a$ and $\sim_b$ are *equivalence relations* on W (reflexive, symmetric and transitive);
- V assigns to each atom p a set $V(p) \subseteq W$ (the worlds where p is true).

We read $w \sim_i v$ as "agent i cannot distinguish w from v". An equivalence relation is the same thing as a *partition* of W into disjoint non-empty blocks, the *equivalence classes*. We write $[w]_i$ for the i-class of w. In what follows we always define $\sim_a$ and $\sim_b$ by listing their classes, and then check that the listed classes really partition W.

2.2 The modal language

Ordinary modal formulas (announcement-free) are built by

$$\theta ::= p \mid \top \mid \neg\theta \mid (\theta \wedge \theta) \mid (\theta \vee \theta) \mid \Diamond_a\theta \mid \Diamond_b\theta,$$

where p ranges over atoms. Truth at a world w of $\mathcal{M}$ is defined as usual:

$$\begin{aligned}
\mathcal{M}, w \models p &\iff w \in V(p),\\
\mathcal{M}, w \models \neg\theta &\iff \mathcal{M}, w \not\models \theta,\\
\mathcal{M}, w \models \theta \wedge \theta' &\iff \mathcal{M}, w \models \theta \text{ and } \mathcal{M}, w \models \theta', \quad (\vee \text{ similarly}),\\
\mathcal{M}, w \models \Diamond_i\theta &\iff \text{there is } v \in [w]_i \text{ with } \mathcal{M}, v \models \theta.
\end{aligned}$$

So $\Diamond_i\theta$ says "agent i considers θ possible". Knowledge is the abbreviation $K_i\alpha := \neg\Diamond_i\neg\alpha$ ("α holds throughout my class"). Expanding this knowledge abbreviation only adds negations, so it does not change the number of occurrences of any atom.

Fact 4 (Modal formulas are constant on classes)**.** If $\eta = \Diamond_i\beta$, then η has the same truth value at all worlds of any i-class.

Proof. If $w \sim_i v$ then $[w]_i = [v]_i$, and the truth of $\Diamond_i\beta$ at a world depends only on its i-class. □

We call a formula of the form $\Diamond_a\beta$ *a-modal* and one of the form $\Diamond_b\beta$ *b-modal*. Fact 4 is the class-invariance property used to estimate cuts of modal leaves in the lower-bound proof.

2.3 Public announcements

PAL adds formulas $\langle\alpha\rangle\beta$. If α is a formula and $\mathcal{M}$ a model, let $\mathcal{M}|\alpha$ be the model obtained by keeping only the worlds where α is true (in $\mathcal{M}$), restricting $\sim_a, \sim_b$ and V to them. Restrictions of equivalence relations are equivalence relations. If its domain is non-empty, $\mathcal{M}|\alpha$ is therefore an S5 model. We also allow the empty restriction as an auxiliary structure with no designated world; this causes no issue below because $\mathcal{M}, w \models \alpha$ guarantees that w survives. Then

$$\mathcal{M}, w \models \langle\alpha\rangle\beta \iff \mathcal{M}, w \models \alpha \text{ and } \mathcal{M}|\alpha, w \models \beta.$$

For a fixed χ write $\langle\chi\rangle^k\beta$ for $\langle\chi\rangle\cdots\langle\chi\rangle\beta$ with k announcements. Define $\mathcal{M}^0 = \mathcal{M}$ and $\mathcal{M}^r = \mathcal{M}^{r-1}|\chi$. Unwinding the definition k times gives:

Fact 5. $\mathcal{M}, w \models \langle\chi\rangle^k\top$ if and only if w is a world of $\mathcal{M}^k$, that is, w survives k successive announcements of χ.

We say that a world w is *deleted in round* r if $w \in \mathcal{M}^{r-1}$ but $w \notin \mathcal{M}^r$.

2.4 Size and occurrences of q

The size $|\theta|$ is the number of nodes in the syntax tree of θ. Repeated subformulas are counted every time they occur, so there is no sharing. For PAL we count each announcement as one binary syntax-tree node, so $|\langle\alpha\rangle\beta| = 1 + |\alpha| + |\beta|$. We write $\mathrm{occ}_q(\theta)$ for the number of occurrences of the atom q in θ. Clearly $\mathrm{occ}_q(\theta) \le |\theta|$, and

$$\mathrm{occ}_q(\neg\theta) = \mathrm{occ}_q(\Diamond_i\theta) = \mathrm{occ}_q(\theta), \qquad \mathrm{occ}_q(\theta \wedge \theta') = \mathrm{occ}_q(\theta \vee \theta') = \mathrm{occ}_q(\theta) + \mathrm{occ}_q(\theta').$$

Note that we only count q. Formulas may use t, p_0, p_1 and any other atoms as often as they like "for free". This makes the lower bound *stronger*.

Finally, $\models_{\mathbf{S5}_2} \Phi \leftrightarrow \theta$ means that Φ and θ have the same truth value at every world of every two-agent S5 model. We will only use this for *one* model, $\mathcal{M}_n$.

3 The shared incidence model $\mathcal{M}_n$

3.1 Definition

Fix $n \geq 1$. The model has two kinds of worlds: *junctions* (think: nodes of a graph) and *ports* (think: edges of a graph, each with a colour 0 or 1).

Junctions. They are arranged in ranks $0, 1, \ldots, n$:

$$J_0 = \{A_0,\ B_0\}, \qquad J_h = \{A_h,\ B_h^0,\ B_h^1\} \quad (1 \leq h \leq n).$$

We call A_h the *good* junction of rank h and the B's *bad* (or faulty) junctions. Note that B_0, with no superscript, is the bad junction of rank 0, while B_h^0 is the colour-0-faulty junction of rank h.

Ports. For $1 \leq h \leq n$, every $x \in J_h$, every $y \in J_{h-1}$ and every colour $c \in \{0, 1\}$, there is exactly one port $\pi^c(x, y)$ ("the colour-c edge from x down to y"), *except* when

$$x = B_h^c \quad \text{and} \quad y = A_{h-1}.$$

That is, the faulty junction B_h^c is missing exactly one edge: its colour-c edge to the good junction below. The tables show which colours of ports exist from a source at rank h to a target at rank $h-1$: the left table is for $h \geq 2$, and the right table covers $h = 1$.

source \ target	A_{h-1}	B_{h-1}^0	B_{h-1}^1
A_h	$0, 1$	$0, 1$	$0, 1$
B_h^0	1 only	$0, 1$	$0, 1$
B_h^1	0 only	$0, 1$	$0, 1$

$(h = 1)$	A_0	B_0
A_1	$0, 1$	$0, 1$
B_1^0	1 only	$0, 1$
B_1^1	0 only	$0, 1$

For a port $\pi = \pi^c(x, y)$ we write $\mathrm{src}(\pi) = x$ (its source) and $\mathrm{tgt}(\pi) = y$ (its target). For a junction x we set $\mathrm{src}(x) = \mathrm{tgt}(x) = x$.

Agents. The equivalence classes are:

$$C_a(x) = \{x\} \cup \{\pi : \pi \text{ a port with } \mathrm{src}(\pi) = x\}, \qquad C_b(y) = \{y\} \cup \{\pi : \pi \text{ a port with } \mathrm{tgt}(\pi) = y\},$$

for junctions x and y. In words: agent a knows the *source* of an edge, and confuses a node with its outgoing edges. Agent b knows the *target* of an edge, and confuses a node with its incoming edges.

Valuation.

$$V(t) = \bigcup_{h=0}^{n} J_h \text{ (all junctions)}, \qquad V(p_c) = \{\text{colour-}c \text{ ports}\}, \qquad V(q) = \{A_0\}.$$

All other atoms (if a proposed equivalent formula happens to use them) are false everywhere. Thus t says "I am a junction", p_c says "I am a port of colour c", and q marks the unique good junction at the bottom.

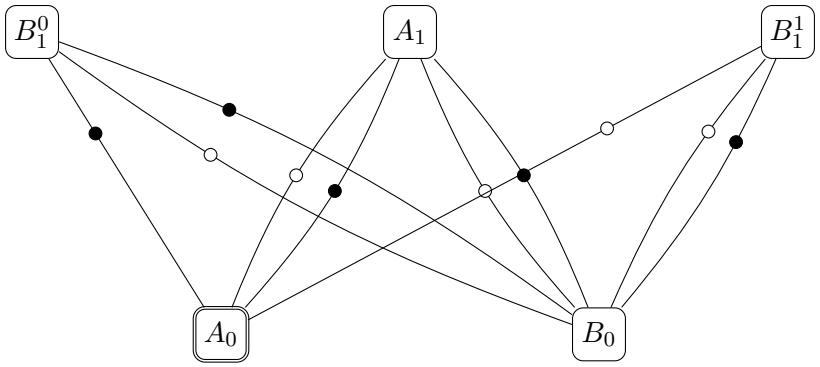


Figure 1: The model $\mathcal{M}_1$. Boxes are junctions; A_0 (double box) is the only q-world. Each small circle on a line is a *port* (a world of its own): white for colour 0, black for colour 1. The a-class of a junction is the junction together with the ports on lines going *down* from it. The b-class of a junction is the junction together with the ports on lines coming *into* it from above. B_1^0 lacks a white port to A_0; B_1^1 lacks a black port to A_0.

3.2 Basic properties

Lemma 6. *The families $\{C_a(x)\}_x$ and $\{C_b(y)\}_y$ (over all junctions x, y) are each partitions of the set of all worlds. Hence $\sim_a$ and $\sim_b$ are equivalence relations and $\mathcal{M}_n$ is a two-agent S5 model.*

Proof. Consider $\{C_a(x)\}$. A junction z lies in $C_a(z)$ and in no other $C_a(x)$, because the only junction in $C_a(x)$ is x. A port π lies in $C_a(\mathrm{src}(\pi))$ and in no other $C_a(x)$, because it has exactly one source. So every world is in exactly one block. The argument for $\{C_b(y)\}$ is identical, with targets instead of sources. □

Fact 7 (Where a world's classes are)**.** For every world w: its a-class is $C_a(\mathrm{src}(w))$ and its b-class is $C_b(\mathrm{tgt}(w))$.

Fact 8 (Which ports exist)**.** Let $1 \le h \le n$.

(P1) A_h has ports of both colours to every junction of rank $h-1$.

(P2) Every junction of rank h has ports of both colours to every *bad* junction of rank $h-1$.

(P3) B_h^c has a colour-$(1-c)$ port to A_{h-1}, but its colour-c ports go *only* to bad junctions of rank $h-1$.

(P4) The only missing ports are $\pi^c(B_h^c, A_{h-1})$. In particular, a port $\pi^c(x, y)$ with y bad always exists.

(P5) Rank-0 junctions have no outgoing ports, so $C_a(A_0) = \{A_0\}$ and $C_a(B_0) = \{B_0\}$. Rank-n junctions have no incoming ports, so $C_b(z) = \{z\}$ for $z \in J_n$.

Size. There are $2+3n$ junctions. Between ranks 1 and 0 there are $3 \cdot 2 \cdot 2 - 2 = 10$ ports. Between ranks h and $h-1$ for $h \ge 2$ there are $3 \cdot 3 \cdot 2 - 2 = 16$ ports. Hence

$$|\mathcal{M}_n| = (3n+2) + 10 + 16(n-1) = 19n - 4.$$

Remark 9 (Why "shared"). Each lower junction is *shared* by all incoming ports: for example A_{h-1} is the target of ports from A_h, B_h^0 and B_h^1. We do not make a fresh copy of the lower part for each parent, which would give a tree of exponential size. Sharing keeps the model of linear size. The price is that agent b, standing at a port into A_h, may consider possible other ports into A_h whose sources lie at rank $h+1$. So a formula can "walk upward" in the model, and the lower-bound proof must show that such upward walks do not help (Section 8.5, the "absorption" step).

4 The short PAL formula

4.1 The announced formula χ

$$\chi \;=\; \big(t \wedge (q \vee (\Diamond_a p_0 \wedge \Diamond_a p_1))\big) \;\vee\; \big(\neg t \wedge \Diamond_b t\big), \qquad \Phi_n \;=\; \underbrace{\langle\chi\rangle\langle\chi\rangle\cdots\langle\chi\rangle}_{2n+1 \text{ announcements}} \top.$$

The displayed formula χ has 15 syntax-tree nodes. Since it is written out $2n+1$ times, the size convention of Section 2.4 gives

$$|\Phi_n| = 1 + (2n+1)(1+|\chi|) = 32n + 17.$$

In particular, $|\Phi_n| = O(n)$ with every occurrence counted separately.

The first disjunct speaks about junctions, the second about ports. The next lemma makes this precise. It holds not only in $\mathcal{M}_n$ but in every model obtained from $\mathcal{M}_n$ by deleting worlds.

Lemma 10 (What χ says)**.** *Let $\mathcal{M}'$ be obtained from $\mathcal{M}_n$ by deleting some worlds (and restricting relations and valuation).*

(a) A junction $x \in \mathcal{M}'$ satisfies χ in $\mathcal{M}'$ iff $x = A_0$, or x still has, in $\mathcal{M}'$, at least one outgoing port of colour 0 and at least one outgoing port of colour 1.

(b) A port $\pi \in \mathcal{M}'$ satisfies χ in $\mathcal{M}'$ iff its target $\mathrm{tgt}(\pi)$ is still in $\mathcal{M}'$.

Proof. (a) At a junction t is true, so χ reduces to $q \vee (\Diamond_a p_0 \wedge \Diamond_a p_1)$. Now q is true only at A_0. In $\mathcal{M}'$ the a-class of x is $C_a(x) \cap \mathcal{M}'$, whose p_c-worlds are exactly the surviving colour-c ports with source x.

(b) At a port t is false, so χ reduces to $\Diamond_b t$. The b-class of π in $\mathcal{M}'$ is $C_b(\mathrm{tgt}(\pi)) \cap \mathcal{M}'$. The only t-world (junction) in $C_b(y)$ is y itself. So $\Diamond_b t$ holds iff $\mathrm{tgt}(\pi) \in \mathcal{M}'$. □

So each announcement of χ does two things simultaneously:

- it deletes every port whose target has already disappeared, and
- it deletes every junction other than A_0 that has lost all its outgoing ports of some colour.

4.2 Round-by-round behaviour

Let $\mathcal{M}^0 = \mathcal{M}_n$ and $\mathcal{M}^r = \mathcal{M}^{r-1}|\chi$. By Fact 5, Φ_n is true at w iff $w \in \mathcal{M}^{2n+1}$.

Define the candidate deletion times, which the next lemma verifies:

$$\delta(B_0) = 1, \qquad \delta(B_h^0) = \delta(B_h^1) = 2h+1 \quad (1 \le h \le n).$$

Thus $\delta(B_0) = 2\cdot 0 + 1$ fits the same pattern, and the candidate time for every bad junction of rank $h-1$ is $2h-1$.

Lemma 11 (Survival Lemma)**.** *For every $r \geq 0$:*

(i) every good junction A_h, and every port whose target is good, lies in $\mathcal{M}^r$;

(ii) a bad junction z lies in $\mathcal{M}^r$ iff $r < \delta(z)$, so z is deleted in round $\delta(z)$;

(iii) a port whose target is a bad junction z lies in $\mathcal{M}^r$ iff $r \leq \delta(z)$, so it is deleted in round $\delta(z) + 1$.

Consequently $\mathcal{M}_n, A_n \models \Phi_n$, while $\mathcal{M}_n, B_n^0 \not\models \Phi_n$ and $\mathcal{M}_n, B_n^1 \not\models \Phi_n$.

Proof. Step 1 (ports follow their targets). By Lemma 10(b), for $r \geq 1$ a port π is in $\mathcal{M}^r$ iff $\pi \in \mathcal{M}^{r-1}$ and $\mathrm{tgt}(\pi) \in \mathcal{M}^{r-1}$. Every original port lies in $\mathcal{M}^0$. Unfolding this recurrence gives, for $r \geq 1$,

$$\pi \in \mathcal{M}^r \iff \mathrm{tgt}(\pi) \in \mathcal{M}^s \text{ for every } 0 \leq s < r \iff \mathrm{tgt}(\pi) \in \mathcal{M}^{r-1}. \qquad (*)$$

The last equivalence uses the fact that the domains only shrink. Thus, if a target is deleted, its ports are deleted exactly one round later. Hence (iii) follows from (ii), and the port part of (i) follows from the junction part of (i). A port may survive even after its source has been deleted: restriction removes worlds individually, and the port clause of χ tests only its target.

Step 2 (junctions, by induction on the rank). By Lemma 10(a), for $r \geq 1$ a junction x is in $\mathcal{M}^r$ iff $x \in \mathcal{M}^{r-1}$ and ($x = A_0$ or x has outgoing ports of both colours in $\mathcal{M}^{r-1}$).

- A_0 satisfies q, so it is never deleted.
- B_0 has no q and no outgoing ports at all, so it is deleted in round $1 = \delta(B_0)$.
- A_h ($h \geq 1$): by induction A_{h-1} is never deleted, so by $(*)$ the ports $\pi^0(A_h, A_{h-1})$ and $\pi^1(A_h, A_{h-1})$ are never deleted. Hence A_h always has outgoing ports of both colours and is never deleted.
- B_h^c ($h \geq 1$): its colour-$(1-c)$ port to A_{h-1} is never deleted. Its colour-c ports exist by Fact 8(P2), and all go to bad junctions of rank $h-1$ by (P3). By induction these are all deleted in round $2h-1$, so by $(*)$ the colour-c ports lie in $\mathcal{M}^s$ exactly for $s \leq 2h-1$. Therefore B_h^c retains both colours at the start of every round $1 \leq r \leq 2h$ and survives all those rounds. At the start of round $2h+1$, its colour-c ports are absent, so this is its first deletion round: $\delta(B_h^c) = 2h+1$.

Incoming ports never matter for a junction's survival, since χ only looks at a junction's a-class, i.e. its outgoing ports.

Conclusion. $A_n \in \mathcal{M}^{2n+1}$, while $\delta(B_n^0) = \delta(B_n^1) = 2n+1$ means both are deleted in round $2n+1$, so they are not in $\mathcal{M}^{2n+1}$. □

Example 12 ($n = 1$; see Figure 1)**.** $\Phi_1 = \langle\chi\rangle\langle\chi\rangle\langle\chi\rangle\top$. Round 1: B_0 disappears (no q, no outgoing ports). Everything else survives. Round 2: the six ports into B_0 disappear, since their b-class no longer contains a junction. Round 3: B_1^0 disappears, because its only colour-0 port went to B_0 and is gone. Symmetrically B_1^1 disappears. A_1 survives, because its ports to A_0 are still there. What remains: A_0, A_1 and the four ports into A_0.

Example 13 ($n = 2$)**.**

worlds	deleted in round
B_0	1
the 6 ports into B_0	2
$B_1^0,\ B_1^1$	3
the 12 ports into B_1^0, B_1^1	4
$B_2^0,\ B_2^1$	5
all other worlds	never

Φ_2 has 5 announcements, so it is true at A_2 and false at B_2^0, B_2^1.

Remark 14 (Why the PAL formula is short). Information about "faultiness" travels upward at a speed of two rounds per rank: one round to delete the ports into a dead junction, and one round to delete the junctions that thereby lost a colour. The same constant-size test χ does this work at every rank, simultaneously. Announcements let us *re-use* χ by iterating it.

5 Why we expect an exponential lower bound

Before the formal proof, here is the natural announcement-free way to separate A_h from B_h^0, B_h^1. Define

$$G_0 = q, \qquad G_h = t \wedge \Diamond_a\big(p_0 \wedge \Diamond_b G_{h-1}\big) \wedge \Diamond_a\big(p_1 \wedge \Diamond_b G_{h-1}\big) \quad (h \geq 1).$$

Read G_h as: "I am a junction, I have a colour-0 port whose target satisfies G_{h-1}, and I have a colour-1 port whose target satisfies G_{h-1}." (Because G_{h-1} starts with t, or is q, it can only be true at a junction. The only junction in the b-class of a port is its target.)

By induction on h, G_h is true at A_h (use its two ports to A_{h-1}) and false at B_h^c, because all colour-c ports of B_h^c go to bad junctions, where G_{h-1} is false. But G_h contains *two* copies of G_{h-1}, so $\mathrm{occ}_q(G_h) = 2^h$.

The outgoing-port patterns of A_h and B_h^c differ in the colour-c port to A_{h-1}. This explains the two branches of the displayed formula G_h, but it does not prove that every separator must have this form: their incoming classes can also differ, and a modal formula can inspect many worlds or move upward through shared junctions. The invariant below handles all such formulas. Together with that invariant, G_n proves exact tightness in the number of q-occurrences for separating A_n from B_n^0, B_n^1. No global equivalence between G_n and Φ_n is claimed.

6 The occurrence-count invariant

From now on all truth values are computed in the *original* model $\mathcal{M}_n$; there are no announcements in θ. For a modal formula θ and world x write

$$[\![\theta]\!](x) \in \{0, 1\}$$

for the truth value of θ at x in $\mathcal{M}_n$ (1 = true).

Definition 15 (Rank discrepancies and the invariant)**.**

$$\begin{aligned} d_0(\theta) &= \big|[\![\theta]\!](A_0) - [\![\theta]\!](B_0)\big| \in \{0,1\}, \\ d_h(\theta) &= \sum_{i=0}^{1} \big|[\![\theta]\!](A_h) - [\![\theta]\!](B_h^i)\big| \in \{0,1,2\} \qquad (1 \le h \le n), \\ w_h &= 2^{h-1} \qquad (1 \le h \le n), \\ N(\theta) &= \max\big(d_0(\theta),\ w_1 d_1(\theta),\ \ldots,\ w_n d_n(\theta)\big). \end{aligned}$$

In words: $d_h(\theta)$ counts how many of the two faulty junctions of rank h are separated from the good one by θ. A separation at rank h is "worth" $w_h = 2^{h-1}$. $N(\theta)$ is the most valuable rank at which θ separates. Note $w_1 = 1$ and $w_{h+1} = 2w_h$.

Example 16. $N(q) = 1$, since q separates A_0 from B_0 and is false at all junctions of rank ≥ 1. For every other atom p, $N(p) = 0$, since p is constant on junctions. For G_h from Section 5, with $1 \le h \le n$, $d_h(G_h) = 2$, so $N(G_h) \ge 2^{h-1} \cdot 2 = 2^h = \mathrm{occ}_q(G_h)$.

Lemma 17 (Invariant Lemma)**.** *Every ordinary modal formula θ satisfies*

$$N(\theta) \ \le\ \mathrm{occ}_q(\theta).$$

The proof (Section 8) is by induction on θ. Atoms and Boolean connectives are easy. The real work is the modal operators: we have to show that a diamond cannot "create" separations that are not already paid for by the q's inside it.

7 Tools: weighted cuts

7.1 Cuts

Definition 18. A *comparison set* E is a finite multiset of triples (x, y, λ) where x, y are worlds of $\mathcal{M}_n$ and $\lambda > 0$ is a weight. We read (x, y, λ) as "compare x with y, with weight λ". The *cut* of a formula α along E is

$$\mathrm{cut}_E(\alpha) = \sum_{(x,y,\lambda)\in E} \lambda \,\big|[\![\alpha]\!](x) - [\![\alpha]\!](y)\big|,$$

the total weight of the comparisons on which α disagrees. The *total weight* of E is $\|E\| = \sum \lambda$; obviously $\mathrm{cut}_E(\alpha) \le \|E\|$.

Stars. The discrepancies are themselves cuts. Define the *rank-h star*

$$S_0 = \{(A_0, B_0, 1)\}, \qquad S_h = \{(A_h, B_h^0, w_h),\ (A_h, B_h^1, w_h)\} \quad (h \ge 1).$$

Then $d_0(\theta) = \mathrm{cut}_{S_0}(\theta)$ and $w_h d_h(\theta) = \mathrm{cut}_{S_h}(\theta)$, so

$$N(\theta) = \max_{0\le h\le n} \mathrm{cut}_{S_h}(\theta).$$

For $k \ge 1$, we say F is a *sub-star* of S_k if every comparison in F is between A_k and some B_k^i (in either order), and for each i the total weight of comparisons between A_k and B_k^i is at most the weight of that pair in S_k. For $k = 0$, a sub-star consists only of comparisons between A_0 and B_0 with total weight at most 1. Clearly

$$F \text{ a sub-star of } S_k \implies \mathrm{cut}_F(\theta) \le \mathrm{cut}_{S_k}(\theta) \le N(\theta). \qquad \text{(Star)}$$

Reversing the endpoints of a comparison leaves every cut unchanged, because the difference is taken in absolute value. Comparisons (u, u, λ) of a world with itself contribute 0 to every cut and can always be ignored.

7.2 Boolean behaviour

Lemma 19. *For every comparison set E:*

$$\mathrm{cut}_E(\neg\alpha) = \mathrm{cut}_E(\alpha), \qquad \mathrm{cut}_E(\alpha \circ \beta) \le \mathrm{cut}_E(\alpha) + \mathrm{cut}_E(\beta) \quad (\circ \in \{\wedge, \vee\}).$$

The same holds for d_h, and hence $N(\neg\alpha) = N(\alpha)$ and $N(\alpha \circ \beta) \le N(\alpha) + N(\beta)$.

Proof. It suffices to check one comparison (x, y). Negation flips both values, so it preserves whether they differ. For $\alpha \wedge \beta$: if α agrees on x, y and β agrees on x, y, then so does $\alpha \wedge \beta$. So whenever $\alpha \wedge \beta$ disagrees (contributing 1), at least one of α, β disagrees (contributing at least 1 on the right). The case of $\vee$ is the same. Multiply by λ and sum.

For N: $\mathrm{cut}_{S_h}(\alpha \circ \beta) \le \mathrm{cut}_{S_h}(\alpha) + \mathrm{cut}_{S_h}(\beta) \le N(\alpha) + N(\beta)$ for every h; take the maximum over h. □

7.3 Boolean skeletons and the transfer principle

Any formula α can be viewed as a Boolean combination (using $\neg, \wedge, \vee$) of *leaves*. A leaf is an occurrence of an atom or of $\top$ that is not inside any diamond, or a *maximal* modal subformula occurrence, i.e. an occurrence of some $\Diamond_i\beta$ that is not inside another diamond. For example, the leaves of $q \wedge \neg(\Diamond_a p_0 \vee \Diamond_b(q \wedge \Diamond_a q))$ are q, $\Diamond_a p_0$ and $\Diamond_b(q \wedge \Diamond_a q)$.

Every occurrence of q in α lies in exactly one leaf occurrence, so

$$\mathrm{occ}_q(\alpha) = \sum_{\text{leaf occurrences } \ell} \mathrm{occ}_q(\ell).$$

Applying Lemma 19 repeatedly gives:

Lemma 20 (Transfer principle)**.** *Let E be a comparison set and α a formula. Suppose*

(T1) every atomic or constant leaf ℓ of α satisfies $\mathrm{cut}_E(\ell) \le \mathrm{occ}_q(\ell)$, and

(T2) every modal leaf η of α satisfies $\mathrm{cut}_E(\eta) \le \mathrm{occ}_q(\eta)$.

Then $\mathrm{cut}_E(\alpha) \le \mathrm{occ}_q(\alpha)$.

Proof. $\mathrm{cut}_E(\alpha) \le \sum_\ell \mathrm{cut}_E(\ell) \le \sum_\ell \mathrm{occ}_q(\ell) = \mathrm{occ}_q(\alpha)$. □

Condition (T1) is automatic when the two worlds in each comparison have the same atomic valuation, since then every atomic leaf has cut 0; the constant $\top$ always has cut 0. In the induction, (T2) will be obtained in the form $\mathrm{cut}_E(\eta) \le N(\eta)$, combined with the induction hypothesis $N(\eta) \le \mathrm{occ}_q(\eta)$.

7.4 Projection: reducing modal leaves to junctions

Fact 4 and Fact 7 give a simple way to compute cuts of modal leaves. If η is a-modal, its value at w equals its value at $\mathrm{src}(w)$, since both lie in $C_a(\mathrm{src}(w))$. If η is b-modal, its value at w equals its value at $\mathrm{tgt}(w)$.

Lemma 21 (Projection)**.** *For a comparison set E let*

$$\mathrm{src}(E) = \{(\mathrm{src}(x), \mathrm{src}(y), \lambda) : (x, y, \lambda) \in E\}, \qquad \mathrm{tgt}(E) = \{(\mathrm{tgt}(x), \mathrm{tgt}(y), \lambda) : (x, y, \lambda) \in E\}.$$

These are multiset images*: each occurrence of a comparison in E contributes one projected occurrence, even if several projected triples coincide. Thus multiplicities and weights are preserved. If η is a-modal then $\mathrm{cut}_E(\eta) = \mathrm{cut}_{\mathrm{src}(E)}(\eta)$. If η is b-modal then $\mathrm{cut}_E(\eta) = \mathrm{cut}_{\mathrm{tgt}(E)}(\eta)$.*

So to bound the cut of an a-modal leaf, replace every port by its source. To bound a b-modal leaf, replace every port by its target. Then compare with the stars using (Star). In particular:

- two ports with the *same source* contribute 0 for every a-modal leaf;
- two ports with the *same target* contribute 0 for every b-modal leaf.

8 Proof of the Invariant Lemma

We prove $N(\theta) \le \mathrm{occ}_q(\theta)$ by induction on the construction of θ. The induction hypothesis (IH) is that $N(\eta) \le \mathrm{occ}_q(\eta)$ for every proper subformula η of θ.

8.1 Atoms, constants and Boolean connectives

- $\theta = q$: q is true at A_0 only, so $d_0(q) = 1$ and $d_h(q) = 0$ for $h \ge 1$. Hence $N(q) = 1 = \mathrm{occ}_q(q)$.
- θ another atom, or $\top$: constant on junctions (every junction satisfies t and no p_c; other atoms are false everywhere). So $N(\theta) = 0 = \mathrm{occ}_q(\theta)$.
- $\theta = \neg\alpha$, $\alpha \wedge \beta$, $\alpha \vee \beta$: by Lemma 19 and the IH, $N(\neg\alpha) = N(\alpha) \le \mathrm{occ}_q(\alpha) = \mathrm{occ}_q(\neg\alpha)$ and $N(\alpha \circ \beta) \le N(\alpha) + N(\beta) \le \mathrm{occ}_q(\alpha) + \mathrm{occ}_q(\beta) = \mathrm{occ}_q(\alpha \circ \beta)$.

8.2 The modal case: general recipe

Let $\theta = \Diamond_i\alpha$, so $\mathrm{occ}_q(\theta) = \mathrm{occ}_q(\alpha)$. Every modal leaf of the Boolean skeleton of α is a proper subformula of θ, even when that leaf is α itself. The IH therefore applies to every such leaf. Since N is a maximum over ranks, it suffices to fix one rank h and prove

$$\mathrm{cut}_{S_h}(\Diamond_i\alpha) \;\le\; \mathrm{occ}_q(\alpha). \tag{Goal}$$

For each positive rank $h \ge 1$ we use the following recipe. The rank-zero cases, whose comparison has weight 1, are handled separately below:

Step 1. Let $I = \{i' \in \{0,1\} : \Diamond_i\alpha$ differs at A_h and $B_h^{i'}\}$. If $I = \emptyset$ there is nothing to prove, since the left side of (Goal) is 0.

Step 2. The value of $\Diamond_i\alpha$ at A_h is a single bit. So either it is *true at A_h and false at every $B_h^{i'}$ with $i' \in I$*, or *false at A_h and true at every $B_h^{i'}$ with $i' \in I$*. We treat these two subcases separately.

Step 3. Using witnesses for the diamond, build a comparison set E in which every comparison has one α-true and one α-false world, and the total weight is $w_h|I|$. Then

$$\mathrm{cut}_E(\alpha) = \|E\| = w_h|I| = \mathrm{cut}_{S_h}(\Diamond_i\alpha).$$

Step 4. Check (T1): compared worlds have the same atoms (except in one rank-0 situation, treated separately).

Step 5. Check (T2): every modal leaf η of α satisfies $\mathrm{cut}_E(\eta) \le N(\eta)$, using Projection and (Star). By the IH, $N(\eta) \le \mathrm{occ}_q(\eta)$.

Step 6. By the Transfer principle, $\mathrm{cut}_{S_h}(\Diamond_i\alpha) = \mathrm{cut}_E(\alpha) \le \mathrm{occ}_q(\alpha)$.

Recall what a diamond means. If $\Diamond_i\alpha$ is *true* at w, there is a *witness*: a world in the i-class of w where α is true. If $\Diamond_i\alpha$ is *false* at w, then α is false at *every* world of the i-class of w. Every comparison we build pairs a witness (where α is true) with a world from a class where $\Diamond_i\alpha$ is false (where α is false). So every comparison disagrees on α, as required in Step 3.

Steps 1, 2 and 6 are the same at every positive rank. We now verify Steps 3–5 for both operators and treat the rank-zero cases separately.

8.3 The a-diamond at rank $h \ge 1$

Here $C_a(A_h) = \{A_h\} \cup \{\pi^c(A_h, y)\}$ and $C_a(B_h^i) = \{B_h^i\} \cup \{\pi^c(B_h^i, y)\}$: the junction together with its ports down to rank $h-1$.

Subcase A1: $\Diamond_a\alpha$ true at A_h. Then α is false throughout $C_a(B_h^i)$ for all $i \in I$. Choose *one* α-true witness $x \in C_a(A_h)$ and use the same x for all $i \in I$. (Why this matters is explained in the b-leaf check below.) For each $i \in I$ add comparisons as follows:

situation	comparisons for this i	src-projection	tgt-projection
(1) $x = A_h$	(A_h, B_h^i, w_h)	(A_h, B_h^i, w_h)	(A_h, B_h^i, w_h)
(2) $x = \pi^c(A_h, y)$ and $\pi^c(B_h^i, y)$ exists	$(x, \pi^c(B_h^i, y), w_h)$	(A_h, B_h^i, w_h)	(y, y): zero
(3) otherwise; then $c = i$, $y = A_{h-1}$, and $h \ge 2$	$(x, \pi^i(B_h^i, B_{h-1}^0), w_{h-1})$, $(x, \pi^i(B_h^i, B_{h-1}^1), w_{h-1})$	(A_h, B_h^i), total weight $2w_{h-1} = w_h$	the lower-rank star S_{h-1}
(3') same, with $h = 1$	$(x, \pi^i(B_1^i, B_0), 1)$	$(A_1, B_1^i, 1)$	$(A_0, B_0, 1)$

Why the comparisons are legitimate. In (2), the partner port exists by assumption. In (3), the only missing port is $\pi^i(B_h^i, A_{h-1})$ (Fact 8), so the case occurs exactly when $c = i$ and $y = A_{h-1}$. The ports $\pi^i(B_h^i, z)$ for bad z exist by (P2). All partners lie in $C_a(B_h^i)$, so they are α-false, while x is α-true. Each i contributes total weight w_h; this uses $w_h = 2w_{h-1}$ for $h \ge 2$ and $w_1 = 1$. So $\mathrm{cut}_E(\alpha) = w_h|I|$.

(T1) Atoms. Compared worlds are either two junctions of rank $h \ge 1$ (all atoms agree; q is false at both) or two ports of the same colour (all atoms agree).

(T2) a-modal leaves η. By Projection, $\mathrm{src}(E)$ consists of comparisons (A_h, B_h^i) with total weight w_h for each $i \in I$: a sub-star of S_h. So $\mathrm{cut}_E(\eta) \le \mathrm{cut}_{S_h}(\eta) \le N(\eta)$.

(T2) b-modal leaves η. Look at the last column.

- If the witness is the junction $x = A_h$, only situation (1) occurs, and $\mathrm{tgt}(E)$ is a sub-star of S_h.
- If the witness is a port $x = \pi^c(A_h, y)$, situation (1) never occurs. Situation (2) contributes zero. Situation (3)/(3') can occur only for the single index $i = c$, because the witness has one fixed colour c. After discarding diagonal comparisons, $\mathrm{tgt}(E)$ is (at most) one copy of S_{h-1}, and $\mathrm{cut}_E(\eta) \le \mathrm{cut}_{S_{h-1}}(\eta) \le N(\eta)$.

Here we see why one common witness is used for all i. With two different port witnesses of colours 0 and 1, both indices could fall into situation (3). The b-leaf would then see *two* copies of S_{h-1}, which $N(\eta)$ need not cover.

Subcase A2: $\Diamond_a\alpha$ false at A_h. Then α is false throughout $C_a(A_h)$. For each $i \in I$, $\Diamond_a\alpha$ is true at B_h^i; choose an α-true witness $x_i \in C_a(B_h^i)$.

situation	comparison	src-projection	tgt-projection
$x_i = B_h^i$	(x_i, A_h, w_h)	(B_h^i, A_h, w_h)	(B_h^i, A_h, w_h)
$x_i = \pi^c(B_h^i, y)$	$(x_i, \pi^c(A_h, y), w_h)$	(B_h^i, A_h, w_h)	(y, y): zero

The partner $\pi^c(A_h, y)$ always exists by (P1) and lies in $C_a(A_h)$, so it is α-false. Atoms agree as before. For a-modal leaves, $\mathrm{src}(E)$ is a sub-star of S_h. For b-modal leaves, $\mathrm{tgt}(E)$ consists of zeros and a sub-star of S_h. Both bounds follow from (Star).

8.4 The a-diamond at rank 0

By (P5), $C_a(A_0) = \{A_0\}$ and $C_a(B_0) = \{B_0\}$. So $[\![\Diamond_a\alpha]\!](A_0) = [\![\alpha]\!](A_0)$ and likewise at B_0. Hence $d_0(\Diamond_a\alpha) = d_0(\alpha) \le N(\alpha) \le \mathrm{occ}_q(\alpha)$ by the IH.

8.5 The b-diamond at rank $h \ge 1$

For $h < n$, $C_b(A_h)$ consists of A_h together with the existing ports into it from rank $h+1$; similarly for $C_b(B_h^i)$. The top rank has no incoming ports.

Top rank $h = n$. By (P5) all b-classes of rank-n junctions are singletons. So $\Diamond_b\alpha$ has the same values as α on J_n, and $\mathrm{cut}_{S_n}(\Diamond_b\alpha) = \mathrm{cut}_{S_n}(\alpha) \le N(\alpha) \le \mathrm{occ}_q(\alpha)$. Assume from now on that $1 \le h \le n-1$.

Subcase B1: $\Diamond_b\alpha$ true at A_h. Choose one α-true witness $x \in C_b(A_h)$. For $i \in I$, α is false throughout $C_b(B_h^i)$.

situation	comparison for this i	src-projection	tgt-projection
$x = A_h$	(A_h, B_h^i, w_h)	(A_h, B_h^i, w_h)	(A_h, B_h^i, w_h)
$x = \pi^c(u, A_h)$	$(x, \pi^c(u, B_h^i), w_h)$	(u, u): zero	(A_h, B_h^i, w_h)

The partner $\pi^c(u, B_h^i)$ exists because every upper source has ports of both colours to every bad target (P2). It lies in $C_b(B_h^i)$, so it is α-false. Atoms agree. For b-modal leaves, $\mathrm{tgt}(E)$ is a sub-star of S_h. For a-modal leaves, $\mathrm{src}(E)$ is zeros plus a sub-star of S_h. Done by (Star).

Subcase B2: $\Diamond_b\alpha$ false at A_h (the backward-movement case). Now α is false throughout $C_b(A_h)$. For each $i \in I$ choose an α-true witness $x_i \in C_b(B_h^i)$.

situation	comparison	src-projection	tgt-projection
(1) $x_i = B_h^i$	(x_i, A_h, w_h)	(B_h^i, A_h, w_h)	(B_h^i, A_h, w_h)
(2) $x_i = \pi^c(u, B_h^i)$ and $\pi^c(u, A_h)$ exists	$(x_i, \pi^c(u, A_h), w_h)$	(u, u): zero	(B_h^i, A_h, w_h)
(3) otherwise; then $u = B_{h+1}^c$	$(x_i, \pi^c(A_{h+1}, A_h), w_h)$	$(B_{h+1}^c, A_{h+1}, w_h)$	(B_h^i, A_h, w_h)

In (3), the missing port into A_h of colour c from rank $h+1$ is exactly $\pi^c(B_{h+1}^c, A_h)$ (Fact 8). The replacement partner $\pi^c(A_{h+1}, A_h)$ exists by (P1). All partners lie in $C_b(A_h)$, so they are α-false. Each comparison has weight w_h, so $\mathrm{cut}_E(\alpha) = w_h|I|$, and

$$\|E\| = w_h|I| \le 2w_h = w_{h+1}.$$

Atoms agree: junction pairs are at rank $h \ge 1$, and port pairs have the same colour.

b-modal leaves. In every situation the tgt-projection is (B_h^i, A_h, w_h), so $\mathrm{tgt}(E)$ is a sub-star of S_h. Done by (Star).

a-modal leaves: the absorption argument. Now $\mathrm{src}(E)$ may contain *upper-rank* comparisons $(B_{h+1}^c, A_{h+1}, w_h)$ from situation (3), besides current-rank comparisons from (1) and zeros from (2). This is the effect of moving "backward" (upward) through the shared junction: the witness is a port coming from the faulty junction B_{h+1}^c above. There are two possibilities.

- If η agrees on every upper-rank comparison (A_{h+1}, B_{h+1}^c) that occurs, those comparisons contribute nothing. What is left is a sub-star of S_h, so $\mathrm{cut}_E(\eta) \le \mathrm{cut}_{S_h}(\eta) \le N(\eta)$.
- If η disagrees on some occurring upper-rank comparison, then $d_{h+1}(\eta) \ge 1$, so

$$N(\eta) \;\ge\; w_{h+1}d_{h+1}(\eta) \;\ge\; w_{h+1} \;=\; 2w_h \;\ge\; \|E\| \;\ge\; \mathrm{cut}_E(\eta).$$

In the second possibility the *entire* cut, including any current-rank comparisons, is absorbed by the single upper-rank separation. This is why the weights double from rank to rank. A separation one rank higher is already worth as much as everything the current rank can ever demand ($|I| \le 2$ comparisons of weight w_h). So walking upward never makes a separation cheaper.

8.6 The b-diamond at rank 0

There is one comparison pair (A_0, B_0) with weight 1. Suppose $\Diamond_b\alpha$ differs at A_0 and B_0 (otherwise there is nothing to prove). We build E consisting of *one* comparison of weight 1, by the same witness rules as before:

situation	comparison	src-proj.	tgt-proj.
true at A_0, witness A_0	$(A_0, B_0, 1)$	$(A_0, B_0, 1)$	$(A_0, B_0, 1)$
true at A_0, witness $\pi^c(u, A_0)$	$(x, \pi^c(u, B_0), 1)$	zero	$(A_0, B_0, 1)$
false at A_0, witness B_0	$(B_0, A_0, 1)$	$(B_0, A_0, 1)$	$(B_0, A_0, 1)$
false at A_0, witness $\pi^c(u, B_0)$, partner $\pi^c(u, A_0)$ exists	$(x, \pi^c(u, A_0), 1)$	zero	$(B_0, A_0, 1)$
false at A_0, witness $\pi^c(B_1^c, B_0)$ (partner missing)	$(x, \pi^c(A_1, A_0), 1)$	$(B_1^c, A_1, 1)$	$(B_0, A_0, 1)$

All partners exist (by (P1), (P2)), and the comparison disagrees on α, so $\mathrm{cut}_E(\alpha) = 1 = d_0(\Diamond_b \alpha)$.

(T1) Atoms: the bottom-rank q-cost. If the comparison is between ports, they have the same colour and all atoms agree. If it is the junction pair (A_0, B_0) in either order, all atoms agree *except* q. Each atomic leaf occurrence of q has cut at most $\|E\| = 1 = \mathrm{occ}_q(q)$, and all other atomic leaves have cut 0. So (T1) holds. This is the modal transfer case where an atomic occurrence of q pays for a nonzero cut, just as q pays for the bottom-rank separation in the atomic base case.

(T2) Modal leaves. For a b-modal leaf η, the target projection is either $\{(A_0, B_0, 1)\}$ or $\{(B_0, A_0, 1)\}$. Both have the same cut as S_0, so

$$\mathrm{cut}_E(\eta) = \mathrm{cut}_{\mathrm{tgt}(E)}(\eta) = \mathrm{cut}_{S_0}(\eta) = d_0(\eta) \le N(\eta).$$

For an a-modal leaf, the source projection is a diagonal comparison (cut 0), the bottom junction pair in either order (cut $d_0(\eta)$), or $(B_1^c, A_1, 1)$. The last is a sub-star of S_1 because $w_1 = 1$, so its cut is at most $w_1 d_1(\eta) \le N(\eta)$.

By the Transfer principle, $d_0(\Diamond_b \alpha) = \mathrm{cut}_E(\alpha) \le \mathrm{occ}_q(\alpha)$.

8.7 Conclusion of the induction

For $\theta = \Diamond_a \alpha$ and for $\theta = \Diamond_b \alpha$ we have shown $\mathrm{cut}_{S_h}(\theta) \le \mathrm{occ}_q(\alpha) = \mathrm{occ}_q(\theta)$ for every rank $h \in \{0, \dots, n\}$. Taking the maximum over h gives $N(\theta) \le \mathrm{occ}_q(\theta)$. Together with the atomic and Boolean cases, this completes the induction and proves Lemma 17. □

9 Completing the proof of the Main Theorem

Proof of Theorem 1. The formula Φ_n has size $O(n)$. Let θ be any ordinary modal formula with $\models_{\mathbf{S5}_2} \Phi_n \leftrightarrow \theta$. Since $\mathcal{M}_n$ is a two-agent S5 model (Lemma 6), θ and Φ_n agree at every world of $\mathcal{M}_n$. By the Survival Lemma 11,

$$[\![\theta]\!](A_n) = 1, \qquad [\![\theta]\!](B_n^0) = [\![\theta]\!](B_n^1) = 0,$$

so $d_n(\theta) = 2$. By the Invariant Lemma 17,

$$|\theta| \;\ge\; \mathrm{occ}_q(\theta) \;\ge\; N(\theta) \;\ge\; w_n\, d_n(\theta) \;=\; 2^{n-1} \cdot 2 \;=\; 2^n. \qquad \square$$

10 Discussion

Separation on one model. Every ordinary modal formula that is true at A_n and false at B_n^0 and B_n^1 in the single model $\mathcal{M}_n$, which has $19n - 4$ worlds, must contain at least 2^n occurrences of q, however many occurrences of other atoms and of the permitted Boolean and a, b-modal operators it uses. Full equivalence with Φ_n is only used to get these three truth values.

Tightness. The formula G_n of Section 5 has exactly 2^n occurrences of q and separates A_n from B_n^0, B_n^1. Thus the lower bound is exact for this separation problem in the measure occ_q. Indeed, for $n \geq 1$ the point A_0 satisfies Φ_n but not G_n, so G_n is not equivalent to Φ_n even everywhere in $\mathcal{M}_n$.

The roles of the design choices.

- *Two colours* give two different faults at each positive rank. The explicit separator G_h uses one branch for each colour. The invariant proves that every separator incurs the same q-occurrence lower bound, without assuming that it has this form.
- *Weights* $w_h = 2^{h-1}$ encode the doubling. In Subcase A1, the missing-port comparisons project to a full rank-$(h-1)$ star: two comparisons of weight w_{h-1} when $h \geq 2$, or the single bottom pair when $h = 1$. Other a-diamond comparisons use current-rank stars or diagonal pairs. In Subcase B2, one upper-rank separation of weight $w_{h+1} = 2w_h$ covers the entire comparison multiset, including any current-rank contributions.
- *One common witness* in Subcase A1 confines the missing-port comparisons to one colour, so their target projections contribute at most one lower-rank star.
- *Shared junctions* make the model small but allow backward movement. The absorption argument in Subcase B2 shows that this movement cannot evade the occurrence bound.
- *The atom* q is the only atom that distinguishes the two rank-0 junctions. The atomic base case gives $N(q) = 1 = \mathrm{occ}_q(q)$; the rank-0 b-diamond case likewise charges each nonzero atomic cut to a q-occurrence. Boolean transfer and projection carry these bounds through the rest of the formula.

Summary of notation.

J_h	junctions of rank h: $\{A_0, B_0\}$ or $\{A_h, B_h^0, B_h^1\}$
$\pi^c(x, y)$	colour-c port from $x \in J_h$ to $y \in J_{h-1}$ (missing iff $x = B_h^c$, $y = A_{h-1}$)
$C_a(x)$, $C_b(y)$	x with its outgoing ports; y with its incoming ports
$\mathrm{src}, \mathrm{tgt}$	source/target of a port (identity on junctions)
χ, Φ_n	announced formula; $2n+1$ announcements of χ followed by $\top$
$\delta(z)$	deletion round of a bad junction: $2h+1$ at rank h
$[\![\theta]\!](x)$	truth value of θ at x in $\mathcal{M}_n$
d_h, w_h, N	rank discrepancy, weights for $h \geq 1$, invariant $\max(d_0, \max_{1 \leq h \leq n} w_h d_h)$
S_h	rank-h star; $N(\theta) = \max_h \mathrm{cut}_{S_h}(\theta)$
$\mathrm{cut}_E(\alpha)$	total weight of comparisons in E on which α disagrees
$\mathrm{occ}_q(\theta)$	number of occurrences of q in θ

References

[1] Jan A. Plaza. Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, editors, *Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems: Poster Session Program*, pages 201–216. Oak Ridge National Laboratory, 1989. Report ORNL/DSRD-24. Reprinted as [2].

[2] Jan Plaza. Logics of public communications. *Synthese*, 158(2):165–179, 2007. doi:10.1007/s11229-007-9168-7.

[3] Carsten Lutz. Complexity and succinctness of public announcement logic. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '06)*, pages 137–143. ACM, 2006. doi:10.1145/1160633.1160657.

[4] Tim French, Wiebe van der Hoek, Petar Iliev, and Barteld Kooi. Succinctness of epistemic languages. In Toby Walsh, editor, *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 881–886. AAAI Press/IJCAI, 2011. doi:10.5591/978-1-57735-516-8/IJCAI11-153.

[5] Tim French, Wiebe van der Hoek, Petar Iliev, and Barteld Kooi. On the succinctness of some modal logics. *Artificial Intelligence*, 197:56–85, 2013. doi:10.1016/j.artint.2013.02.003.

[6] Hans van Ditmarsch, Jie Fan, Wiebe van der Hoek, and Petar Iliev. Some exponential lower bounds on formula-size in modal logic. In Rajeev Goré, Barteld Kooi, and Agi Kurucz, editors, *Advances in Modal Logic*, volume 10, pages 139–157. College Publications, 2014. `https://www.aiml.net/volumes/volume10/Ditmarsch-Fan-Hoek-Iliev.pdf`.