Research Article

Current-State Interfaces and the Future Quotient in a Finite Icosahedral (AKN-Type) Attachment Geometry: Untouchable Differences and Persistent History Parity

Masaru Numagaki[1]

1. Medcure Co., Ltd, Japan

On a frozen 691-cell patch of AKN-type rhombohedral geometry with a single-face attachment grammar, a companion paper constructed finite current-state codes that decide the complete legal future, without determining whether the normalised code $R_K^{(2)}$ among them is minimal. Here we ask which differences between current states act on the legal future, and compare the partition of states by the normalised code $R_K^{(2)}$ with the future quotient, the partition by equality of extension languages.

Assuming that every face belongs to at most two catalogue cells (K1, an audited finite property of the frozen catalogue), we prove in Lean that $R_K^{(2)}$ is an exact complete invariant of the one-step legal interface: two catalogue states have the same normalised code exactly when they have the same available cells and the same legal moves. Together with an exhaustive enumeration, this yields all identifications within eight insertions of the seed. The future quotient is strictly coarser even among reachable states that still have legal moves: a preregistered computation on complete descendant DAGs certifies 5,384,241 future-equivalent pairs separated by $R_K^{(2)}$. Under the same assumption we prove that if a cell is available in a reachable state s but not in a state t, and some legal continuation of s touches it, then an explicit word distinguishes the futures of s and t; hence the availability differences absorbed by future equivalence lie in cells that the available side can never touch, and an exhaustive audit of all certified nonterminal pairs found no exception. Finally, an outcome-blind, preregistered test shows that a history-derived parity, evaluable from the current state, does not factor through the future quotient, even between states with identical normalised code and size; by an existing integration law, such a difference persists along every common legal continuation.

The future quotient itself is classical; the contribution is its explicit realisation in a finite attachment geometry, with formally verified general theorems and computationally certified reachable witnesses. No physical claim is made.

Corresponding author: Masaru Numagaki, m.numagaki@medcure.co.jp

 doi.org/10.32388/8C43EK

1. Introduction

1.1. Problem

A companion paper, Waterway I [1], fixed a frozen 691-cell patch of AKN-type rhombohedral geometry and a single-face attachment grammar on it, and constructed explicit finite codes of the current state that decide the complete legal future: the residue R_{cat} and its normalisation $R_K^{(2)}$, which erases permanently dead slots. Equal codes imply equal extension languages. Since $R_K^{(2)}$ identifies reachable states that R_{cat} separates [1], Main Theorem G(ii), R_{cat} is not Nerode-minimal; whether $R_K^{(2)}$ is, that is, whether equal extension languages imply equal $R_K^{(2)}$, was left open [1], Section 6.

This paper asks a sharper question: *which differences between current states act on the legal future?* The question is asked for two kinds of difference. The first is geometric: which catalogue cells are still available. The second is history-derived: the explicit potential Φ of [2], a parity that, up to its value at the seed, counts non-base-aligned (HIGH) insertions modulo 2 but can be evaluated from the current state alone. The reference partition throughout is the future quotient $\sim_F$, the partition of states by equality of complete extension languages.

1.2. Main results

The results form one chain. Status tags (Section 2.5): [M] Lean theorem; [M*] Lean theorem whose hypothesis is the audited frozen catalogue certificate; [M$_{\mathrm{paper}}$] short mathematical argument combining formal results, not itself formalised; [E] frozen, preregistered computation or audit.

1. *The normalised code is an exact invariant of the one-step interface* (Theorem 1) [M] under K1; [M*] for the frozen catalogue. On catalogue states, under K1, $R_K^{(2)}(s) = R_K^{(2)}(t)$ if and only if s and t have the same available catalogue cells and the same legal moves; R_{cat} is injective. Together with an exhaustive enumeration, the $R_K^{(2)}$-identifications among the $720{,}473$ states within eight insertions of the seed are exactly 562 classes (Corollary 1.1) [M*] + [E].
2. *The future quotient is strictly coarser* (Theorem 2) [E]. Among reachable states that still admit a legal move, there are future-equivalent pairs that $R_K^{(2)}$ separates; on the examined set there are $5{,}384{,}241$ such pairs in $280{,}816$ classes, each certified by exact comparison of complete canonical futures. By Theorem 1, every such pair has the same legal moves and different available cells (Proposition 1, under K1).
3. *Touchable differences separate futures* (Theorem 3) [M*]. If a cell is available in a reachable state s and not in a state t, and some legal continuation of s touches it, an explicit distinguishing word can be built from a prefix of that continuation. Hence the availability differences that survive future equivalence lie in cells that the available side can never touch (Corollaries 3.1 and 3.2). An exhaustive audit of all $5{,}384{,}241$ pairs of Theorem 2 found no exception [E].
4. *A current-state parity survives beyond the future quotient* (Theorem 4) [E]. In a preregistered, outcome-blind test, Φ splits future-equivalence classes in both strata defined in advance as structural, including pairs with identical normalised code and identical number of present cells. By the integration law of [2], such a parity difference is preserved along every common legal continuation (Corollary 4.1) [M$_{\mathrm{paper}}$].

Read together: under K1 the code $R_K^{(2)}$ is exactly the one-step interface (1); the complete legal future does not determine all of it (2); on reachable states, an availability difference between future-equivalent states lies in a cell that the available side can never touch (3), the converse being left open here (see the note added to Section 5.4); and beyond the future quotient there can remain a one-bit difference, readable from the current state and never visible to legality, that is preserved along every common legal continuation (4). Figure 1 summarises the chain.

doi.org/10.32388/8C43EK

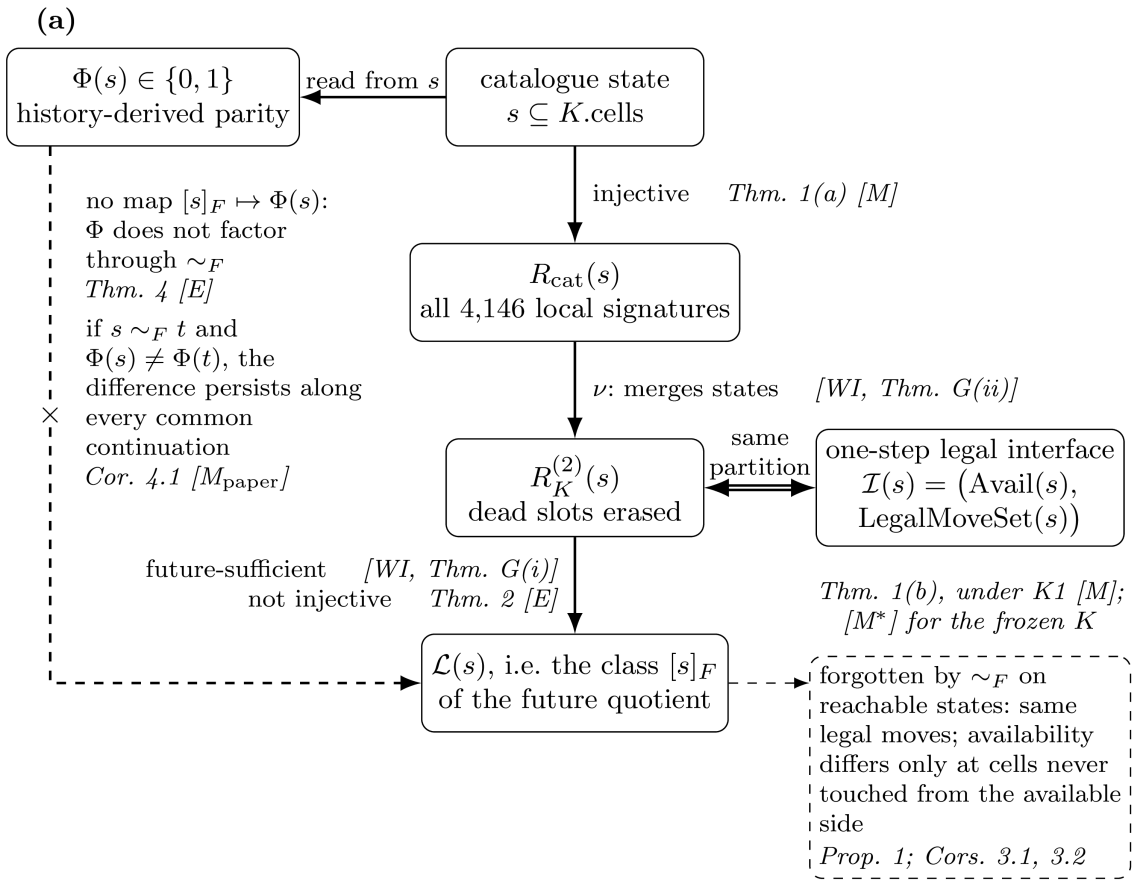


Figure 1. Overview of the results. (a) The chain of maps on a catalogue state. R_{cat} is injective; the normalisation ν merges states; $R_K^{(2)}$ determines the complete legal future and, under K1, induces the same partition as the one-step interface $\mathscr{I}$, but it is not minimal: the future quotient is strictly coarser. The differences that the future quotient forgets are confined to cells that the available side can never touch; this is a necessary condition, and its converse is not a result of this paper (see the note in Section 5.4). The parity Φ is read from the current state but is not a function of the future-equivalence class. WI: Waterway I [1]. Italic labels give the result and, where it is a single result, its evidence tag (Section 2.5). (b) Schematic nesting of the partitions. The diagram

combines in one picture situations found in different examined sets; it is not a single observed class, and the numbers and sizes of ellipses and dots do not represent any count.

1.3. Position

The future quotient itself is not new; it belongs to the Myhill–Nerode tradition [3][4] and, for this deterministic partial transition system, coincides with strong bisimulation [5][6]. The contribution here is an explicit realisation in a finite AKN-type attachment geometry, together with formally verified general theorems and computationally certified reachable witnesses showing which geometric and history-derived distinctions are preserved, identified, or remain invisible under exact future equivalence. Related work and the novelty boundary are discussed in Section 7.

1.4. Method

General statements are proved in Lean 4 [7] with Mathlib [8], on top of the archive of [1], or follow from Lean theorems by short arguments given in the text. The catalogue-specific hypotheses, that every face is a face of at most two catalogue cells (K1) and that the seed row is a catalogue cell, enter as a frozen certificate audited outside Lean. Existence statements on reachable states are established by computation under a preregistration fixed before the decisive quantities were computed; the depth-eight grouping and the terminal witness, which had been observed earlier, were reproduced as replication and control targets. The class lists were frozen before the parity Φ was evaluated on any of their states, except for one control pair known in advance. The scripts of the production stages were frozen before their production runs (with two recorded exceptions of procedure, Section 8.2), and every output is recorded in a hash ledger (Section 8).

1.5. Claim firewall

- *Scope of computations.* The counts in Theorems 2 and 4 and in the audits concern the examined sets only; the existence conclusions concern reachable states. No frequency in the reachable state space is claimed.
- *Catalogue facts.* The face-owner bound and the membership of the seed row are audited finite facts about the frozen source file, not Lean theorems.
- *Mechanism.* No mechanism is claimed for the parity splits of Theorem 4; in particular, they are not attributed to untouchable cells.
- *Open converse.* Whether untouchable availability differences together with equal legal moves imply future equivalence is not established in this paper (Section 5.4; see the note added there).
- *No physical claim.* “History” means the combinatorial insertion word; nothing is claimed about physical memory, erasure, or irreversibility.

Section 2 recalls the setting. Sections 3–6 prove Theorems 1–4. Section 7 discusses related work, and Section 8 the scope and provenance. Appendix A lists the formal names and the claim firewall.

2. Setting

We recall the objects of [1] that are used here; apart from two Lean proofs noted below (the inclusion of reachable states in the catalogue and K2), nothing in this section is new.

2.1. Substrate and grammar

A D1b cell is a rhombohedron given by a base address $n \in \mathbb{Z}^6$ and a set T of three of the six directions; its six faces are $n + \sigma e_y + [0,1]^{T \setminus \{y\}}$ with $y \in T$ and $\sigma \in \{0,1\}$. The frozen catalogue K is a fixed set $K.\mathrm{cells}$ of 691 accepted cells of an icosahedral cut-and-project construction of AKN type [9], stored in the file `t7_ref2_tiles.json`; its cells are referred to by row numbers. The six basis vectors of $\mathbb{Z}^6$ are projected onto the six fivefold directions of an icosahedron [1], Section 2.1; the rotation group of the icosahedron, isomorphic to A_5, permutes the six direction lines and hence the face orientations and the cell types. No result of this paper uses this action, and neither the catalogue nor any reachable state is claimed to be invariant under it. A state is a finite set s of cells, and $\mathrm{inc}_s(f)$ is the number of cells of s having f as a face. The move alphabet is the set M_K of slots $i = (c, f)$ with $c \in K.\mathrm{cells}$ and f a face of c. A slot (c, f) is *legal* at s if $c \notin s$, $\mathrm{inc}_s(f) = 1$, and no other face of c has positive incidence; the update is $U(s, i) = s \cup \{c\}$. We write $s \cdot w$ for the state reached by a legal word w. All reachable states are taken from the seed $\{165\}$; they are subsets of $K.\mathrm{cells}$ (this inclusion is a Lean theorem, given that row 165 is a catalogue cell).

2.2. Extension languages and the future quotient

The extension language $\mathscr{L}(s)$ is the set of words legal from s, and $s \sim_F t$ if $\mathscr{L}(s) = \mathscr{L}(t)$. This is a Myhill–Nerode-type equivalence on states and, for this deterministic partial system, coincides with strong bisimulation. Future equivalence is preserved by common legal continuations: if $s \sim_F t$ and $w \in \mathscr{L}(s)$, then $s \cdot w \sim_F t \cdot w$. The set $\mathrm{LegalMoveSet}(s) \subseteq M_K$ of legal slots is the one-letter part of $\mathscr{L}(s)$, so future-equivalent states have the same legal moves.

The availability set is defined only for catalogue cells:

$$\mathrm{Avail}(s) = \{c \in K.\mathrm{cells} : \ c \notin s, \ \text{no two faces of } c \text{ are occupied in } s\},$$

where a face is occupied if it has positive incidence. Availability is a cell-level object; it differs from the slot-level set $\mathrm{LegalMoveSet}(s)$ both in type and in criteria [1], Section 7.

2.3. The codes

For a slot $i = (c, f)$, the local signature $\rho_s(i) = (P, I, O)$ records whether c is present (P), the incidence class 0, 1, or ≥ 2 of f (I), and whether another face of c has positive incidence (O); the slot is legal exactly when $\rho_s(i) = (0, 1, 0)$. The residue $R_{\mathrm{cat}}(s)$ is the table of all 4,146 local signatures. A signature is *dead* if $P = 1$, $I = \geq 2$, or $O = 1$; deadness is permanent. The normalisation ν replaces every dead signature by a fixed value $\rho_\dagger$, and $R_K^{(2)}(s)(i) = \nu(R_{\mathrm{cat}}(s)(i))$.

Three results of [1] are used:

- [1], Main Theorem A: $R_{\mathrm{cat}}(s) = R_{\mathrm{cat}}(t)$ implies $s \sim_F t$;

- [1], Main Theorem G(i): $R_K^{(2)}(s) = R_K^{(2)}(t)$ implies $s \sim_F t$;
- [1], Main Theorem G(ii), under an audited local contact certificate: two distinct reachable states at depth six, $b \cup \{163\}$ and $b \cup \{206\}$ with $b = \{92, 165, 345, 394, 507, 551\}$, have equal $R_K^{(2)}$ but different R_{cat}; this pair serves as the control pair of Section 6.

All three are Lean theorems (the last with the local contact certificate as hypothesis). By the last, R_{cat} is not Nerode-minimal on reachable states; whether $R_K^{(2)}$ is was left open.

2.4. Catalogue properties

Two incidence properties are used.

- (K1) Every face is a face of at most two catalogue cells.
- (K2) Two distinct cells share at most one face.

K2 holds for every pair of D1b cells, independently of any catalogue, and is proved in Lean. K1 is a property of the frozen catalogue. It enters the Lean statements through the *frozen catalogue certificate*, which consists of a face-owner bound, from which K1 is derived in Lean, and the fact $165 \in K.\mathrm{cells}$; both are finite facts about the frozen source file, audited independently and not proved in Lean.

2.5. Evidence tags

Each result carries a tag.

- M : a Lean theorem for an arbitrary catalogue object satisfying the hypotheses stated in the theorem, such as K1. All Lean modules of this paper were built in the author's environment (Lean `v4.34.0-rc1`), and the axiom audit of their main theorems reports only `propext`, `Classical.choice`, and `Quot.sound` (Section 8.1).
- M* : a statement about the frozen catalogue obtained by instantiating a Lean theorem whose catalogue hypothesis is the frozen catalogue certificate. The certificate, and the identification of the frozen source file with a catalogue object of the Lean statements (the 691 rows are not encoded in Lean), are audited, not proved in Lean.
- $\mathrm{M}_{\mathrm{paper}}$: a short mathematical argument that combines formal results and is given in the text but not itself formalised.
- E : a result of a frozen computation or audit, run under the preregistration of Section 8 with frozen scripts and hashed outputs.
- O : an open problem.

3. What $R_K^{(2)}$ identifies

This section characterises the identifications made by the normalised code $R_K^{(2)}$ of [1], Section 6. A characterisation was stated there as a paper proposition [1], Proposition 3; here it is proved in Lean, in a form that reads as a statement about the one-step legal interface. Throughout, $s, t \subseteq K.\mathrm{cells}$ are catalogue states, M_K is the set of move slots $i = (c, f)$ with $c \in K.\mathrm{cells}$ and f a face of c, and

$$\mathrm{LegalMoveSet}(s) = \{\, i \in M_K : \ i \text{ is legal at } s \,\}.$$

A cell c is *fresh* at s if $c \notin s$ and every face of c has incidence 0 in s.

3.1. Slot values of the normalised code

Lemma 1 (Slot trichotomy). *For every catalogue object K, every state s, and every slot $i = (c, f) \in M_K$, the value $R_K^{(2)}(s)(i)$ is exactly one of $(0,1,0)$, $(0,0,0)$, $\rho_\dagger$, and*

$$R_K^{(2)}(s)(i) = (0,1,0) \iff i \text{ is legal at } s, \qquad R_K^{(2)}(s)(i) = (0,0,0) \iff c \text{ is fresh at } s.$$

Consequently $R_K^{(2)}(s) = R_K^{(2)}(t)$ if and only if every slot is legal at s exactly when it is legal at t, and every catalogue cell is fresh at s exactly when it is fresh at t.

The lemma uses no property of the catalogue. It records formally what the paper proof of [1], Proposition 3 uses: the normalisation ν sends a present cell and an overconnected cell to the same value $\rho_\dagger$ on every slot, so $R_K^{(2)}$ does not distinguish "present" from "permanently blocked".

3.2. Availability from the normalised code

Passing from freshness to availability uses the property, checked for the frozen catalogue by an independent audit, that every face is a face of at most two catalogue cells (K1).

Lemma 2 (Availability recovery). *Assume K1. For a catalogue state s and a catalogue cell c,*

$$c \in \mathrm{Avail}(s) \iff c \text{ is fresh at } s \text{ or } c \text{ has a slot in } \mathrm{LegalMoveSet}(s).$$

Equivalently, $c \in \mathrm{Avail}(s)$ if and only if some slot of c has a value of $R_K^{(2)}(s)$ different from $\rho_\dagger$. The direction "$\Leftarrow$" holds without K1.

Without K1, the recovery can fail in principle: an absent cell with exactly one occupied face of incidence at least two is available, yet all its slots are dead.

3.3. Theorem 1

Write the *one-step legal interface* of a catalogue state as

$$\mathscr{I}(s) = \big(\mathrm{Avail}(s),\ \mathrm{LegalMoveSet}(s)\big).$$

Theorem 1 (The normalised code exactly determines the one-step legal interface). *Let $s, t \subseteq K$.* cells.

a. *$R_{\mathrm{cat}}(s) = R_{\mathrm{cat}}(t)$ implies $s = t$.*

b. *If K1 holds, then*

$$R_K^{(2)}(s) = R_K^{(2)}(t) \iff \mathscr{I}(s) = \mathscr{I}(t).$$

Proof. (a) A slot of a cell c has a presence bit 1 exactly when c is present, and every catalogue cell has a slot. (b) By Lemma 1, equality of $R_K^{(2)}$ is agreement of legality on all slots and of freshness on all catalogue cells. Agreement of legality on all slots is $\mathrm{LegalMoveSet}(s) = \mathrm{LegalMoveSet}(t)$, since every legal move is a slot of M_K. A fresh cell has no legal slot, so, by Lemma 2, a catalogue cell is fresh exactly when it is available and has no legal slot. Hence, given equal legal-move sets, agreement of freshness at a catalogue cell is equivalent to agreement of availability there. □

Status

Lemmas 1 and 2 and Theorem 1 are proved in Lean for an arbitrary catalogue object satisfying the stated hypotheses, K1 being a hypothesis in Lemma 2 and Theorem 1(b) [M]. The availability and legality predicates are those of the archive of [1]; the bridge to the abstract attachment system used in the proofs is itself proved exact. When instantiated on the frozen catalogue, K1 is supplied by a face-owner certificate, a finite fact about the frozen source file that is audited independently and not proved in Lean; the instantiated statements are therefore of status [M*]. Part (a) needs no hypothesis, although [1], Proposition 3 stated it under K1. Formal names are listed in Appendix A.

Theorem 1 says that $R_K^{(2)}$ is an exact, complete invariant of the one-step legal interface: it induces exactly the same partition of catalogue states as $\mathscr{I}$, so two catalogue states have the same normalised code exactly when they offer the same available cells and the same legal moves. It does not say that equal interfaces imply equal futures by themselves; that implication holds because equal $R_K^{(2)}$ implies equal futures [1], Main Theorem G(i).

3.4. All identifications within depth eight

Corollary 1.1 (Identifications within depth eight)**.** *Among the* $720{,}473$ *states reachable from* $\{165\}$ *by at most eight legal insertions, the classes of* $R_K^{(2)}$ *with at least two states are exactly* 562 *classes:* 527 *of size two,* 31 *of size three,* 3 *of size four, and one of size five.*

Status

The corollary combines a Lean theorem with an exhaustive computation ([M*] + [E]). The Lean part is the forward direction of Theorem 1(b) for states reachable from $\{165\}$: under the frozen catalogue certificate (K1 and $165 \in K.\,\mathrm{cells}$), every $R_K^{(2)}$-collision has the same interface $\mathscr{I}$. The computational part enumerates all states within depth eight and groups them by $\mathscr{I}$; a digest of $\mathscr{I}$ is used only as an index, and every candidate group is confirmed by exact comparison. Since states with the same interface have the same digest, no collision is missed. The per-class recomputation of $R_K^{(2)}$ over all $4{,}146$ slots, reported in [1], Section 6 and repeated here under a frozen script, confirms that $R_K^{(2)}$ is constant on each of the 562 classes and that R_{cat} is pairwise distinct on each. The remaining assumptions are the audited catalogue certificate and the agreement of the computational model with the formal definitions of availability and legality, which is supported by the independent reimplementation of stage S0 (all states within depth five) and by the audits of Section 8, but is not a Lean theorem.

The 562 classes form the state set S-A1 of Section 6. All their states have legal moves; the shallowest identifications occur at depth six, and the witness pair of [1], Main Theorem G is one of them. In 397 classes, all states have the same number of present cells; in the other 165, the number varies by one.

3.5. What is not claimed

- Theorem 1 does not say that $R_K^{(2)}$ is Nerode-minimal; Section 4 shows that it is not.
- Corollary 1.1 concerns depth at most eight only.
- K1 is not proved in Lean for the frozen catalogue; it enters as an audited certificate.

4. $R_K^{(2)}$ is not Nerode-minimal

Theorem 1 identifies the partition induced by $R_K^{(2)}$ on catalogue states: it is the partition by the one-step legal interface $\mathscr{I}$. The future quotient $\sim_F$ is the partition induced by equality of complete extension languages, and by [1], Main Theorem G(i) it is at least as coarse as the partition by $R_K^{(2)}$. This section shows that it is strictly coarser on reachable states, already among states that still have legal moves. Together with Theorem 1, this moves the question from codes to futures: the remaining gap between $R_K^{(2)}$ and $\sim_F$ consists of interface differences not distinguished by the complete extension language.

4.1. Where the gap can lie

Proposition 1 (Shape of the gap). *Assume K1. Let $s, t \subseteq K.\mathrm{cells}$ with $s \sim_F t$ and $R_K^{(2)}(s) \neq R_K^{(2)}(t)$. Then $\mathrm{LegalMoveSet}(s) = \mathrm{LegalMoveSet}(t)$ and $\mathrm{Avail}(s) \neq \mathrm{Avail}(t)$.*

Proof. Every one-letter word is legal from s exactly when it is legal from t, so the legal-move sets agree. If the available catalogue cells also agreed, then $\mathscr{I}(s) = \mathscr{I}(t)$, and Theorem 1(b) would give $R_K^{(2)}(s) = R_K^{(2)}(t)$. □

This was anticipated in [1], Section 6 as a consequence of the then unformalised [1], Proposition 3; with Theorem 1 it now rests on a Lean theorem and the one-line argument above (status $[\mathrm{M}_{\mathrm{paper}}]$). Every pair in the gap is therefore an *availability gap*: same legal moves, different available cells. Section 5 gives a necessary condition on the availability differences that can occur in it.

4.2. Terminal states

The simplest source of such pairs is terminality. Two terminal states have the same extension language $\{\varepsilon\}$, while $R_K^{(2)}$ can still record which catalogue cells remain available.

Proposition 2 (Terminal witness). *The terminal states reached from $\{165\}$ by the seeded random legal runs with seeds 0 and 1 (of 356 and 350 moves) are future equivalent, and their residues $R_K^{(2)}$ differ in 198 of the $4{,}146$ slots. They have 17 and 22 available catalogue cells, respectively.*

Status

The witness was observed in an exploratory run before the preregistration and is reproduced here by a frozen procedure [E]: both words were replayed move by move from $\{165\}$ with legality checked at each step; terminality (no legal slot among the $4{,}146$) was confirmed by the grammar core and by an independent reimplementation; and the residues were compared over all slots, the first differing slot being recorded. It is a reproduction of an observed witness, not a preregistered discovery. A Lean certificate is possible in principle but has not been produced.

Terminal collapse is a weak form of the gap: the future has become trivial, and the differences recorded by $R_K^{(2)}$ concern cells that can no longer be inserted at all. The next result shows that the gap is not confined to terminal states.

4.3. Theorem 2

Theorem 2 (Nonterminal non-minimality). *There are states s, t reachable from $\{165\}$, both admitting a legal move, such that $s \sim_F t$ and $R_K^{(2)}(s) \neq R_K^{(2)}(t)$. Within the examined set S-DAG, such pairs occur in $280{,}816$ future-equivalence classes and number*

5,384,241.

The examined set and the certificate.

Anchors were taken 4, 8, 12, and 16 steps before the end of 64 seeded random legal runs from $\{165\}$ (256 anchors). The set S-DAG consists of the complete descendant DAGs of the 245 anchors whose DAGs could be explored completely (see below). Since legal insertions only add cells, each descendant DAG is finite and acyclic. On a complete DAG the canonical future

$$F(x) = \{\, (i,\ F(U(x,i))) :\ i \in \mathrm{LegalMoveSet}(x) \,\}$$

is computed from the terminal nodes upwards, and two states are identified exactly when their canonical futures are equal (C-DAG). Equality of canonical futures is equality of extension languages. Identity of canonical futures is decided by exact comparison of their encodings, never by a hash value. Of the 256 anchors, 245 were explored completely; the 11 anchors whose DAGs exceeded 200,000 nodes were discarded and issued no certificate. The complete DAGs contain 3,237,517 distinct states and 1,285,298 distinct futures; 695,474 classes contain at least two states. In each class, nonterminal pairs with different $R_K^{(2)}$ (*A2′ pairs*) and with equal $R_K^{(2)}$ (*A1-deep pairs*, 6,574,730 in all) were counted separately; a class containing at least one A2′ pair (respectively A1-deep pair) is called an A2′ class (respectively A1-deep class), and these classes form the strata A2′ and A1-deep. A class contains up to 52 distinct values of $R_K^{(2)}$.

Status

Theorem 2 is a result of the preregistered computation [E]. The preregistration recorded the existence of A2′ pairs as expected but uncertain, and fixed in advance that a negative outcome would have been reported as “not found in the examined set”. The future equivalence of each pair rests on the C-DAG certificate. The scripts were frozen before the production runs, the outputs were frozen with SHA-256 hashes, and all are listed in the provenance ledger (Section 8).

Consistency with Proposition 1

An independent audit of the frozen classes (Section 5.3) found, over all 280,816 classes with A2′ pairs, no class with unequal legal-move sets among its nonterminal members and no A2′ pair with equal available sets, as Proposition 1 requires.

An example

The A2′ witness of Section 6.3 is one such pair: two states with 356 and 357 present cells, recorded in the descendant DAG of the anchor four steps before the end of run 0, with the same legal moves, available sets differing exactly in row 622, and residues differing in 6 slots.

Corollary 2.1. *On states reachable from $\{165\}$, the partition induced by $R_K^{(2)}$ is strictly finer than the future quotient, even when restricted to states that admit a legal move.*

Status

The partition by $R_K^{(2)}$ refines the future quotient by [1], Main Theorem G(i) [M]; the refinement is strict by the witness pairs of Theorem 2 [E]. The corollary is an existence statement and makes no claim about frequency.

4.4. What is not claimed

- No statement is made about the frequency of the gap outside the examined set; S-DAG consists of states close to the end of runs, and the counts are not estimates for the reachable state space.
- Whether the gap occurs among states within depth eight (the set S-A1 of Section 3.4) was not examined; S-A1 consists of $R_K^{(2)}$-classes and certifies future equivalence only through equality of $R_K^{(2)}$.
- Theorem 2 is not formalised in Lean; its certificate is computational.

5. Touchable differences separate futures

By Proposition 1, every future-equivalent pair that $R_K^{(2)}$ separates differs in its available catalogue cells. This section gives a necessary condition on the availability differences that can survive future equivalence between reachable states: they lie in cells that the available side can never touch. The mechanism is constructive. If a difference cell can be touched along some legal continuation, a distinguishing word can be written down explicitly.

5.1. Touching

A cell c is *touched* at a state x if some face of c has positive incidence in x. It is *permanently untouchable from* s if it is touched at no state $s \cdot w$ with w legal from s (including $w = \varepsilon$). An available cell c that is touched at s has exactly one occupied face g under K1, and the slot (c, g) is legal at s. In particular, an available cell that is not fresh carries a legal slot (Lemma 2).

Besides K1 the arguments use the property that two distinct catalogue cells share at most one face (K2). K2 holds for every pair of D1b cells, independently of the catalogue, and is proved in Lean.

5.2. Theorem 3

Theorem 3 (Constructive separation). *Assume the frozen catalogue certificate. Let s be reachable from $\{165\}$, let t be any state, and let $c \in K.\mathrm{cells}$ with $c \in \mathrm{Avail}(s)$ and $c \notin \mathrm{Avail}(t)$. If w is legal from s and some face of c has positive incidence in $s \cdot w$, then there are a prefix u of w and a move slot j inserting c such that $u\,j$ is legal from s and not legal from t.*

Idea of the proof. If c is already touched at s, take u empty: as noted in Section 5.1, the available touched cell c carries a legal slot (c, g). Otherwise let u be the shortest prefix of w after which some face g of c has positive incidence, and let u' be u without its last move. Before the first touch, no move inserts c or changes the incidence of any face of c; hence c remains fresh along u', and therefore available. The last move of u inserts a cell that shares at most one face with c (K2), so after u exactly one face g of c is occupied, with incidence exactly one, since all faces of c had incidence zero before the last move. Hence (c, g) is legal at $s \cdot u$. From t, either u is already illegal, or c is still unavailable after u, because availability is never regained. In both cases $u\,(c, g)$ is not legal from t. □

Status

Theorem 3 is proved in Lean ([M*] for the frozen catalogue), first for an abstract single-face attachment system with K1 and K2 as hypotheses and then, through a definition bridge proved exact, for the D1b semantics of [1]. The only catalogue hypothesis exposed is the frozen catalogue certificate (K1 and $165 \in K.\mathrm{cells}$), a finite fact about the frozen source file that is

audited independently and not proved in Lean. The state t is arbitrary; reachability of s is used only to place s inside the catalogue. In its catalogue-contained form ($s \subseteq K.\mathrm{cells}$, with the face-owner bound as the only catalogue hypothesis), the theorem is a Lean theorem for an arbitrary catalogue object, as for Theorem 1 [M]; the form stated here is its instantiation on the frozen catalogue [M*].

Corollary 3.1 (Touch corollary). *Assume the frozen catalogue certificate. Let s be reachable from $\{165\}$, $s \sim_F t$, and $c \in K.\mathrm{cells}$ with $c \in \mathrm{Avail}(s)$ and $c \notin \mathrm{Avail}(t)$. Then c is permanently untouchable from s. In particular c is fresh at s.*

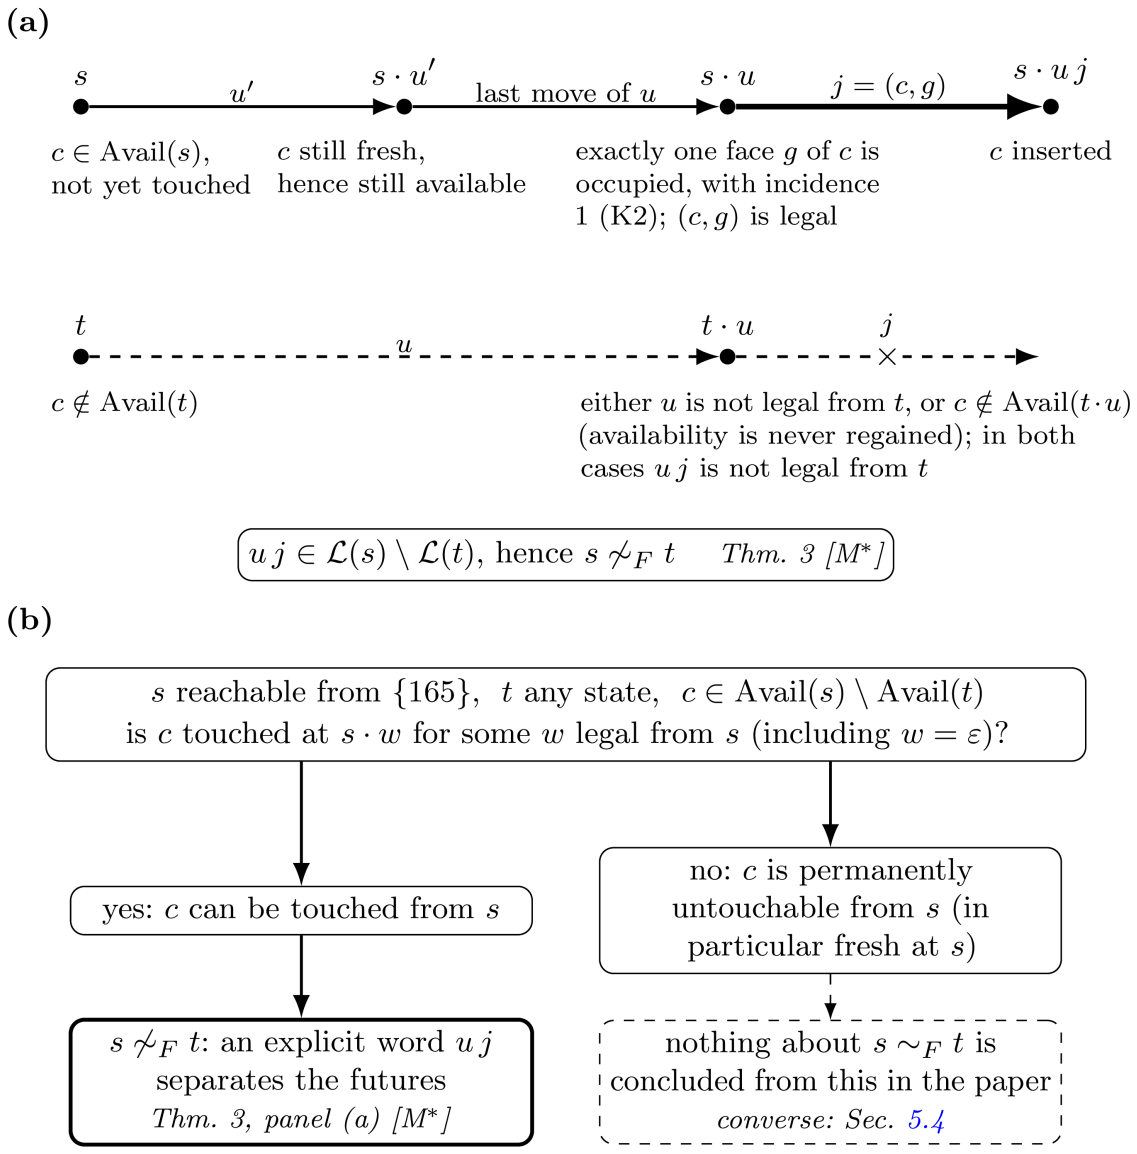
Proof. If some legal w from s touched c, Theorem 3 would give a word $u\,j$ legal from s and not from t, contradicting $s \sim_F t$. The case $w = \varepsilon$ gives freshness. □

Corollary 3.1 is also proved in Lean, in both forms. Combined with Proposition 1, it localises the gap between $R_K^{(2)}$ and $\sim_F$.

Corollary 3.2 (Localisation of the gap). *Assume the frozen catalogue certificate. Let s, t be reachable from $\{165\}$ with $s \sim_F t$ and $R_K^{(2)}(s) \neq R_K^{(2)}(t)$. Then $\mathrm{LegalMoveSet}(s) = \mathrm{LegalMoveSet}(t)$, the symmetric difference of the available catalogue cells is nonempty, and every cell of $\mathrm{Avail}(s) \setminus \mathrm{Avail}(t)$ is fresh at s and permanently untouchable from s, and symmetrically for t.*

Corollary 3.2 combines Lean theorems (Theorem 1, Corollary 3.1, and the inclusion of reachable states in the catalogue) by the short arguments given above; the combination itself is not separately formalised (status [$\mathrm{M}_{\mathrm{paper}}$] from [M*]).

Thus the available catalogue cells of two future-equivalent reachable states can differ only at cells that are untouched now and remain untouched along every legal continuation of the side on which they are available. The cells by which the two states differ as sets need not be of this kind: in the A2′ witness of Section 6.3, the rows 247, 287, and 622 are present in one state and absent from the other, and only row 622 is an availability difference. Theorem 3 answers affirmatively the question whether an availability difference in a cell that can still be touched from the available side eventually shows up in legality: such a difference cannot exist between future-equivalent states. Figure 2 illustrates the construction and the resulting case split.



Hence, if $s \sim_F t$, only the right branch can occur (*Cor. 3.1 [M*]*). Audit on S-DAG: in all 5,384,241 pairs of Theorem 2 (19,209,838 difference cells, counted in both directions), every availability-difference cell lies in the right branch, for the side on which it is available (*Sec. 5.3 [E]*).

Figure 2. Touchable and untouchable availability differences. (a) The construction in the proof of Theorem 3, drawn schematically (no catalogue geometry is depicted). Along a legal continuation w of s that touches c, let u be the shortest prefix after which a face of c is occupied and u' the prefix without its last move; if c is already touched at s, then u is empty. The slot $j = (c, g)$ is legal after u from s but not from t. (b) The case split behind Corollary 3.1. The left branch separates the futures. The right branch is only a necessary condition for future equivalence: this paper does not claim that a difference in an untouchable cell is invisible to the future. The converse, for catalogue states with equal legal moves in which every availability difference is untouchable from its available side, was proved later for the abstract system (note added in Section 5.4) and is not used here.

Remark (where the residues differ)

Under the hypotheses of Corollary 3.2, every slot at which $R_K^{(2)}(s)$ and $R_K^{(2)}(t)$ differ is a slot of a cell in the symmetric difference of $\mathrm{Avail}(s)$ and $\mathrm{Avail}(t)$; its value is $(0,0,0)$ on the side where the cell is available and $\rho_\dagger$ on the other. Indeed, by Lemma 1 and the equality of the legal-move sets, a differing slot (c,f) has value $(0,0,0)$ on one side, say at s, so that c is fresh at s, and value $\rho_\dagger$ at t. If c were available at t, Lemma 2 would make c fresh at t, contradicting the value $\rho_\dagger$, or give c a legal slot at t, hence at s, which is impossible for a fresh cell. By Corollary 3.1, c is therefore permanently untouchable from s. (Status [$\mathrm{M}_{\mathrm{paper}}$] from [M] and [M*].)

5.3. Audits on the examined states

Two preregistered audits test the implementation of these statements on the frozen data. They do not prove the theorem, which is a Lean theorem, but check that the computational model and the formal definitions agree where both apply.

Exhaustive audit on S-DAG

For every state x of a complete descendant DAG, the set $T(x)$ of cells touched at some legal continuation of x was computed exactly: $T(x)$ is the set of cells touched at x if x is terminal, and the union of $T(y)$ over the children y of x otherwise (incidence only grows, so this is the union over terminal descendants). For every A2′ pair (s,t) of Theorem 2 it was checked that

$$\big(\mathrm{Avail}(s)\setminus\mathrm{Avail}(t)\big)\cap T(s)=\varnothing \quad\text{and}\quad \big(\mathrm{Avail}(t)\setminus\mathrm{Avail}(s)\big)\cap T(t)=\varnothing.$$

Over all 5,384,241 A2′ pairs, with 19,209,838 difference cells counted in both directions, there was no violation [E]. In words: each availability-difference cell in a future-equivalent A2′ pair is untouchable from the available side. This is a statement about the availability differences of these pairs; the audit did not separately examine the cells by which the two states of a pair differ as sets.

Finite-horizon audit near the seed

Among the 720,473 states within depth eight, states were grouped by legal-move set, and the first 2,000 pairs with equal legal-move sets and different available sets were taken in a preregistered order. These pairs carry no certificate of future equivalence; the audit tests the constructive form of Theorem 3. For each of the 4,022 difference cells, the legal words of length at most three from the available side were searched for a word touching the cell. For the 3,413 cells touched within this horizon, the word $u\,j$ of Theorem 3 was built and verified to be legal from the available side and not legal from the other side, with no failure [E]. The remaining 609 cells were not touched within the horizon; by the preregistered rule they count neither for nor against the theorem. This audit is not fully blind: a development run had previewed one pair of the kind examined, and this is recorded as a protocol deviation. Candidate pairs were indexed by 128-bit digests under the assumption that no digest collision occurs among the explored states; every selected pair was re-checked exactly.

A descriptive probe

For the 218 pairs in which no difference cell was touched within the horizon, the sets of legal words of length at most two agreed in all cases. Agreement within a finite horizon is not future equivalence, and no conclusion is drawn from it.

5.4. The converse

The converse of Corollary 3.2 is open: if $\mathrm{LegalMoveSet}(s) = \mathrm{LegalMoveSet}(t)$ and every availability difference lies in a cell that is permanently untouchable from its available side, is $s \sim_F t$? This does not follow by the same argument. A difference cell c present in t occupies faces that it shares with neighbouring cells, and such a neighbour may become legal in the future of t without becoming legal in the future of s. A correct converse must therefore control the influence of a difference cell on its neighbours, and the right form of that condition is not known. Equivalently, it is open whether the coarsening of $R_K^{(2)}$ obtained by also merging the slots of permanently untouchable cells is future sufficient and Nerode-minimal on reachable states. Such a coarsening is defined through reachability of the future and is not a locally computable code.

Note added

After the work reported here was completed, the converse was proved in Lean in subsequent work (in preparation), for the catalogue-relative abstract attachment system of Appendix A.1 and under K1 alone: if s and t are catalogue states with $\mathrm{LegalMoveSet}(s) = \mathrm{LegalMoveSet}(t)$, and every cell that is available in one of them and not in the other is permanently untouchable from the state in which it is available, then $s \sim_F t$ (`waterwayN2ConverseV1`; built in the author's environment, with an axiom audit reporting only `propext`, `Classical.choice`, and `Quot.sound`). That module is not part of the formal code of this paper, its instantiation in the D1b names of the frozen catalogue is not given here, and the question of Nerode minimality raised in the preceding paragraph is not addressed by it. The note is not used anywhere else in this paper.

5.5. What is not claimed

- The catalogue certificate (K1 and $165 \in K.\mathrm{cells}$) is not proved in Lean.
- The audits concern the examined states only.
- The cells by which two future-equivalent states differ as sets (present in one, absent from the other) are not claimed to be untouchable; the remark above concerns only the slots at which the residues $R_K^{(2)}$ differ.
- The converse in Section 5.4 is not proved in this paper (see the note added there).

6. History parity and the future quotient

6.1. The potential

We use the explicit potential of the companion paper [2] without reproving its properties.[1] For a point $b \in \mathbb{Z}^6$ let $p(b) = \sum_{i=1}^{6} b_i \bmod 2$. For a catalogue state $s \subseteq K.\mathrm{cells}$,

$$\Phi(s) = \sum_{c \in s} p(\mathrm{base}\, c) \; + \sum_{f:\ \mathrm{inc}_s(f) \geq 2} p(\mathrm{base}\, f) \pmod 2,$$

and for a move slot $i = (c, f)$ let $\sigma(i) = 0$ if the face f is based at the base point of c (base-aligned, LOW) and $\sigma(i) = 1$ otherwise (non-base-aligned, HIGH); σ is the relative-side move label of [2]. Two facts from [2] are used.

i. $\sigma(i)$ depends only on the catalogue and the slot, not on the state; $\Phi(s)$ is a function of the current state alone.

ii. For every word w legal from s,

$$\Phi(s \cdot w) = \Phi(s) + \sum_{i \in w} \sigma(i) \pmod 2. \tag{1}$$

Thus Φ is a history-derived but current-state evaluable parity. On states reachable from $\{165\}$, (1) expresses $\Phi(s)$ as $\Phi(\{165\})$ plus the number of non-base-aligned moves of any legal word reaching s, modulo 2; since Φ is a function of the current state, the parity of this count is independent of the legal word used to reach the state.

The question of this section is whether Φ factors through the future quotient $q_F : S \to S/\sim_F$ on the set S of states reachable from $\{165\}$, that is, whether $s \sim_F t$ implies $\Phi(s) = \Phi(t)$ on reachable states. Since Φ is a function of the current state, this is a comparison of two partitions of the same state set; no stored history is required to evaluate Φ.

6.2. Design of the test

The test was preregistered. The preregistration fixed the state sets, the certificates of future equivalence, the decision rules, the strata, and the rule for selecting the witnesses reported below. Except for the pre-existing control pair, no value of Φ on the frozen evaluation sets was inspected before preregistration and class freezing. The control pair is the witness pair of [1], Main Theorem G, for which $\Phi(s) = \Phi(t)$ was known beforehand; it belongs to stratum A1-P below and served as a gate (the run would have stopped had it not reproduced $\Delta\Phi = 0$).

Certificates of future equivalence

Two states are treated as future equivalent only under one of three certificates: equality of $R_K^{(2)}$ over all $4{,}146$ slots, which implies future equivalence by [1], Main Theorem G(i) (C-R2); exact equality of canonicalised futures computed on the complete finite descendant DAGs of both states (C-DAG); or terminality of both states (C-TERM). Agreement of finite-horizon signatures is not a certificate.

State sets and strata

- *S-A1*: all 562 $R_K^{(2)}$-classes with at least two states among the states reachable from $\{165\}$ within eight insertions (C-R2). By Corollary 1.1 these are exactly the $R_K^{(2)}$-collisions within depth eight. Stratum *A1-P*: the 397 classes whose states all have the same number of present cells; stratum *A1-S*: the remaining 165.
- *S-DAG*: the complete descendant DAGs of 245 anchors taken $4, 8, 12, 16$ steps before the end of 64 seeded random legal runs (C-DAG); 11 anchors exceeding $200{,}000$ nodes were discarded without issuing any certificate. Stratum *A2′*: the A2′ classes, which contain a nonterminal pair with different $R_K^{(2)}$; stratum *A1-deep*: the A1-deep classes, which contain a nonterminal pair with equal $R_K^{(2)}$ (Section 4.3); *TERM-DAG*: the class of terminal states found in these DAGs.
- *S-TERM*: the 64 distinct terminal states of the same runs (C-TERM).

Decision rules

A certified class *splits* if it contains two states with different Φ; in A2′ a split requires a pair with different $R_K^{(2)}$ and different Φ, in A1-deep a pair with equal $R_K^{(2)}$ and different Φ. The first gate (A0) asks whether any certified class in any set splits. A split in A1-P or in A2′ is called a *structural* separation. A split occurring only among terminal states is attributed to the collapse of the future to $\{\varepsilon\}$ and is not called structural. Absence of splits would have been reported only as the absence of counterexamples in the examined classes.

Execution

The class lists were frozen, with SHA-256 hashes, before Φ was computed for any of their states other than the control pair. The evaluation script was frozen after development runs restricted to the control pair and to states within depth five, and was executed once. For every examined state, Φ was computed from the current state and compared with the path value given by (1) along the recorded legal word; the run stops on any disagreement. All hashes are listed in the provenance ledger (Section 8).

6.3. Result

Theorem 4 (History-parity separation from the future quotient). *On the examined reachable states, Φ does not factor through $\sim_F$. More precisely:*

a. *In stratum A1-P there are states s, t reachable from $\{165\}$ with*

$$R_K^{(2)}(s) = R_K^{(2)}(t), \qquad |s| = |t|, \qquad \Phi(s) \neq \Phi(t);$$

in particular $s \sim_F t$, $\mathrm{LegalMoveSet}(s) = \mathrm{LegalMoveSet}(t)$, and $\mathrm{Avail}(s) = \mathrm{Avail}(t)$.

b. *In stratum A2′ there are nonterminal states $s \sim_F t$ with $R_K^{(2)}(s) \neq R_K^{(2)}(t)$ and $\Phi(s) \neq \Phi(t)$.*

Status

Theorem 4 is a result of the preregistered computation [E]. Within it, the future equivalence in (a) rests on [1], Main Theorem G(i) and the equality of the one-step interface in (a) on Theorem 1 ([M*] for the frozen catalogue); the future equivalence in (b) rests on the C-DAG certificate.

The A1-P witness

The witness selected by the preregistered rule (least depth, then lexicographic order) is the pair

$$s = \{0, 46, 165, 206, 255, 299, 436, 507\}, \quad t = \{0, 117, 165, 206, 230, 255, 299, 436\},$$

reached by the legal words

$$\begin{aligned} s &: \ 165 \to 0 \to 206 \to 299 \to 507 \to 436 \to 46 \to 255, \\ t &: \ 165 \to 0 \to 206 \to 299 \to 117 \to 230 \to 255 \to 436, \end{aligned}$$

where each arrow inserts the named row through its unique occupied face. Both states lie at depth seven, have eight present cells, and have the same normalised residue $R_K^{(2)}$ over all $4{,}146$ slots; hence they have the same available catalogue cells, the

same legal moves, and the same extension language. Nevertheless $\Phi(s) = 0$ and $\Phi(t) = 1$. The two states share the first three insertions after the seed and the rows 255 and 436; they differ in the pair $\{46, 507\}$ versus $\{117, 230\}$.

The A2′ witness

The witness was selected by the preregistered rule, in the operational form fixed in the frozen evaluation script before Φ was computed: first by the earliest point, in the ordered cumulative registry of anchors by (D, seed), by which both members had appeared; then by lexicographic order. Both members were first recorded in the descendant DAG of the anchor four steps before the end of run 0. Its two states have 356 and 357 present cells; one contains row 247, the other rows 287 and 622. They have the same legal moves and both admit a legal move; their available sets differ exactly in row 622, their residues $R_K^{(2)}$ differ in 6 slots, and $\Phi = 1$ and $\Phi = 0$, respectively. The full words are given in the machine-readable witness record of the supplementary data.

Terminal states

The terminal states of runs 0 and 1, which are the witness pair of Section 4, have $\Phi = 0$ and $\Phi = 1$. As stated in the decision rules, terminal splits are not counted as structural.

Extent within the examined sets

Table 1 reports every quantity fixed by the preregistration: for each stratum, the number of certified classes, the number that split, and the number of split pairs. The first gate A0 is a global yes/no gate over all sets, and it is positive; the A0 row of the table reports its aggregate over S-DAG. Splits occur in every stratum. The companion statistic $\Psi(s) = |s| \bmod 2$ is constant on each A1-P class by the definition of the stratum, so the 80 A1-P splits are not explained by the parity of the number of present cells. For the other strata, the last column gives the number of classes on which Ψ is not constant; the overlap of Φ-splits and Ψ-splits was not part of the preregistered analysis and is not reported here.

Stratum	classes	split	fraction	split pairs	Ψ varies
A1-P (primary)	397	80	0.20	88	0
A1-S (secondary)	165	43	0.26	54	165
A2′	280,816	197,002	0.70	2,806,374	215,057
A1-deep (secondary)	563,260	228,033	0.40	2,658,333	200,976
TERM (64)	1	1	1.00	1,020	1
TERM-DAG (2,164)	1	1	1.00	1,169,995	1
A0 (S-DAG, all)	695,474	357,473	0.51	6,634,702	295,230

Table 1. Splitting of certified future-equivalence classes by Φ. Counts refer to the examined sets only and are not estimates for the reachable state space. TERM and TERM-DAG: number of terminal states in parentheses. The row A0 covers all classes of S-DAG; the classes of S-A1 and S-TERM are those of the rows above it.

Checks

The control pair reproduced $\Delta\Phi = 0$. The current-state value of Φ agreed with the path value (1) on every examined state (1,164 states of S-A1, 64 of S-TERM, 2,647,693 class members of S-DAG); the literal face-incidence evaluation of Φ was additionally compared with an equivalent pair-sum evaluation on all states of S-A1 and S-TERM and on a fixed sample of 2,656 states of S-DAG. The A2′ witness was rebuilt from its words, and its residues were recomputed in full.

6.4. Persistence along common futures

Corollary 4.1 (Persistence of a parity difference). *Let $s, t \subseteq K$.* cells *with $s \sim_F t$, and let w be a word legal from s. Then w is legal from t, $s \cdot w \sim_F t \cdot w$, and*

$$\Phi(s \cdot w) - \Phi(t \cdot w) = \Phi(s) - \Phi(t) \pmod 2.$$

In particular, if $\Phi(s) \neq \Phi(t)$, then $\Phi(s \cdot w) \neq \Phi(t \cdot w)$ for every common legal continuation w.

Proof. Since $\mathscr{L}(s) = \mathscr{L}(t)$ and $w \in \mathscr{L}(s)$, also $w \in \mathscr{L}(t)$. For every word v, $v \in \mathscr{L}(s \cdot w)$ iff $wv \in \mathscr{L}(s)$ iff $wv \in \mathscr{L}(t)$ iff $v \in \mathscr{L}(t \cdot w)$; hence $s \cdot w \sim_F t \cdot w$. Applying (1) to s and to t with the same word gives the identity. □

Status

Corollary 4.1 holds for all states of the catalogue transition system to which the integration law (1) applies, not only for the examined ones; it is a mathematical statement, not an empirical generalisation beyond the examined sets. It follows from the integration law (1) of [2], a Lean theorem of the archive on which the modules of this paper are built (built in the author's environment with them), by the elementary argument above, which is not separately formalised (status $[\mathrm{M}_{\mathrm{paper}}]$ from $[\mathrm{M}]$). It is not a finding of the computation.

Combining Theorem 4 with Corollary 4.1: two current states with the same complete legal future can carry a one-bit difference that is readable from the current state, that never appears in any legality judgement, and that is preserved along every common legal continuation. In the A1-P witness the two states even have the same normalised residue and the same number of present cells, so the bit is neither part of the one-step legal interface nor the parity of the state size. Figure 3 shows the witness and the persistence along a common continuation.

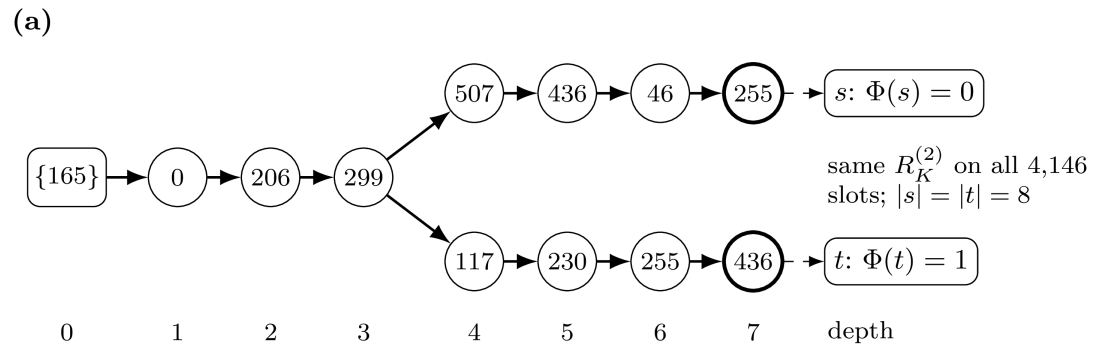


Figure 3. Persistent history parity. (a) The A1-P witness of Theorem 4(a), drawn from its two legal words (Section 6.3). Each circle is a catalogue row inserted through its unique occupied face; rows are zero-based catalogue indices. The two words share the first three insertions after the seed and end at depth seven in states with the same normalised code and the same number of present cells, yet Φ differs. Only the endpoint values of Φ are shown; no catalogue geometry is depicted. (b) Schematic of Corollary 4.1: along any common legal continuation the two states remain future equivalent and their Φ-difference is unchanged. The parity is never visible to legality; it is a function of the current state that is not a function of the future-equivalence class.

6.5. What is not claimed

- No statement is made about reachable states outside the examined sets; the fractions in Table 1 are not estimates for the reachable state space.

- No mechanism for the splits is claimed. In particular, it is not claimed that a Φ-difference is caused by a difference in permanently untouchable cells (Section 5); the A2′ witness exhibits both, but their relation was not analysed.
- Corollary 4.1 is a consequence of the integration law of [2], not an empirical finding.
- "History" refers to the combinatorial insertion word. No statement is made about physical memory, erasure, or irreversibility.

7. Related work and novelty boundary

The objects compared in this paper are classical; the results concern what these objects do in one explicit finite geometry. We list the neighbouring lines of work in order of proximity and state, for each, what is shared and what is not.

Myhill–Nerode equivalence and bisimulation

Future equivalence $\sim_F$ is the state version of right-language (Myhill–Nerode) equivalence [3][4], and for a deterministic partial transition system it coincides with trace equivalence and strong bisimulation [5][6]. Computing the coarsest stable partition of a finite system is the classical partition-refinement problem [10]. Nothing here is new as a construction. What the present paper adds is a description, in the terms of a specific geometry, of how the explicit codes R_{cat} and $R_K^{(2)}$ relate to this quotient on reachable states: under K1, $R_K^{(2)}$ is exactly the one-step interface (Theorem 1), and the quotient is strictly coarser (Theorem 2).

Follower sets and future covers

In symbolic dynamics, the follower set of a word w is the set of words u with wu in the language of the shift [11]; equality of follower sets is called future equivalence, and its classes are the states of the future cover [12]. The extension language $\mathscr{L}(s)$ is the analogue of a follower set indexed by geometric states rather than by words, and the notation $\sim_F$ follows this usage. No novelty is claimed for the equivalence itself.

Extender sets

For multidimensional subshifts, the extender set of a finite pattern is the set of configurations on the complement with which it combines to a point of the shift [13]. This is a static completion notion. It is closer to the forced-completion and empire questions for tilings than to the sequential extension language studied here.

Abstract interpretation and exact abstraction

The passage $R_{\mathrm{cat}} \to R_K^{(2)} \to \sim_F$ is a chain of successively coarser exact abstractions of the legality semantics [14]. Theorems 1 and 2 locate the first step precisely (dead slots) and show that the second step is proper. We do not identify future sufficiency with completeness in the sense of abstract interpretation.

Causal states and stored information

In computational mechanics, causal states partition pasts by their conditional futures [15], and crypticity measures stored information that is not revealed by future behaviour [16]. Theorem 4 with Corollary 4.1 gives a deterministic analogue of state information that is not revealed by future behaviour: a one-bit current-state parity that the complete legal future does not determine and that every common continuation preserves. The settings differ, since crypticity is an information measure for stochastic processes and Φ is a parity on a deterministic finite system. We do not identify the two.

Dead transitions in Petri nets

Deciding whether a transition can ever fire again from a reachable marking is a reachability question, distinct from structural sufficient conditions [17]. Permanently untouchable cells (Section 5) are defined by a reachability condition of the same kind, while dead slots of $R_K^{(2)}$ are detected locally. This is an analogy only: cells, slots, and state-dependent legality do not correspond one-to-one to Petri-net transitions.

AKN-type tilings, matching rules, empires, and growth

The substrate is an icosahedral cut-and-project patch of AKN type [9]. Matching rules and the empire of a patch concern the global tilings compatible with a patch [18][19], and quasicrystal growth rules concern attachment processes that build a tiling [20][21]. *The AKN connection concerns the icosahedral cut-and-project substrate only. The sequential one-cell attachment grammar studied here is an additional finite transition rule on a fixed patch and is not claimed to be a standard AKN matching rule or growth rule* (see also [1], Section 2.4).

Novelty boundary

The future quotient, partition refinement, follower sets, and the relative-side potential are established. The contribution of this paper consists of structural results in one explicit finite attachment geometry: an exact characterisation of the identifications made by $R_K^{(2)}$ (Theorem 1); certified reachable witnesses that the future quotient is strictly coarser, already among states with legal moves (Theorem 2); a constructive separation theorem that confines the availability differences absorbed by future equivalence between reachable states to cells permanently untouchable from the available side (Theorem 3); and an outcome-blind test showing that a history-derived, current-state evaluable parity does not factor through the future quotient (Theorem 4). The general statements are formally verified or follow from formal results by short arguments given in the text; the existence statements are certified by preregistered, audited computation.

8. Scope, provenance, and conclusion

8.1. Formal verification

The Lean modules of this paper extend the archive of [1]. They consist of an abstract single-face attachment system with the touch lemma, a definition bridge proved exact between that system and the D1b semantics, the catalogue-free proof of K2, the connection of the face-owner certificate to K1, the constructive separation theorem with its corollaries, and the four

modules of Theorem 1. All were built in the author's environment with Lean `v4.34.0-rc1`; the final build of the archive completed with 9,022 jobs. An axiom audit of the main Lean theorems of the formal bundle, as listed in the build records, reports only `propext`, `Classical.choice`, and `Quot.sound`; not every auxiliary name of Appendix A was audited individually. No module on which the results depend uses `sorry`, `admit`, a new `axiom`, or `native_decide`. The formal names are listed in Appendix A.

8.2. Preregistered computation

All production runs behind Theorems 2 and 4, Corollary 1.1, and the audits of Section 5.3 were executed under a frozen preregistration. Some target structures had been observed beforehand: the depth-eight grouping S-A1 and the terminal witness pair were treated as preregistered replication or control targets rather than as new discoveries. The nonterminal A2′ classification and the parity test were evaluated under the frozen, outcome-blind protocol. The preregistration fixed the state sets, the certificates of future equivalence, the strata, the decision rules, the predictions, and the rules for selecting witnesses. Its SHA-256 hash is `6ea3a02a…73f0`. The computation proceeded in a chain of stages, with outputs frozen with hashes; the scripts of stages S1–S5 were frozen before their production runs:

- S0 the grammar core, validated on all 5,095 states within depth five against an independent reimplementation;
- S1 the 562 classes within depth eight;
- S2 the 64 terminal runs and the 256 anchors;
- S3 the complete descendant DAGs, exact canonical futures, and the classification into strata;
- S4 the touch audits and the terminal witness;
- S5 the parity test, run once after the final class manifest had been frozen.

Two steps departed from this order of freezing. The validation gate of S0 was run before its script was frozen; the script hashes recorded at that run are unchanged. The final class manifest, which re-verifies the hashes of the frozen class lists and bundles them without new computation, was written and run in one step. The general statement of Theorem 1 was proved in Lean before the final class manifest was frozen; the class lists of S1 and S3 had been frozen earlier. The parity Φ was not evaluated on any state of the frozen evaluation sets before that point, with the exception of the control pair. All 75 files of the chain, together with the three frozen external inputs on which the production runs operate (the catalogue, the grammar core, and the archive of [2]), are listed with their hashes in a machine-readable ledger; the hashes were recomputed and agree with the hashes in the frozen records wherever such a record exists. Six further inputs fixed by the preregistration (earlier exploratory scripts and their outputs, used as expected values in gates) and the supplementary witness record of Section 6.3 are listed with their hashes in a dated addendum to the ledger, which leaves the ledger itself unchanged; these hashes were also recomputed and agree.

Deviations

Two deviations from the preregistered development range were recorded before the results they could affect. A timing run used an unregistered seed and recorded only the time per move and the final number of cells. A development run of the finite-horizon audit previewed one pair of the kind later examined, so that audit is not fully blind (Section 5.3). Neither

deviation involved any evaluation of Φ or the terminal witness. Other procedural changes made before the affected production runs, such as the replacement of a stage script after a run was stopped for memory and a change of storage format, are recorded in the frozen records.

8.3. Remaining assumptions

1. The frozen catalogue certificate (K1 and $165 \in K.\mathrm{cells}$) is audited independently, not proved in Lean; so is the identification of the frozen source file with a catalogue object of the Lean statements, since the 691 rows are not encoded in Lean.
2. The agreement of the computational model with the formal definitions of availability, legality, and $R_K^{(2)}$ is supported by the independent reimplementation used in stage S0 (on all states within depth five) and in the terminality check of Proposition 2, and by the consistency checks of Section 5.3; it is not a Lean theorem.
3. Candidate pairs of the finite-horizon audit were indexed by 128-bit digests under the assumption of no collision; this affects only which candidates were found, since every selected pair was re-checked exactly.
4. All computational conclusions concern the examined sets.

8.4. Open problems

1. *The converse of Corollary 3.2* (Section 5.4): do equal legal moves and availability differences confined to permanently untouchable cells imply future equivalence? A positive answer requires control of the influence of a difference cell on its neighbours. A note added to Section 5.4 records that the converse has since been proved in Lean for the abstract attachment system under K1; its instantiation for the frozen catalogue is not given here.
2. *The mechanism of the parity splits.* Which geometric features of a future-equivalence class determine whether Φ splits it? In particular, how are parity splits related to untouchable availability differences in stratum $\mathrm{A2}'$? Any analysis of this question will be exploratory and reported separately from the preregistered results.
3. *The gap near the seed.* Whether future-equivalent states within depth eight are separated by $R_K^{(2)}$ was not examined.
4. *Formal certificates for the witnesses.* The terminal and nonterminal witnesses of Section 4 and the parity witnesses of Section 6 are certified computationally; Lean certificates are possible in principle.
5. *Growing catalogues.* All results concern the frozen catalogue. Systems in which the catalogue grows are outside the present scope.

8.5. Conclusion

On a frozen AKN-type patch with a single-face attachment grammar, four facts about the relation between current states and complete legal futures have been established.

- Under K1, the normalised code $R_K^{(2)}$ is an exact complete invariant of the one-step legal interface (Theorem 1).
- The future quotient is strictly coarser, already among reachable states that still have legal moves (Theorem 2).
- On reachable states, every availability difference between future-equivalent states lies in a cell that the available side can never touch; this is a necessary condition, and whether it is sufficient for future equivalence together with equal legal

moves is not established here (Theorem 3; see the note added to Section 5.4).

- Beyond the future quotient, a parity readable from the current state can survive: it splits examined future-equivalence classes in both structural strata, and a parity difference persists along every common continuation (Theorem 4, Corollary 4.1).

The construction thus provides an explicit finite example of a current-state distinction that the complete legal future does not determine and that every common legal continuation preserves. No more general principle is claimed.

Appendix A. Formal names and claim firewall

A.1. Formal correspondence

The availability predicates of the formal sources (that of the archive of [1] and that of the abstract attachment system, which are proved equivalent) take an arbitrary cell as argument, whereas $\mathrm{Avail}(s)$ in the text is the set of available catalogue cells (Section 2.2). The formal statements therefore quantify over $c \in K.\,\mathrm{cells}$ where the text compares the sets $\mathrm{Avail}(s)$; the two forms are equivalent. The exact bridges between the abstract system and the archive of [1] are `waterwayIID1bIncBridgeV1`, `waterwayIID1bLegalBridgeV1`, `waterwayIID1bWordLegalBridgeV1`, `waterwayIID1bFutEqBridgeV1`, and `waterwayIID1bAvailBridgeV1`. Namespaces: `WaterwayII.Cat` for the abstract system, `A5DiscreteFoundation` otherwise.

Statement	Formal name(s)	Hypotheses
Lemma 1	`waterwayIID1bR2SlotTrichotomyV1`, `waterwayIID1bR2SlotEqLiveOneIffV1`, `waterwayIID1bR2SlotEqLiveZeroIffV1`, `waterwayIID1bR2SlotEqDeadIffV1`, `waterwayIID1bR2EqIffLegalFreshV1`; every catalogue cell has a slot: `waterwayIID1bCellHasFaceV1`	none
Lemma 2	`availIffFreshOrLegal`, `waterwayIID1bAvailIffLiveSlotV1`, `waterwayIID1bAvailIffLiveSlotOfCertificateV1`; direction $\Leftarrow$: `availOfFresh`, `availOfLegal`	K1 (or face-owner certificate); $\Leftarrow$: none
Theorem 1(a)	`waterwayIID1bRcatEqImpEqV1`	none
Theorem 1(b)	`waterwayIID1bR2EqIffAvailLegalEqV1`, `waterwayIID1bR2ExactlyDeterminesOneStepInterfaceV1`, `waterwayIID1bR2ExactlyDeterminesOneStepInterfaceOfCertificateV1`	K1 (or face-owner certificate)
Corollary 1.1	`waterwayIID1bR2CollisionSameKeyReachableV1` (Lean part)	frozen catalogue certificate; enumeration [E]
Proposition 1	Theorem 1(b) and a one-line argument	K1
Corollary 2.1	[1], Main Theorem G(i) and the witnesses of Theorem 2 [E]	none
Theorem 3	`waterwayIID1bConstructiveSeparationV1` (reachable form), `waterwayIID1bConstructiveSeparationOfSubsetV1` (catalogue-contained form); abstract form `distinguishingWord`, `firstTouch`	frozen catalogue certificate; face-owner certificate and $s \subseteq K$. cells; abstract: K1, K2
Corollary 3.1	`waterwayIID1bTouchCorollaryV1`, `waterwayIID1bTouchCorollaryCurrentV1`; abstract form `touchLemma`, `touchLemmaCurrent`	frozen catalogue certificate; abstract: K1, K2
Corollary 3.2	Theorem 1(b), Corollary 3.1, reachable $\subseteq K$, and a short argument	frozen catalogue certificate
K2	`waterwayIID1bSharedFaceUniqueV1`, `waterwayIID1bK2ForCatalogueV1`	none
K1 from certificate	`waterwayIID1bK1OfCertificateV1`	face-owner certificate
Reachable $\subseteq K$	`waterwayIID1bReachableFromSeedSubsetCatalogueV1`	$165 \in K$. cells
Used from Waterway I [1]		

Statement	Formal name(s)	Hypotheses
Main Theorem A	`differenceBarrierD1bCompleteCatalogueResidueIsFutureSufficientV1`	none
Main Theorem G(i)	`differenceBarrierD1bCompleteCatalogueExtendedNormalizedResidueIsFutureSufficientV1`	none
Main Theorem G(ii)	`differenceBarrierD1bB2H4ReachableStrictCoarseningWitnessV1`	local contact certificate
Used from [2]		
Φ	`differenceBarrierD1bRelativeSideExplicitPotentialV1`	—
σ	`differenceBarrierD1bRelativeSideMoveLabelV1`	—
Eq. (1)	`differenceBarrierD1bRelativeSideParityEndpointReadableV1`, obtained from `differenceBarrierD1bRelativeSidePotentialLegalWordV1` (hypothesis: integrability of the potential) and `differenceBarrierD1bRelativeSideExplicitPotentialIntegratesV1`	none

Table A1. Statements of this paper and their formal names.

A.2. Claim firewall

Claim	Status	Where
Slot trichotomy; $R_K^{(2)}$ equality as agreement of legality and freshness	[M]	Lemma 1
Availability recovery	[M] under K1	Lemma 2
R_{cat} injective on catalogue states	[M]	Theorem 1(a)
$R_K^{(2)}$ is an exact complete invariant of the one-step interface	[M] under K1	Theorem 1(b)
The 562 classes are all identifications within depth eight	[M*] + [E]	Corollary 1.1
Future-equivalent pairs separated by $R_K^{(2)}$ have equal legal moves and different availability	[$\mathrm{M}_{\mathrm{paper}}$] under K1	Proposition 1
Terminal witness (seeds 0, 1)	[E] (reproduction)	Proposition 2
Nonterminal non-minimality, 5,384,241 certified pairs	[E]	Theorem 2
$R_K^{(2)}$ strictly finer than $\sim_F$ on reachable states	[M] ([1]) + [E]	Corollary 2.1
Constructive separation	[M*]	Theorem 3
Availability differences between future-equivalent states (s reachable) are permanently untouchable from the available side	[M*]	Corollary 3.1
Localisation of the gap	[$\mathrm{M}_{\mathrm{paper}}$] from [M*]	Corollary 3.2
Differing slots of $R_K^{(2)}$ belong to availability-difference cells	[$\mathrm{M}_{\mathrm{paper}}$] from [M], [M*]	Remark in Section 5.2
No touch violation on all 5,384,241 A2′ pairs; finite-horizon separation checks	[E] (audit)	Section 5.3
Converse of the localisation	[O] in this paper; proved later for the abstract system (note added)	Section 5.4
Φ splits future-equivalence classes in A1-P and A2′	[E]	Theorem 4
Terminal split (not structural)	[E]	Section 6.3
Persistence of a parity difference along common futures	[$\mathrm{M}_{\mathrm{paper}}$] from [M]	Corollary 4.1
Mechanism of the parity splits	[O]; not claimed	Section 8.4
Frequencies outside the examined sets	not claimed	Sections 4, 6
K1 and $165 \in K$. cells for the frozen catalogue; identification of the source file with a Lean catalogue object	external audit	Sections 2.4, 8.3
AKN-type substrate; grammar not an AKN matching rule	citation; not claimed	Section 7
Physical memory, erasure, irreversibility	not claimed	Section 1.5

Table A2. Every claim with its evidential status. [M]: Lean theorem (hypotheses as stated); [M*]: Lean theorem with the audited frozen catalogue certificate as hypothesis; [E]: frozen preregistered computation or audit; [M_{paper}]: short mathematical argument combining formal results, not itself formalised; [O]: open.

Statements and Declarations

Funding

This research received no external funding.

Potential Competing Interests

The author declares no competing interests.

Ethics

Not applicable. This work is theoretical and formal and does not involve human participants, animals, clinical data, or identifiable personal data.

Data and Code Availability

No experimental or participant data were generated. The Lean sources, the frozen catalogue and its face-owner certificate, the preregistration, the scripts, the frozen outputs, the result ledger with its dated addendum, the witness record, and Supplementary Provenance S1 are deposited at Zenodo, https://doi.org/10.5281/zenodo.24053439; each archive is identified below by its SHA256 digest. The Lean sources are provided in the archive

`WATERWAY_II_LEAN.zip`

`f37a4f4e584faaaee570f64bccd4ee8d9306291f3d206c5c785c8239273e3391`

It contains the recursive `A5` import closure of the modules of this paper and of the Waterway I theorems cited in Appendix A (47 Lean files, importing only Mathlib outside `A5`), an entry point, an axiom audit, a theorem map, and a SHA256 manifest. Of these files, 37 are byte-identical to files of the Waterway I archives [1], deposited at https://doi.org/10.5281/zenodo.23427992; the closure contains no module of the parity-register branches, and the formal names of Φ, σ, and Eq. (1) belong to the archive of [2]. The computational record is provided in the archives

`WATERWAY_II_P3_PIPELINE.zip`

`5e4b375c219d880fb83a28585a4b611082d70ac43ba567e5ed84dee75823963e`

`s3_nodes.sqlite.gz`

`98d488a6d0e93627efc250f9cac7627049c63c537e32d1737c72f743528dca41`

`WATERWAY_II_PRE_P0_INPUTS.zip`

`8fd3af6761917c33a87270671294c4a063b1a08f194254421010d9937a8e57ac`

together with the result ledger `P3_RESULT_LEDGER.json` (SHA256 `6d9c79c497040e18b6395f7c56e58306089b1a8b53f2930a3796c9310f7a4625`), its dated addendum `P3_RESULT_LEDGER_ADDENDUM_2026-10-13.json` (SHA256 `b5397555f3a1a25629377dedd2081b514bda3473384f76c629112d49fba616c8`), and the witness record `WATERWAY_II_SEC6_WITNESS_RECORD_v1.json` (SHA256 `f4b5ee1e44f560d8e05afde8d4e13fe6167cc5616a9111ee8cd51fde3a6bb46d`). Every file of the computational chain is identified by its digest in the ledger or in the addendum, and the deposit includes a script, without third-party dependencies, that checks all of them. The frozen catalogue `t7_ref2_tiles.json` (SHA256 `aca99e0c9b55398d7dd0cd35831996fdad5e9250aa4a4bcb23961329b6fd60f9`) is byte-identical to the file of the Waterway I deposit; the face-owner certificate and its standalone audit script are deposited with it. The preregistration and the stage records are in Japanese; Section 8 and Appendix A summarise them.

Author Contributions

M. Numagaki is the sole author and is responsible for the conception of the study, the mathematical analysis, the formalisation strategy, the design and freezing of the preregistration, provenance decisions, interpretation of the results, and writing and revision of the manuscript.

Use of Generative AI

OpenAI ChatGPT was used during manuscript preparation for English translation and copyediting, LaTeX formatting, organisation and readability improvements, and assistance in drafting or revising formal-code text. Anthropic Claude was used during the final revision for consistency checks of the manuscript against the companion papers, the Lean source files, and the catalogue data; for re-running the supplied audit scripts; for locating bibliographic details; and for drafting revised English passages and LaTeX edits. The mathematical claims, source identities, formal-proof acceptance decisions, build-status assertions, scientific interpretations, and conclusions were reviewed and approved by the author. The author takes full responsibility for the manuscript and the accompanying formal source package.

Acknowledgements

The author thanks the Lean 4 and Mathlib developer communities for the open-source formal tools used in this work [7][8].

Footnotes

[1] Formal names in the companion archive: `differenceBarrierD1bRelativeSideExplicitPotentialV1` (the potential), `differenceBarrierD1bRelativeSideMoveLabelV1` (the move label), `differenceBarrierD1bRelativeSideParityEndpointReadableV1` (the integration law (1)). See Appendix A.

References

1. a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p, q, r, s, t, u, v, w, x, y, z, ^, ^, ^, ^, ^, ^, ^, ^, ^, ^, ^Numagaki M (2026). "Exact Future Representation, Reachable Compression, and Accessibility Deformation on a Finite AKN-Type Rhombohedral Patch." Manuscript v14_r1. Formal code, frozen catalogue, and computational audits: Zenodo. doi:10.5281/zenodo.23427992.
2. a, b, c, d, e, f, g, h, i, jNumagaki M (2026). "Exact Current-State Compilation of a History-Derived Parity in a Finite State-Transition Realization on an Ammann–Kramer–Neri-Type Rhombohedral Substrate." Qeios. Preprint v2. doi:10.32388/L7NTKX.2.
3. a, bMyhill J (1957). "Finite Automata and the Representation of Events." In: Fundamental Concepts in the Theory of Systems. WADC Technical Report 57-624, Wright Air Development Center. p. 112–137.
4. a, bNerode A (1958). "Linear Automaton Transformations." Proc Am Math Soc. **9**(4):541–544. doi:10.1090/S0002-9939-1958-0135681-9.
5. a, bPark DMR (1981). "Concurrency and Automata on Infinite Sequences." In: Theoretical Computer Science. Lecture Notes in Computer Science; vol. 104. p. 167–183. doi:10.1007/BFb0017309.
6. a, bMilner R (1989). Communication and Concurrency. Prentice Hall International Series in Computer Science. New York: Prentice Hall.
7. a, bde Moura L, Ullrich S (2021). "The Lean 4 Theorem Prover and Programming Language." In: Platzer A, Sutcliffe G, editors. Automated Deduction – CADE 28. Lecture Notes in Computer Science; vol. 12699. Cham: Springer. p. 625–635. doi:10.1007/978-3-030-79876-5_37.
8. a, bThe mathlib Community (2020). "The Lean Mathematical Library." In: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs. New York, NY, USA: Association for Computing Machinery. p. 367–381. doi:10.1145/3372885.3373824.
9. a, bKramer P, Neri R (1984). "On Periodic and Non-Periodic Space Fillings of E^{m} Obtained by Projection." Acta Crystallogr A. **40**(5):580–587. doi:10.1107/S0108767384001203.
10. ^Paige R, Tarjan RE (1987). "Three Partition Refinement Algorithms." SIAM J Comput. **16**(6):973–989. doi:10.1137/0216062.
11. ^Lind D, Marcus B (1995). An Introduction to Symbolic Dynamics and Coding. Cambridge: Cambridge University Press.
12. ^Rupe A, Crutchfield JP (2022). "Algebraic Theory of Patterns as Generalized Symmetries." Symmetry. **14**(8):1636. doi:10.3390/sym14081636.
13. ^Ormes N, Pavlov R (2016). "Extender Sets and Multidimensional Subshifts." Ergodic Theory Dyn Syst. **36**(3):908–923. doi:10.1017/etds.2014.71.
14. ^Cousot P, Cousot R (1977). "Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints." In: Proceedings of the ACM Symposium on Principles of Programming Languages (POPL '77). New York, NY, USA: ACM. p. 238–252. doi:10.1145/512950.512973.
15. ^Shalizi CR, Crutchfield JP (2001). "Computational Mechanics: Pattern and Prediction, Structure and Simplicity." J Stat Phys. **104**(3–4):817–879. doi:10.1023/A:1010388907793.
16. ^Crutchfield JP, Ellison CJ, Mahoney JR (2009). "Time's Barbed Arrow: Irreversibility, Crypticity, and Stored Information." Phys Rev Lett. **103**(9):094101. doi:10.1103/PhysRevLett.103.094101.

17. ^ *Murata T (1989). "Petri Nets: Properties, Analysis and Applications." Proc IEEE. **77**(4):541–580. doi:10.1109/5.24143.*
18. ^ *Katz A (1988). "Theory of Matching Rules for the 3-Dimensional Penrose Tilings." Commun Math Phys. **118**(2):263–288. doi:10.1007/BF01218580.*
19. ^ *Hammock D, Fang F, Irwin K (2018). "Quasicrystal Tilings in Three Dimensions and Their Empires." Crystals. **8**(10):370. doi:10.3390/cryst8100370.*
20. ^ *Onoda GY, Steinhardt PJ, DiVincenzo DP, Socolar JES (1988). "Growing Perfect Quasicrystals." Phys Rev Lett. **60**(25):2653–2656. doi:10.1103/PhysRevLett.60.2653.*
21. ^ *Hann CT, Socolar JES, Steinhardt PJ (2016). "Local Growth of Icosahedral Quasicrystalline Tilings." Phys Rev B. **94**(1):014113. doi:10.1103/PhysRevB.94.014113.*

Supplementary data: available at https://doi.org/10.32388/8C43EK

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.