THE HARDEST INSTANCE OF $\mathrm{PCP}[2,w]$: ZHAO'S TABLE, THE VALUE 12 AT $w=7$, AND THE EXTREMAL SUBCLASS AT $w=10$

KOREA SUPERINTELLIGENCE LABS

Abstract. $\mathrm{PCP}[2,w]$ is the set of Post correspondence instances with two tiles, each tile a pair of nonempty binary words of length at most w. Ling Zhao's 2002 search tabulated, for $w \le 6$, the largest shortest-solution length over the class — the values $2,4,6,8,10$ — and recorded Lorentz's conjecture that the maximum is always attained by the instance $(1^n0\,/\,1)$, $(1\,/\,01^n)$ with $n = w-1$, whose optimum is $2n = 2w-2$. We prove that this family has shortest solution exactly $2w-2$ at every width and that its optimal solution is unique; we turn two of the three unsolvability filter types Zhao attributes to Lorentz, and a divisibility law for solution lengths, into theorems; we make every cell of Zhao's published table an exhaustive theorem inside a fixed depth window; we exhibit an instance of $\mathrm{PCP}[2,7]$ whose shortest solution has length exactly 12, the value the conjecture predicts at the first width Zhao did not compute; and we settle the conjecture, inside that same depth window, on the *extremal* subclass — both tiles maximally imbalanced, which is where the conjectured maximiser lives — for every $w \le 10$, one width beyond every exhaustive computation over the full class that we know of. Lorentz's conjecture itself remains open. Every theorem and proposition below is machine-checked in Lean 4; the two upper bounds that are not — the full class at $w = 7,8,9$, and the extremal subclass at $w = 11$ — are labelled as external computations where they appear.

1. Introduction

An instance of Post's correspondence problem is a finite list of *tiles*, each tile a pair of nonempty words over a fixed alphabet; a *solution* is a nonempty finite sequence of tiles, repetitions allowed, whose top strings concatenate to the same word as its bottom strings. Solvability is undecidable in general. It becomes decidable when the number of tiles is fixed at two: this is the theorem of Ehrenfeucht, Karhumäki and Rozenberg [4], simplified for the binary case by Halava, Harju and Hirvensalo [5].

Fixing the number of tiles at two turns the decision problem into a quantitative one. Write $\mathrm{PCP}[2,w]$ for the set of instances with exactly two tiles, all four strings nonempty binary words of length at most w, and $\sigma(T)$ for the length — the number of tiles laid, with repetition — of a shortest solution of a solvable instance T. The question is

$$\text{how large can } \sigma(T) \text{ be, over the solvable } T \in \mathrm{PCP}[2,w]\,? \tag{1}$$

Ling Zhao [1, 2] answered (1) by exhaustive search for $w = 2,\dots,6$, obtaining $2,4,6,8,10$, and recorded a conjecture of Lorentz [3]. Verbatim, [1, §5.2], with the display reproduced as it is set there:

> One special phenomenon we observed is that the hardest instances in PCP subclasses with size 2 can all be represented in the following form:
>
> $$\begin{pmatrix} 1^n0 & 1 \\ 1 & 01^n \end{pmatrix}$$
>
> It is not hard to prove that the optimal length of this kind of instances is $2n$. Lorentz conjectured that such an instance might always be the hardest one in $PCP[2,n+1]$ in the general case. If the conjecture is true, it will lead to a much simpler proof that $PCP[2]$ is decidable than the existing one [2]. Our experimental results support the conjecture in the cases from $PCP[2,2]$ to $PCP[2,6]$.

The display is a 2×2 array whose *columns* are the two tiles: the first is $(1^n0\,/\,1)$ and the second is $(1\,/\,01^n)$. That reading is not an inference from the value $2n$; [2, Table 6.8] prints the five instances themselves, from $(10/1), (1/01)$ at $w = 2$ to $(1^5 0/1), (1/01^5)$ at $w = 6$, against optimal lengths $2, 4, 6, 8, 10$. So the conjecture predicts the answer to (1) to be $2w-2$ at every width, attained by one explicit family.

Lorentz's paper was obtained as well, and everything in it about this family is one paragraph, [3, pp. 227–228], verbatim:

> We do suspect, though, that for size 2, the most difficult instances have length two less than double the width and these maximal instances all have the same form. For example, for width 6 and size 2 the following PCP has solution length 10.
>
> $$\begin{pmatrix} 100000 & 0 \\ 0 & 000001 \end{pmatrix}$$
>
> For larger widths, the pattern is simply to extend the number of 0's in the two longer strings. For small widths it is easy to prove that these are indeed the most difficult PCP's. If this can be proven in general, it would provide a simple proof that the PCP problem with size 2 is decidable. Compare with the difficult proof of [3].

The bracketed markers in these two quotations are the sources' own reference numbers: Zhao's [2] and Lorentz's [3] are both [4], and Zhao's thesis numbers [3] as [8]. His instance and Zhao's differ by two of the symmetries Lorentz himself lists — "We may exchange 0's and 1's to form an equivalent PCP. We may reverse all the strings in an instance." — and so have the same solution lengths. Three points about the attribution, since what this paper claims to add depends on them. *One*: the value $2n$ is asserted in both sources and proved in neither. Lorentz asserts it only inside the conjecture, as "length two less than double the width", hedged by "we do suspect", and gives no argument. *Two*: what Lorentz calls easy to prove is a different statement — that these instances are *maximal*, and only "for small widths". His own exhaustive results for size 2 reach width 3 ([3, PCP (9) and (10)], of shortest solution lengths 2 and 4), so "small widths" is bounded by that. *Three*: in [3] the value and the maximality are a single conjecture; the split into an easy value and a separate conjecture is Zhao's presentation. Zhao's thesis [2, §6.9] records the connection as "Lorentz's paper [8] mentioned it and conjectured that…".

What this paper settles. Sections 2–3 set up the class, a bounded search that decides solvability within a given number of tiles, and the balance argument that makes such a search finite in practice; the balance argument also yields two of the three $O(1)$ unsolvability filter types Lorentz uses — the prefix/postfix filter and the length balance filter, named in [3, §2] and printed as [2, Lemmas 2.4.3 and 2.4.4] — and a divisibility law for solution lengths of two-tile instances that Zhao runs on one example. These are known statements; what is added is proofs, uniform in the instance, and a scope correction: as printed, the filters are stated for a class from which tiles with equal top and bottom string are excluded, and we identify exactly what the extra instances contribute (Theorem 3.5).

Section 4 proves that Lorentz's family has shortest solution exactly $2w-2$ at *every* width $w \geq 2$, and that its optimal solution is unique. The value is the one Zhao and Lorentz assert; the proof is short, and we give it because neither source gives one. Nothing in it evaluates a search.

Section 5 makes each of Zhao's five published cells a two-sided theorem: for $w \leq 6$, some instance of $\mathrm{PCP}[2, w]$ has shortest solution exactly $2w-2$, and no instance of $\mathrm{PCP}[2, w]$ has a shortest solution of length in $[2w-1, 24]$. The second half is a finite exhaustive computation over the class — 252,047,376 ordered instances at $w = 6$ — and the depth window 24 is part of the statement, not an omitted hypothesis.

Section 6 is the first width Zhao did not compute: $\mathrm{PCP}[2, 7]$ contains an instance whose shortest solution has length exactly 12, the conjectured value (Theorem 6.1). The matching upper bound at $w = 7, 8, 9$ was obtained by a separate program and is recorded here as a computation outside the formal development (Remark 6.2).

Section 7 isolates a subclass. Write $\Delta(t) = |\mathbf{u}(t)| - |\mathbf{d}(t)|$; over $\mathrm{PCP}[2, w]$ the largest possible imbalance is $|\Delta| = w-1$, and an instance is *extremal* when both its tiles attain it. Lorentz's conjectured maximiser is extremal, so "the hardest extremal instance of $\mathrm{PCP}[2, w]$ has shortest solution $2w-2$"

is his conjecture restricted to the smallest natural subclass containing it. We prove that restricted statement — with the same depth window as in Section 5, and no claim about instances whose shortest solution exceeds it — for every $w \le 10$ (Theorems 7.2 and 7.3). At $w = 10$ we found no exhaustive statement of any kind about PCP$[2, 10]$ in the literature we searched: Zhao's table stops at $w = 6$, and the exhaustive computation over the full class we are aware of stops at $w = 9$. Section 7 records what that search covered and on what date; it is a not-found, not a proof of absence.

What is *not* settled is stated in Section 9 and is worth stating twice: nothing here bears on instances outside the extremal subclass at $w \ge 7$, and Lorentz's conjecture remains open.

Related formal work. Machine-checked work on Post's correspondence problem has so far concerned decision problems rather than optima. Forster, Heiter and Smolka [7] verify in Coq a reduction of a string rewriting problem to PCP and reductions of PCP to the intersection and the palindrome problem for context-free grammars. Omori and Minamide [6] prove instances of PCP$[3, 4]$ unsolvable by finding inductive invariants over the associated transition systems, generating for each instance a proof checkable in Isabelle/HOL, and report all but two of the instances they treat as settled; [2, Table 6.10] records 3,170 instances of that subclass that Zhao's own solver left undecided. We looked for a machine-checked shortest-solution *length* for any family of instances and found none (Section 7 records what was searched, and when).

Provenance. The publisher's copy of [3] is paywalled, but the Internet Archive holds LNCS 2063 openly and that is the copy read here; the page and equation numbers cited from it are the printed ones. Both attributions this paper makes to Lorentz — the conjecture above, and the filters of Section 3 — therefore rest on his own text rather than on Zhao's report of it, and where the two disagree we have said so. Three facts about [3] are used negatively, and each was checked against the chapter itself: it contains no theorem, lemma or proof (none of those words occurs anywhere in its fifteen pages); it contains no table and no figure; and its exhaustive results for two-tile instances stop at width 3.

2. Instances, solutions, and a bounded search

Let $\Sigma = \{0, 1\}$, let Σ^* be the set of binary words with empty word ε, and $\Sigma^+ = \Sigma^* \setminus \{\varepsilon\}$. A *tile* is a pair $t = (\mathbf{u}(t), \mathbf{d}(t))$ of words, written $(\mathbf{u}(t) \,/\, \mathbf{d}(t))$; an *instance* is a finite list $T = (t_0, \dots, t_{m-1})$ of tiles.

Definition 2.1. A *solution* of T is a nonempty finite sequence $\ell = (i_1, \dots, i_k)$ of indices $i_j < m$ with

$$\mathbf{u}(t_{i_1})\,\mathbf{u}(t_{i_2}) \cdots \mathbf{u}(t_{i_k}) \;=\; \mathbf{d}(t_{i_1})\,\mathbf{d}(t_{i_2}) \cdots \mathbf{d}(t_{i_k}).$$

Its *length* is $|\ell| = k$, the number of tiles laid. T is *solvable* if it has a solution, and then $\sigma(T)$ denotes the least length of one.

Definition 2.2. PCP$[2, w]$ is the set of instances $T = (t_0, t_1)$ of exactly two tiles with $\mathbf{u}(t_j), \mathbf{d}(t_j) \in \Sigma^+$ and $|\mathbf{u}(t_j)|, |\mathbf{d}(t_j)| \le w$ for $j = 0, 1$.

There are $2^{w+1} - 2$ nonempty binary words of length at most w, hence $(2^{w+1} - 2)^4$ instances in PCP$[2, w]$ as just defined — 1,296 at $w = 2$, 810,000 at $w = 4$, 252,047,376 at $w = 6$ and 4,162,314,256 at $w = 7$.

Remark 2.3. Definition 2.2 is a strict superset of the class Zhao tabulates, and his class is defined by a formula rather than reconstructed here. Writing $\mathrm{Str}(w)$ for the nonempty binary words of length at most w and $\mathrm{Pair}(w)$ for the pairs of words from $\mathrm{Str}(w)$ that are *not identical*, [2, Definitions 2.4.4–2.4.6, §2.4.4] puts

$$|\mathrm{Str}(w)| = 2^{w+1} - 2, \qquad |\mathrm{Pair}(w)| = |\mathrm{Str}(w)|\big(|\mathrm{Str}(w)| - 1\big), \qquad |\mathrm{PCP}[s, w]| = P^s_{|\mathrm{Pair}(w)|} - P^s_{|\mathrm{Pair}(w-1)|},$$

with $P^s_n = n(n-1)\cdots(n-s+1)$ — that is, *ordered s-tuples of distinct tiles, no tile having $\mathbf{u} = \mathbf{d}$, with maximum string length exactly w*, which is also what his Definition 2.4.6 calls the "nontrivial and nonredundant" instances. At $s = 2$ the formula reads $T_w(T_w - 1) - T_{w-1}(T_{w-1} - 1)$ with $T_w = W_w^2 - W_w$ and

$W_w = 2^{w+1} - 2$, and it reproduces his five printed totals 868; 32,072; 723,088; 13,543,712; 233,747,008 for $w = 2, \dots, 6$ [2, Table 2.2, and the first column of Table 2.3] exactly. Definition 2.2 allows equal tiles, allows $\mathbf{u} = \mathbf{d}$, and includes the smaller widths. A statement about the superset is therefore formally stronger; Theorem 3.5 says the extra instances all have shortest solution 1 when they are solvable at all, which is why the maximum is unaffected.

Configurations and the step relation. A *configuration* is $\mathrm{up}(s)$ or $\mathrm{dn}(s)$ for a word s: the top row leads the bottom row by s, or trails it by s. Laying one more tile transforms configurations as follows, writing $x \sqsubseteq y$ for "x is a prefix of y":

$$\mathrm{step}\big(t, \mathrm{up}(s)\big) = \begin{cases} \mathrm{up}\big((s\,\mathbf{u}(t))_{\geq |\mathbf{d}(t)|}\big) & \mathbf{d}(t) \sqsubseteq s\,\mathbf{u}(t), \\ \mathrm{dn}\big(\mathbf{d}(t)_{\geq |s\,\mathbf{u}(t)|}\big) & s\,\mathbf{u}(t) \sqsubseteq \mathbf{d}(t), \\ \text{undefined} & \text{otherwise,} \end{cases}$$

and symmetrically on $\mathrm{dn}(s)$, where $x_{\geq k}$ is x with its first k letters removed. A sequence ℓ is a solution of T exactly when laying its tiles from $\mathrm{up}(\varepsilon)$ is everywhere defined and ends at $\mathrm{up}(\varepsilon)$.

The bounded search. Fix an instance T and a depth n. Define a Boolean $\mathrm{reach}(T, n, c)$ by $\mathrm{reach}(T, 0, c) = \text{false}$ and

$$\mathrm{reach}(T, n+1, c) = \neg\,\mathrm{pruned}(T, n+1, c) \;\wedge \bigvee_{\substack{t \in T \\ \mathrm{step}(t,c)\ \text{defined}}} \Big(\mathrm{step}(t, c) = \mathrm{up}(\varepsilon) \;\vee\; \mathrm{reach}\big(T, n, \mathrm{step}(t, c)\big)\Big),$$

with pruned as in (2) below, and put $\mathrm{hasSol}(T, n) = \mathrm{reach}(T, n, \mathrm{up}(\varepsilon))$.

Proposition 2.4 (the search decides bounded solvability)**.** *For every instance T and every n,*

$$\mathrm{hasSol}(T, n) = \text{true} \iff T \textit{ has a solution } \ell \textit{ with } |\ell| \leq n.$$

Both directions are used. The false direction turns a finished search into the statement "no solution of length at most n", the lower half of $\sigma(T) = k$. The true direction turns a finished search into an actual solution, which is what lets an exhaustive sweep certify that every instance with a long solution also has a short one without exhibiting millions of index sequences.

3. Balance, the prune, and Lorentz's filters

Definition 3.1. The *balance* of a configuration is $\beta(\mathrm{up}(s)) = |s|$ and $\beta(\mathrm{dn}(s)) = -|s|$, and the *imbalance* of a tile is $\Delta(t) = |\mathbf{u}(t)| - |\mathbf{d}(t)|$. Put $g^{\downarrow}(T) = \max\big(0, \max_{t \in T}(-\Delta(t))\big)$ and $g^{\uparrow}(T) = \max\big(0, \max_{t \in T} \Delta(t)\big)$.

Theorem 3.2 (one tile moves the balance by its imbalance)**.** *If $\mathrm{step}(t, c) = c'$ then $\beta(c') = \beta(c) + \Delta(t)$. Consequently, if laying ℓ from c is everywhere defined and ends at c', then $\beta(c') = \beta(c) + \sum_j \Delta(t_{i_j})$, and for every such ℓ,*

$$\sum_j \Delta(t_{i_j}) \;\leq\; |\ell| \cdot g^{\uparrow}(T), \qquad -\sum_j \Delta(t_{i_j}) \;\leq\; |\ell| \cdot g^{\downarrow}(T).$$

Proof sketch. The first claim is a case analysis on the four branches of step; in each branch the new overhang is an explicit suffix whose length the prefix hypothesis determines. The rest is induction on ℓ together with $\Delta(t) \leq g^{\uparrow}(T)$ and $-\Delta(t) \leq g^{\downarrow}(T)$ for $t \in T$. □

A solution ends at $\mathrm{up}(\varepsilon)$, where the balance is 0. So from $\mathrm{up}(s)$ a solution using at most n further tiles must return $|s|$ of lead, and one tile returns at most $g^{\downarrow}(T)$; this is the prune used in Proposition 2.4:

$$\text{(2)} \qquad \mathrm{pruned}(T, n, \mathrm{up}(s)) = \big[\, n \cdot g^{\downarrow}(T) < |s| \,\big], \qquad \mathrm{pruned}(T, n, \mathrm{dn}(s)) = \big[\, n \cdot g^{\uparrow}(T) < |s| \,\big].$$

Without it a depth-n search over two tiles is a tree with 2^n leaves and instances built from a single letter never terminate.

Corollary 3.3. *If ℓ is a solution of T then $\sum_j \Delta(t_{i_j}) = 0$.*

The three filters. Call a tile t *startable* if $\mathbf{u}(t)$ and $\mathbf{d}(t)$ are prefix-comparable ($\mathbf{u}(t) \sqsubseteq \mathbf{d}(t)$ or $\mathbf{d}(t) \sqsubseteq \mathbf{u}(t)$), and *endable* if they are suffix-comparable. For two tiles put

$$\mathrm{sgn}(t_0, t_1): \quad \mathbf{u}(t_0) = \mathbf{d}(t_0) \ \vee\ \mathbf{u}(t_1) = \mathbf{d}(t_1) \ \vee\ \Delta(t_0)\Delta(t_1) < 0.$$

Theorem 3.4 (Lorentz's filters)**.** *Let $\ell = (i_1, \ldots, i_k)$ be a solution of T. Then t_{i_1} is startable and t_{i_k} is endable. In particular, if $T = (t_0, t_1)$ is a solvable two-tile instance then at least one of t_0, t_1 is startable, at least one is endable, and* $\mathrm{sgn}(t_0, t_1)$ *holds.*

Proof sketch. For the first tile: the two branches of step at $\mathrm{up}(\varepsilon)$ test exactly $\mathbf{d}(t) \sqsubseteq \mathbf{u}(t)$ and $\mathbf{u}(t) \sqsubseteq \mathbf{d}(t)$, so a tile that can be laid first is startable. For the last tile: the two rows end flush, so $P\,\mathbf{u}(t) = Q\,\mathbf{d}(t)$ for the two words P, Q built from ℓ minus its last entry; if $|\mathbf{u}(t)| \le |\mathbf{d}(t)|$ this forces $Q \sqsubseteq P$ and then $\mathbf{u}(t)$ is a suffix of $\mathbf{d}(t)$, and symmetrically. For the sign condition: by Corollary 3.3 the sum of the imbalances over ℓ vanishes; if neither tile has $\mathbf{u} = \mathbf{d}$ and the two imbalances had the same strict sign, the sum of a nonempty list of terms of one strict sign would not vanish; the remaining case, one tile of imbalance 0, is Theorem 3.5. □

These are the first two of the three filter types Lorentz uses. He introduces all three in [3, §2] — "There are three filters we consider" — as the prefix filter ("There is a symmetric postfix filter as well"), the length balance filter, and the element balance filter, each with a worked instance that the filter rejects. The first two he states under an explicit hypothesis: "(We assume no trivial solutions of length 1 where both strings of the pair are the same.)" at the prefix filter, and "Again, assuming no trivial solutions of length 1" at the length balance filter. Zhao prints the same two as [2, Lemmas 2.4.3 and 2.4.4]: "A solvable instance must have one pair where one string is the other's proper prefix and another pair where one string is the other's proper postfix" and "A solvable instance must have one pair whose top string is longer than the bottom one and another one with its bottom string longer than the top one". (The sentence in Chapter 4 of the thesis that attributes the three filter types to Lorentz cites them in the short form "Lemmas 2.3, 2.4 and 2.5 in Chapter 2"; the lemmas themselves are numbered 2.4.3–2.4.5, in §2.4.2.) Both sources thus state them for a class from which a tile with $\mathbf{u} = \mathbf{d}$ is excluded — Lorentz by hypothesis, Zhao by class convention (Remark 2.3) — and over the superset of Definition 2.2 the comparability is not proper and the sign condition acquires the two disjuncts above. The following theorem is what closes that gap: it says that an instance carrying a length-preserving tile is solvable only by a solution of length one, which is exactly the triviality Lorentz's parenthesis sets aside.

Theorem 3.5 (a length-preserving tile trivialises the instance)**.** *Let $T = (t_0, t_1)$ be solvable and suppose $|\mathbf{u}(t_j)| = |\mathbf{d}(t_j)|$ for some j. Then $\mathbf{u}(t_j) = \mathbf{d}(t_j)$ for some j, and $\sigma(T) = 1$.*

Proof sketch. The first tile of any solution is startable, and a startable tile whose two strings have equal length has them equal. A short sign argument shows the balanced tile is the one that can be laid first, in each case. A tile with $\mathbf{u} = \mathbf{d}$ is by itself a solution of length one. □

So an instance of PCP$[2, w]$ carrying a length-preserving tile — one of those that Zhao's class convention (Remark 2.3) excludes — is never a maximiser of (1), at any width.

A divisibility law. Zhao runs the following argument on one example instance ([2, §2.4.3]: "the length of its solution is $f_1 + f_2 + f_3 = 23x$ which must be a multiple of 23"). In closed form, uniform in the instance, it reads:

Theorem 3.6 (divisibility of solution lengths)**.** *Let $T = (t_0, t_1)$ with $\Delta(t_0) = -p < 0 < q = \Delta(t_1)$ and let $g = \gcd(p, q)$. Then every solution ℓ of T satisfies*

$$\frac{p+q}{g} \;\Big|\; |\ell|, \qquad \textit{hence} \qquad |\ell| \;\ge\; \frac{p+q}{g}.$$

In particular if $p = q$ then every solution of T has even length.

Proof sketch. Let a, b count the occurrences of t_0, t_1 in ℓ. Corollary 3.3 gives $ap = bq$, so (a, b) is an integer multiple of the primitive vector $(q/g, p/g)$, and $|\ell| = a+b$ is a multiple of $p/g+q/g = (p+q)/g$. The lower bound follows because a solution is nonempty. □

Remark 3.7. Theorem 3.6 is a structural statement, not a useful filter for the searches of this paper: for $T \in \mathrm{PCP}[2, w]$ one has $p, q \le w - 1$, so $(p + q)/g \le 2(w - 1) \le 20$ for $w \le 11$, below the depth window 24 used throughout, and the bound never excludes anything. The neighbouring "GCD rule" of [1, §3.1.1] is a statement about configuration lengths rather than solution lengths.

4. Lorentz's family, at every width

Definition 4.1. For $w \ge 2$ put $n = w - 1$ and

$$L_w \;=\; \big((0\,/\,1\,0^n)\,,\;(0^n 1\,/\,0)\big) \;\in\; \mathrm{PCP}[2, w].$$

L_w is Zhao's $(1^n 0\,/\,1), (1\,/\,01^n)$ with the two letters exchanged and the tiles listed in the other order; equivalently, it is Lorentz's own instance $(1\,0^n\,/\,0), (0\,/\,0^n 1)$ — [3, PCP (20)] at $w = 6$ — with the two rows exchanged. Each of these operations preserves solvability and all solution lengths. This orientation is the one our search programs report as the maximiser. At $w = 7$, $L_7 = \big((0\,/\,1000000), (0000001\,/\,0)\big)$.

Theorem 4.2 (the value asserted by Zhao and Lorentz)**.** *For every $w \ge 2$, L_w is solvable and $\sigma(L_w) = 2w - 2$. Consequently, for every $w \ge 2$ there is an instance of* $\mathrm{PCP}[2, w]$ *whose shortest solution has length exactly* $2w - 2$.

Proof sketch. Write $n = w - 1$, $t_0 = (0\,/\,1\,0^n)$, $t_1 = (0^n 1\,/\,0)$, so $\Delta(t_0) = -n$ and $\Delta(t_1) = +n$.

Upper bound. Laying t_1 n times and then t_0 n times gives top word $(0^n 1)^n 0^n$ and bottom word $0^n (1\,0^n)^n$; these agree by the rotation identity $(xy)^k x = x(yx)^k$ with $x = 0^n$, $y = 1$. So $\sigma(L_w) \le 2n$.

Equal counts. By Corollary 3.3 the two tile counts a, b of any solution satisfy $bn = an$, and $n \ge 1$, so $a = b$ and $|\ell| = 2b$.

The opening is forced. From a configuration up(s) with s empty or beginning with 0, tile t_0 cannot be laid: its bottom string begins with 1 while the top row, after $\mathbf{u}(t_0) = 0$ is appended, still begins with 0, and words with different first letters are prefix-incomparable, so both branches of step fail. From up$(0\,s')$, tile t_1 is always legal and yields up$(s'\,0^n 1)$: it strips one 0 from the front and appends $0^n 1$ at the back. Starting from up(ε), therefore, the first tile is t_1 and leaves up$(0^{n-1}1)$, and an induction on k shows that while the overhang still has $k \ge 1$ zeros before its first 1, the next tile is again forced to be t_1. Hence the first n entries of any solution are all t_1 and $b \ge n$.

With the equal-count step, $|\ell| = 2b \ge 2n = 2w - 2$. No bounded search is evaluated anywhere in this argument. □

Theorem 4.3 (the optimal solution is unique)**.** *For every $w \ge 2$ the only solution of L_w of length $2w-2$ is $\ell^\star_w = 1^{w-1}0^{w-1}$, that is, tile t_1 laid $w-1$ times followed by tile t_0 laid $w-1$ times. Moreover every solution of L_w, optimal or not, has even length.*

Proof sketch. At length $2n$ both tile counts are n by the equal-count step, and the first n entries are t_1 by the forced opening; so the remaining n entries contain no t_1 and are all t_0. The parity statement is the equal-count step alone, which constrains every solution and not only the optimal one — this is why the family's value moves in steps of two, and it makes $2w - 3$ an impossible answer rather than merely a refuted one. □

Theorem 4.3 is the column headed "number of optimal solution" in [2, Table 6.8], which reads 1 in each of its five two-tile rows, extended to all widths. Theorem 4.2 is the value both sources assert, and neither proves it: in [3] it sits inside the conjecture, and [1, 2] call a proof easy without giving one.

Proposition 4.4 (independent confirmation at nine widths)**.** *For $w = 2, \ldots, 10$ the equality $\sigma(L_w) = 2w - 2$ also follows from a direct evaluation of the bounded search of Section 2: an explicit solution of length $2w - 2$ together with* $\mathrm{hasSol}(L_w, 2w - 3) = \mathrm{false}$.

Proof sketch. Nine finite computations, each a search to depth $2w-3 \leq 17$ with the prune (2) active, evaluated by the proof kernel rather than by compiled code. The two routes — Proposition 4.4 and Theorem 4.2 — share only the definitions, so their agreement is a genuine check on both. The same comparison was run at $w = 11, \ldots, 16$, where the search returns $20, 22, 24, 26, 28, 30$, matching the general theorem at six widths that neither route had previously visited. $\square$

5. The class swept exhaustively: $w \leq 6$

An exhaustive computation over PCP$[2, w]$ can only certify what a bounded search certifies, so the depth window appears in the statement.

Definition 5.1. For w, L, N write $H(w, L, N)$ for the conjunction of

(i) some $T \in \text{PCP}[2, w]$ satisfies $\sigma(T) = L$; and
(ii) every $T \in \text{PCP}[2, w]$ that has a solution of length at most N has one of length at most L.

Equivalently: L is attained, and no instance of PCP$[2, w]$ has a shortest solution of length in the interval $[L+1, N]$. Instances whose shortest solution exceeds N are not excluded by $H(w, L, N)$ and no claim is made about them.

Throughout, $N = 24$.

Theorem 5.2 (completeness of the enumeration)**.** *Let W_w be the list of nonempty binary words of length at most w. If, for all $u_0, d_0, u_1, d_1 \in W_w$, the instance $\big((u_0/d_0), (u_1/d_1)\big)$ has a solution of length at most L or has none of length at most N, then $H(w, L, N)$'s clause (ii) holds.*

Proof sketch. Every $T \in \text{PCP}[2, w]$ *is* one of the enumerated lists: destructure T as (t_0, t_1) and observe that each of its four strings lies in W_w. No symmetry reduction is used, so nothing has to be proved about the group acting on instances; the price is a factor of about two in the number of instances visited. $\square$

Theorem 5.3 (Zhao's first three cells)**.** $H(2, 2, 24)$*,* $H(3, 4, 24)$ *and* $H(4, 6, 24)$ *hold.*

Proof sketch. Clause (i) is Proposition 4.4 at $w = 2, 3, 4$. Clause (ii) is the exhaustive computation of Theorem 5.2 over 1,296, 38,416 and 810,000 ordered instances respectively, each cell a pair of bounded searches to depths L and 24. The $w = 2$ computation is carried out by the proof kernel; $w = 3$ and $w = 4$ are delegated to compiled evaluation. $\square$

For $w = 5, 6$ the same enumeration is too slow, and the filters of Section 3 are put in the loop: by Theorem 3.4 a solvable instance has a startable tile, so the outer loop runs over startable tiles only (454 of 3,844 at $w = 5$; 1,158 of 15,876 at $w = 6$) and each cell is tested in both tile orders, which covers every solvable instance with no symmetry argument; inside a cell the endable and sign tests run before any search. Separately, the two quantities $g^{\uparrow}(T), g^{\downarrow}(T)$ of (2), which the search of Section 2 recomputes at every node, are hoisted out of the recursion.

Theorem 5.4 (PCP$[2, 5]$)**.** $H(5, 8, 24)$ *holds.*

Proof sketch. Clause (i) is Proposition 4.4 at $w = 5$. Clause (ii) is the filtered exhaustive computation: 3,490,352 cells, of which 896,080 reach a search, against the 14,776,336 ordered instances of the unfiltered enumeration; 39.5 seconds of compiled evaluation. Soundness of the filtered enumeration is Theorem 3.4 together with the both-orders sweep; the hoisting is justified by an induction identifying the two searches. $\square$

Theorem 5.5 (PCP$[2, 6]$, and Zhao's table in full)**.** $H(6, 10, 24)$ *holds; hence*

$$H(2, 2, 24) \;\wedge\; H(3, 4, 24) \;\wedge\; H(4, 6, 24) \;\wedge\; H(5, 8, 24) \;\wedge\; H(6, 10, 24).$$

Proof sketch. As for Theorem 5.4: 36,768,816 cells and 7,465,392 searches against 252,047,376 ordered instances, 5 minutes 47 seconds of compiled evaluation. $\square$

Every value Zhao published for two-tile subclasses is thus a two-sided statement about the superset of Definition 2.2, inside the window $N = 24$.

6. Width seven

Theorem 6.1 (PCP$[2,7]$ attains the conjectured value)**.** *There is an instance $T \in$ PCP$[2,7]$ with $\sigma(T) = 12$; explicitly, $T = L_7 = \big((0\,/\,1000000), (0000001\,/\,0)\big)$. Moreover every solution of L_7 has length at least* 12.

Proof sketch. $L_7 \in$ PCP$[2,7]$ because all four strings are nonempty of length at most 7. The value is Theorem 4.2 at $w = 7$, and independently Proposition 4.4: the sequence $\ell_7^\star = 1^6 0^6$ is a solution of length 12, and the bounded search at depth 11 finishes empty. Both are kernel computations; neither delegates to compiled code. □

$12 = 2 \cdot 7 - 2$ is the value Lorentz's conjecture predicts at the first width Zhao did not compute. The matching upper bound is not part of the formal development:

Remark 6.2 (an external computation). A separate program, written in C and not part of the machine-checked development, enumerates every unordered pair of distinct tiles over $\{0,1\}$ with string lengths $1, \ldots, w$ and finds the shortest solution of each by branch and bound over configurations, to the depth cap 24, using the filters of Theorem 3.4. What it establishes is the same window-bounded statement as Definition 5.1: no instance has a shortest solution of length in $[2w-1, 24]$. It reproduces Zhao's five published values and reports the maximum $2w-2$ at $w = 7, 8, 9$ as well: at $w = 7$, 2,081,124,870 pairs of which 12,363,552 reach a search, 0.8 s on eight cores; $w = 8$ in 5.2 s and $w = 9$ in 1 min 21 s. At every width $3 \le w \le 9$ it reports exactly four maximising tile pairs, and they are the four images of L_w under exchanging the two letters and swapping the two rows; at $w = 2$ it reports 14, consistent with Zhao's footnote that PCP$[2,2]$ has three hardest instances up to isomorphism. Running the same program with the filters disabled at $w = 2, \ldots, 5$ returns identical counts and maxima. These numbers are computations, not theorems, and nothing in Sections 2–7 depends on them.

7. The extremal subclass

Over PCP$[2,w]$ both strings of a tile are nonempty of length at most w, so $|\Delta(t)| \le w - 1$, with equality exactly when the pair of string lengths is $(1,w)$ or $(w,1)$.

Definition 7.1. $T = (t_0, t_1) \in$ PCP$[2,w]$ is *extremal* if $|\Delta(t_0)| = |\Delta(t_1)| = w - 1$. Write Ext$(w)$ for the set of extremal instances, and $H^{\mathrm{Ext}}(w, L, N)$ for the two clauses of Definition 5.1 with PCP$[2,w]$ replaced by Ext(w).

L_w is extremal, its two tiles having shapes $(1,w)$ and $(w,1)$. So Ext(w) is the smallest natural subclass of PCP$[2,w]$ containing Lorentz's conjectured maximiser, and $H^{\mathrm{Ext}}(w, 2w-2, 24)$ is his conjecture restricted to it and to the depth window. The restriction is what makes the question reachable. There are 2^{w+1} tiles of shape $(1,w)$ and 2^{w+1} of shape $(w,1)$, and only a pair with one tile of each shape can be solvable (Theorem 3.4), so a sweep of Ext(w) visits the $2^{w+1} \cdot 2^{w+1} = 4^{w+1}$ such pairs in two tile orders each — 16,777,216 pairs at $w = 11$ — against the $(2^{w+1}-2)^4 \approx 2.8 \times 10^{14}$ ordered instances of the full class at that width.

Theorem 7.2 (the conjecture on Ext(w), $w \le 9$)**.** *$H^{\mathrm{Ext}}(w, 2w-2, 24)$ holds for $w = 2, 3, \ldots, 9$.*

Proof sketch. Clause (i) is Theorem 4.2, since $L_w \in$ Ext(w). For clause (ii): an extremal tile never has $\mathbf{u} = \mathbf{d}$, its two string lengths being 1 and $w \ge 2$, so by Theorem 3.4 a solvable extremal instance pairs a tile of shape $(1,w)$ with a tile of shape $(w,1)$. The computation therefore sweeps the 4^{w+1} such pairs in both tile orders, one pair of bounded searches per cell. These eight values are also implied by the external full-class computation of Remark 6.2; what is added here is that they are theorems and that the subclass has been isolated. □

Theorem 7.3 (the conjecture on Ext(10))**.** *$H^{\mathrm{Ext}}(10, 18, 24)$ holds: Lorentz's instance L_{10} has shortest solution length $18 = 2 \cdot 10 - 2$, and no extremal instance of* PCP$[2,10]$ *has a shortest solution of length in* $[19, 24]$. *Together with Theorem 7.2,*

$$H^{\mathrm{Ext}}(7,12,24) \ \wedge\ H^{\mathrm{Ext}}(8,14,24) \ \wedge\ H^{\mathrm{Ext}}(9,16,24) \ \wedge\ H^{\mathrm{Ext}}(10,18,24).$$

Proof sketch. Exactly as Theorem 7.2, at $w = 10$: $4^{11} = 4{,}194{,}304$ pairs of one tile of each shape, each tested in both tile orders, each cell a pair of bounded searches to depths 18 and 24. The evaluation is delegated to compiled code and took 4 minutes 16 seconds (and 574 seconds on a repeat run under higher machine load: the same check, a factor 2.2 apart). □

Theorem 7.3 is the one width in this paper that is beyond every exhaustive computation over the full class we are aware of: Zhao's published table stops at $w = 6$, and the program of Remark 6.2 stops at $w = 9$, its $w = 10$ outer loop being about 1.5×10^{11} iterations.

We looked for an exhaustive statement of any kind about PCP$[2, 10]$, by anyone, and found none. What that means precisely, as of 2026-08-28. The three papers that work on this question were read in full: [3], whose exhaustive results for two-tile instances stop at width 3 (PCP (9) and PCP (10) there, and nothing for $w \geq 4$), and [1, 2], whose tables stop at $w = 6$. Beyond those: a fixed-string scan of a local 85 GB arXiv text corpus for the strings a paper on this question would have to contain — among them `PCP[2,` in three spacings, "Lorentz conjectured", and the title of [3] in three casings — returned a single match in the whole corpus, and that one a false positive; web queries for a table of PCP$[2, w]$ past $w = 6$ returned only [1, 2] and the decidability line [4, 5]; of the works citing [3], none but Zhao's own discusses two-tile instances at all; and $14, 32, 78, 168, 358, 736, 1502, 3032$, the solvable extremal instance counts of Remark 7.4 and the most distinctive integer sequence here, is not in the OEIS. Two of those searches carry a positive control, because a query that silently returns nothing reads exactly like one that finds nothing: the OEIS query was run alongside $2, 3, 5, 7, 11, 13, 17$, which returns A000040, and the corpus scan alongside a file-level count of "post correspondence problem", which hits 124 of the corpus' part files and so establishes that the corpus does contain PCP literature. This is a not-found on a dated search rather than a claim that no such statement exists; what it does not rest on any more is an unread source.

Remark 7.4 (external confirmation, and $w = 11$). A second C program, again outside the formal development, sweeps $\mathrm{Ext}(w)$ directly for $w \leq 11$ (the 127-bit configuration word it uses suffices exactly that far). It agrees with Theorems 7.2 and 7.3 at every width, reports the maximum $2w - 2$ at $w = 11$ as well (20, in 0.88 s), and reports exactly four maximising instances at every $w = 3, \ldots, 11$ — L_w and its images under exchanging the letters and swapping the rows. The numbers of solvable extremal instances it counts are $14, 32, 78, 168, 358, 736, 1502, 3032, 6102, 12240$ for $w = 2, \ldots, 11$. The corresponding formal statement at $w = 11$ is not available: the compiled evaluation was stopped after 30 minutes 51 seconds without producing a result.

8. Controls

Exhaustive statements of the shape $H(w, L, N)$ fail silently if the check they rest on passes for the wrong reason — a bound so large that nothing tests it, a filter that skips instances, a hypothesis that is vacuous. The following are recorded as part of the development.

Proposition 8.1 (controls for the unfiltered sweep)**.** *(i) The sweep of Theorem 5.2* fails *at $w = 4$ with $L = 5$ and at $w = 3$ with $L = 3$, so the values 6 and 4 are not artefacts of a test that passes for every L. (ii) $\sigma(L_7)$ is refuted from both sides: it is neither 13 (a solution of length 12 exists) nor 11 (the depth-11 search finishes empty). (iii) The length bound in Definition 2.2 is load-bearing: $L_8 \notin$ PCP$[2, 7]$ by exactly that bound, while $\sigma(L_8) = 14 > 12$. (iv) The instance $\big((0/1), (1/0)\big)$ has no solution* of any length *— certified not by a bounded search but by an inductive invariant over configurations — and the bounded search at depth 24 also returns false on it, so the branch of the sweep that passes unsolvable instances is telling the truth. (v) The prune (2) is tight at the maximiser: the solution of L_{10} passes an overhang of 81 symbols with nine tiles left and $g^{\downarrow}(L_{10}) = 9$, so the allowance $9 \cdot 9 = 81$ is met exactly, and the prune fires at 82. (vi)* PCP$[2, 7]$ *is inhabited and contains solvable instances, so clause (ii) of Definition 5.1 is not an implication with an empty antecedent.*

Proposition 8.2 (controls for the parametric theorem)**.** *For every $w \geq 2$, neither $2w-1$ nor $2w-3$ is the shortest-solution length of L_w. At $w = 1$ the family degenerates: L_1 is the instance $\big((0/1), (1/0)\big)$*

of Proposition 8.1(iv), which has no solution at all, so the hypothesis $w \geq 2$ in Theorems 4.2 and 4.3 is load-bearing for a structural reason rather than by an off-by-one. The forced opening of Theorem 4.2 is not vacuous: the optimal solution does begin with $w-1$ copies of t_1. Finally, the search route and the inductive route agree at $w = 11, \ldots, 16$, widths neither had visited.

Proposition 8.3 (controls for the filtered and extremal sweeps)**.** *The filtered sweep fails at $w = 5$ with $L = 7$, and the extremal sweep fails at $w = 7$ with $L = 11$, at $w = 9$ with $L = 15$ and at $w = 10$ with $L = 17$. The filtered sweep agrees with the unfiltered sweep of Theorem 5.2 at $w = 2, 3, 4$, on values of L where the answer is true* and *on values where it is false, and $H(4, 6, 24)$ is re-derived through the filtered sweep — this is the control that would catch a filter that wrongly skips instances. At $w = 4$ the extremal maximum equals the full-class maximum, so the subclass does not lose the maximum where both are known.* $\mathrm{Ext}(w)$ *is inhabited and contains a solvable instance at every width.*

9. What is not settled

(1) **Lorentz's conjecture remains open.** Theorem 4.2 is the attained half only. The not-exceeded half at general w is the conjecture itself; the forced-opening argument is about one family and says nothing about the other instances of $\mathrm{PCP}[2, w]$.

(2) $\mathrm{PCP}[2, 7]$**, full class, formally.** Theorem 6.1 is one-sided. The filtered enumeration at $w = 7$ is 364,128,304 cells, about 57 minutes of interpreted evaluation at the per-cell cost measured at $w = 6$; a compiled-library evaluation is the obvious route and was not taken here. Remark 6.2 reports the value the external program obtains.

(3) $\mathrm{Ext}(11)$ **formally.** The statement $H^{\mathrm{Ext}}(11, 20, 24)$ is the same shape as Theorem 7.3 at four times the size; the evaluation was stopped after 30 minutes 51 seconds. Splitting the outer list into blocks is the natural route.

(4) $\mathrm{Ext}(w)$ **for** $w \geq 12$**, externally.** The program of Remark 7.4 refuses beyond $w = 11$: a reachable extremal configuration has overhang length a multiple of $w - 1$ bounded by half the depth cap, so a 127-bit configuration word suffices exactly while $(c/2)(w-1) + w < 126$ for the depth cap c in use, and $w = 12$ at the cap $c = 2w - 2 = 22$ needs 133 bits.

(5) **Uniqueness of the extremal maximiser.** Remark 7.4 reports exactly four maximising instances at every $w = 3, \ldots, 11$; the corresponding theorem needs a sweep computing the exact shortest length per instance rather than a two-sided bound, and is not proved here. Theorem 4.3 is uniqueness of the optimal *solution*, a different statement.

(6) **The third filter.** Lorentz's third filter, the element balance filter, is not proved here. The two sources state it differently, and neither form is covered by Theorem 3.4. In [3, §2] it is an unsolvability test, the letter analogue of the length balance filter: "If one of the two solution strings has an excess number of 1's say, and there is no pair in the PCP instance that will allow the other solution string to gain 1's then this filter will reject the instance as having no solution." In [2, Lemma 2.4.5] it is instead a simplification rule — for a fixed letter x, if the top string of every tile contains no fewer x's than its bottom string, then every tile whose top string contains strictly *more* x's may be deleted without changing solvability, and symmetrically with the rows exchanged — and Zhao notes that in that form it "cannot be directly used to prove instances unsolvable". The letter-counting equations behind both are [2, §2.4.3], the same subsection Theorem 3.6 comes from.

(7) **The depth window.** Every statement $H(w, L, 24)$ and $H^{\mathrm{Ext}}(w, L, 24)$ leaves open the existence of instances whose shortest solution exceeds 24. Removing the window for the full class is, for w large, the decidability question itself.

10. Verification

Every theorem, proposition and corollary in this paper has been formally verified in Lean 4 (toolchain `v4.32.0`); the only claims that have not are the two computations labelled as external, Remarks 6.2 and 7.4, and nothing else here rests on them. The development is free of `sorry` and depends on no library beyond *Batteries* — *Mathlib* is not imported by any module of this family — and the repository reference is to be added at submission. Proposition 2.4, Theorem 3.2, the prefix and suffix halves of Theorem 3.4, Theorems 3.5, 4.2, 4.3 and 5.2 and Proposition 4.4 depend on no axioms beyond `propext` and `Quot.sound`; the sign half of Theorem 3.4, Theorem 3.6 and the completeness proofs behind the filtered and the extremal sweeps additionally use `Classical.choice`, through case splits on propositions with no decidability instance in scope. What is delegated to compiled evaluation, through Lean's `native_decide` and its associated axiom, is exactly the large finite sweeps: one Boolean each for $w = 3, 4$ in Theorem 5.3, for Theorem 5.4, for Theorem 5.5, for each width of Theorem 7.2 and for Theorem 7.3, plus one for each control of Proposition 8.3 and for the two failing sweeps of Proposition 8.1(i); the $w = 2$ sweep of Theorem 5.3, all of Section 4, Theorem 6.1, and parts (ii)–(vi) of Proposition 8.1 together with all of Proposition 8.2, are carried out by the proof kernel. Every value quoted from a sweep was first produced by an independent implementation in C (Remarks 6.2 and 7.4), written against the same step relation, before the corresponding formal statement was attempted; the full-class program reproduces Zhao's five published values, and the two programs agree with each other and with the formal statements wherever they overlap. Timings are wall-clock on a shared machine and vary by a factor of about two with load.

References

[1] L. Zhao, *Tackling Post's Correspondence Problem*, in: Computers and Games (CG 2002), J. Schaeffer, M. Müller and Y. Björnsson (eds.), Lecture Notes in Computer Science 2883, Springer, 2003, pp. 326–344; doi:10.1007/978-3-540-40031-8_22. The copy read here is the author's PDF, `webdocs.cs.ualberta.ca/~games/PCP/paper/CG2002.pdf`; the section and equation numbers cited above are its own.

[2] L. Zhao, *Solving and Creating Difficult Instances of Post's Correspondence Problem*, MSc thesis, University of Alberta, 2002; `webdocs.cs.ualberta.ca/~games/PCP/thesis/pcp.pdf`.

[3] R. J. Lorentz, *Creating Difficult Instances of the Post Correspondence Problem*, in: Computers and Games (CG 2000), T. A. Marsland and I. Frank (eds.), Lecture Notes in Computer Science 2063, Springer, 2001, pp. 214–228; doi:10.1007/3-540-45579-5_14. The copy read here is the openly archived scan of the LNCS 2063 volume at `archive.org/details/springer_10.1007-3-540-45579-5`; page, PCP and equation numbers cited above are the printed ones. (Zhao's reference list, and several papers following it, cite this work at pages 145–159; the printed Springer volume has 214–228, and its first page carries the line "T.A. Marsland and I. Frank (Eds.): CG 2000, LNCS 2063, pp. 214–228, 2001".)

[4] A. Ehrenfeucht, J. Karhumäki and G. Rozenberg, *The (generalized) Post correspondence problem with lists consisting of two words is decidable*, Theoretical Computer Science 21(2) (1982) 119–144.

[5] V. Halava, T. Harju and M. Hirvensalo, *Binary (generalized) Post correspondence problem*, Theoretical Computer Science 276(1–2) (2002) 183–204; doi:10.1016/S0304-3975(01)00157-8.

[6] A. Omori and Y. Minamide, *Further Tackling Post Correspondence Problem and Proof Generation*, in: Proceedings of the 14th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2025), ACM, pp. 231–242; doi:10.1145/3703595.3705886.

[7] Y. Forster, E. Heiter and G. Smolka, *Verification of PCP-Related Computational Reductions in Coq*, in: Interactive Theorem Proving (ITP 2018), J. Avigad and A. Mahboubi (eds.), Lecture Notes in Computer Science 10895, Springer, 2018, pp. 253–269; doi:10.1007/978-3-319-94821-8_15; arXiv:1711.07023.