


ON THE ROTATION DISTANCE BETWEEN BINARY TREES

PATRICK DEHORNOY

ABSTRACT. We develop combinatorial methods for computing the rotation distance between binary trees, *i.e.*, equivalently, the flip distance between triangulations of a polygon. As an application, we prove that, for each n , there exist size n trees at distance $2n - O(\sqrt{n})$.

If T, T' are finite binary rooted trees, one says that T' is obtained from T by one rotation if T' coincides with T except in the neighbourhood of some inner node where the branching patterns respectively are



Under the standard correspondence between trees and bracketed expressions, a rotation corresponds to moving a pair of brackets using the associativity law. If two trees T, T' have the same size (number of inner nodes), one can always transform T to T' using finitely many rotations. The *rotation distance* $\text{dist}(T, T')$ is the minimal number of rotations needed in the transformation, and $d(n)$ will denote the maximum of $\text{dist}(T, T')$ for T, T' of size n . Then $d(n)$ is the diameter of the n th associahedron, the graph K_n whose vertices are size n trees and where T and T' are adjacent if and only if $\text{dist}(T, T')$ is one.

There exists a one-to-one correspondence between size n trees and triangulations of an $(n + 2)$ -gon. Under this correspondence, a rotation in a tree translates into a *flip* of the associated triangulation, *i.e.*, the operation of exchanging diagonals in the quadrilateral formed by two adjacent triangles. So $d(n)$ is also the maximal flip distance between two triangulations of an $(n + 2)$ -gon.

In [17], using a simple counting argument, D. Sleator, R. Tarjan, and W. Thurston prove the inequality $d(n) \leq 2n - 6$ for $n \geq 11$ and, using an argument of hyperbolic geometry, they prove $d(n) \geq 2n - 6$ for $n \geq N$, where N is some ineffective (large) integer. A brute force argument gives $d(n) = 2n - 6$ for $11 \leq n \leq 19$. It is natural to conjecture $d(n) = 2n - 6$ for $n \geq 11$, and to predict the existence of a combinatorial proof. After [17], various related questions have been addressed [13, 16, 9, 11], or [2] for a general survey, but it seems that no real progress has been made on the above conjecture.

The aim of this paper is to develop combinatorial methods for addressing the problem and, more specifically, for proving lower bounds on the rotation distance between two trees. At the moment, we have no complete determination of the value of $d(n)$, but we establish a lower bound in $2n - O(\sqrt{n})$ that is valid for each n .

Theorem. *For $n = 2m^2$, we have $d(n) \geq 2n - 2\sqrt{2n} + 1$.*

1991 *Mathematics Subject Classification.* 05C12, 20F38, 52B20.

Key words and phrases. binary tree, rotation distance, triangulations, flip, Thompson's group.

We shall develop two approaches, which correspond to two different ways of specifying a rotation in a tree. The first method takes the *position* of the subtree that is rotated into account. This viewpoint naturally leads to introducing a partial action of Thompson's group F on trees and to expressing the rotation distance between two trees T, T' as the length of the element of F that maps T to T' with respect to a certain family of generators. This approach is very natural and it easily leads to a lower bound in $\frac{3}{2}n + O(1)$ for $d(n)$. However, due to the lack of control on the geometry of the group F , it seems difficult to obtain higher lower bounds in this way.

The second approach takes *names*, rather than positions, into account: names are given to the leaves of the trees, and one specifies a rotation using the names of certain leaves that characterize the considered rotation. This approach leads to partitioning the associahedron K_n into regions separated by sort of discriminant curves. Then, one proves that two trees T, T' are at distance at least ℓ by showing that any path from T to T' through K_n necessarily intersects at least ℓ pairwise distinct discriminant curves. Progressively refining the approach finally leads to the lower bound $2n - O(\sqrt{n})$. No obstruction a priori forbids to continue up to $2n - 6$ but a few more technical ingredients will probably be needed.

The paper is organized as follows. After setting the framework in Section 1, we develop the approach based on positions in Section 2, and use it to deduce lower bounds for $d(n)$ that lie in $n + O(1)$, and then in $\frac{3}{2}n + O(1)$. Section 3 presents the approach based on names, introducing the so-called covering relation, a convenient way of describing the shape of a tree in terms of the names attributed to its leaves. This leads to a new proof for a lower bound in $\frac{3}{2}n + O(1)$. In Section 4, we introduce collapsing, an operation that consists in erasing some leaves in a tree, and use it to improve the previous bound to $\frac{5}{3}n + O(1)$. Finally, in Section 5, applying the same method in a more tricky way, we establish the $2n - O(\sqrt{n})$ lower bound.

We use \mathbb{N} for the set of all nonnegative integers.

1. TREES, ROTATIONS, AND TRIANGULATIONS

1.1. **Trees.** All trees we consider are finite, binary, rooted, and ordered (for each inner node: the associated left and right subtrees are identified). We denote by \bullet , the tree consisting of a single vertex. If T_0, T_1 are trees, $T_0 \wedge T_1$ is the tree whose left subtree is T_0 and right subtree is T_1 . The *size* $|T|$ of a tree T is the number of symbols \wedge in the (unique) expression of T in terms of \bullet and \wedge . Thus $\bullet \wedge ((\bullet \wedge \bullet) \wedge \bullet)$ is a typical tree, usually displayed as

$$(1.1) \quad \begin{array}{c} \wedge \\ / \quad \backslash \\ \bullet \quad \wedge \\ \backslash \quad / \\ \bullet \quad \bullet \end{array}.$$

Its decomposition comprises three carets, so its size is 3.

Certain special trees will play a significant role, namely those such that, for each inner node, only one of the associated subtrees may have a positive size. Such a tree is completely determined by a sequence of 0's and 1's, called its *spine*.

Definition 1.1. (See Figure 1.) For α a finite sequence of 0's and 1's, the *thin tree* with *spine* α , denoted $\langle \alpha \rangle$, is recursively defined by the rules

$$(1.2) \quad \langle \emptyset \rangle = \bullet, \langle 0 \rangle = \langle 1 \rangle = \bullet \wedge \bullet, \text{ and } \langle 0\alpha \rangle = \langle \alpha \rangle \wedge \bullet, \langle 1\alpha \rangle = \bullet \wedge \langle \alpha \rangle \text{ for } \alpha \neq \emptyset.$$

For instance, the tree of (1.1) is thin, with spines 100 and 101. Defining the spine so that it is not unique may appear surprising, but, in this way, $\langle \alpha \rangle$ has size n for α of length n , which will make statements simpler.




FIGURE 1. Thin trees, with their spines in bold: from left to right, $\langle 1001011100 \rangle$, which is also $\langle 100101101 \rangle$, the right comb of size 5, the left comb of size 4, the right zigzag of size 9, the left zigzag of size 8.

Some particular families of thin trees will often appear in the sequel, namely
 - *right combs* $\langle 111\dots \rangle$ and their counterparts *left combs* $\langle 000\dots \rangle$,
 - *right zigzags* $\langle 101010\dots \rangle$ and *left zigzags* $\langle 010101\dots \rangle$ —see Figure 1.

1.2. **Rotations.** For each vertex v (inner node or leaf) in a tree T , there exists a unique subtree of T with root in v —see (2.1) below for a more formal definition. This subtree will be called the v -subtree of T . For instance, if v is the root of T , then the v -subtree of T is T itself. Then we can define the rotations mentioned in the introduction as follows.

Definition 1.2. (See Figure 2.) If T, T' are trees, we say that T' is obtained from T by a *positive rotation*, or, equivalently, that (T, T') is a *positive base pair*, if there exists an inner node v of T such that this v -subtree of T has the form $T_1 \wedge (T_2 \wedge T_3)$ and T' is obtained from T by replacing the v -subtree with $(T_1 \wedge T_2) \wedge T_3$. In this case, we say that (T', T) is a *negative base pair*.





FIGURE 2. A base pair: T' is obtained from T by rotating some subtree of T that can be expressed as $T_1 \wedge (T_2 \wedge T_3)$ to the corresponding $(T_1 \wedge T_2) \wedge T_3$ (positive rotation)—or vice versa (negative rotation).

By construction, rotations preserve the size of a tree. Conversely, it is easy to see that, if T and T' are trees with the same size, there exist a finite sequence of rotations that transforms T into T' —see for instance Remark 2.16 below—so a natural notion of distance appears.

Definition 1.3. If T, T' are equal size trees, the *rotation distance* between T and T' , denoted $\text{dist}(T, T')$, is the minimal number of rotations needed to transform T into T' . For $n \geq 1$, we define $d(n)$ to be the maximum of $\text{dist}(T, T')$ for T, T' of size n .

By definition, we have $\text{dist}(T, T') = 1$ if and only if (T, T') is a base pair. As proved in [17, Lemma 2], the inequality $\text{dist}(T, T') \leq 2n - 6$ holds for all size n trees T, T' for $n > 10$, and, therefore, we have¹ $d(n) \leq 2n - 6$ for $n > 10$.

1.3. Triangulations. As explained in [17], there exists a one-to-one correspondence between the triangulations of an $(n + 2)$ -gon and size n trees: having chosen a distinguished edge, one encodes a triangulation by the dual graph, a tree that becomes rooted once a distinguished edge has been fixed (see Figure 3). Then performing one rotation corresponds to performing one flip in the associated triangulations, this meaning that some pattern

is replaced with , *i.e.*,

the diagonals are exchanged in the quadrilateral made by two adjacent triangles. So the rotation distance between two trees of size n is also the flip distance between the corresponding triangulations of an $(n + 2)$ -gon, and the number $d(n)$ is the maximal flip distance between two triangulations of an $(n + 2)$ -gon.




FIGURE 3. Coding a triangulation of an $(n + 2)$ -gon by a size n tree: choose a distinguished edge E , and hang the graph dual to the triangulation under the vertex corresponding to E .

1.4. Associahedra. For each n , we have a binary relation on size n trees, namely being at rotation distance 1. It is natural to introduce the graph of this relation.

Definition 1.4. (See Figure 4.) For $n \geq 1$, the *associahedron* K_n is the (unoriented) graph whose vertices are size n trees and whose edges are base pairs.

The number of vertices of K_n is $\frac{1}{2n} \binom{2n}{n}$, the n th Catalan number, and every vertex in K_n has degree $n - 1$. The fact that any two trees of the same size are connected by a sequence of rotations means that K_n is a connected graph. For all size n trees T, T' , the number $\text{dist}(T, T')$ is the edge-distance between T and T' in K_n , and the number $d(n)$ is the diameter of K_n .

The name “associahedron” stems from the fact that, when we decompose trees as iterated \wedge -products, performing a rotation at v means applying the associativity law $x \wedge (y \wedge z) = (x \wedge y) \wedge z$ to the v -subtree.

Taking the sign of rotations into account, *i.e.*, distinguishing whether associativity is applied from $x \wedge (y \wedge z)$ to $(x \wedge y) \wedge z$, or in the other direction, amounts to

¹contrary to [17], where notation changes from Section 2.3, we stick to the convention that n (and not $n - 2$) denotes the size of the reference trees




FIGURE 4. The associahedron K_4 : the vertices are the fourteen trees of size 4, or, equivalently, the fourteen triangulations of a hexagon, and the edges connect trees lying at rotation distance 1 or, equivalently, triangulations lying at flip distance 1. The diameter $d(4)$ of K_4 is 4. The thin trees $\langle 1100 \rangle$ and $\langle 0011 \rangle$ (framed) are typical examples of size 4 trees at distance 4.


orienting the edges of associahedra, as shown in Figure 5 below. This orientation defines a partial ordering on K_n , which admits the right comb $\langle 1^n \rangle$ as a minimum and the left comb $\langle 0^n \rangle$ as a maximum. This partial ordering is known to be a lattice, the Tamari lattice [19, 10, 8, 18]. Let us mention that alternative lattice orderings on K_n are constructed in [15] and [12].

2. USING POSITIONS

Hereafter we address the problem of establishing lower bound for the rotation distance $\text{dist}(T, T')$ between two trees of size n , *i.e.*, to prove that any path from T to T' through K_n has length at least ℓ for some ℓ . Any such result immediately implies a lower bound for the diameter of the corresponding associahedron.

To this end, we have to analyze the rotations that lead from T to T' and, for that, we need a way to specify a rotation precisely. Exactly as in the case of permutations and their decompositions into transpositions (instances of commutativity), we can specify a rotation (*i.e.*, an instance of associativity) by taking into account either the position where the rotation occurs, or the names of the elements that are rotated. In this section, we develop the first approach, based on positions.

2.1. The address of a rotation. The position of a vertex v in a tree T can be unambiguously specified using a finite sequence of 0's and 1's that describes the path from the root of T to v , using 0 for forking to the left and 1 for forking to the right. Such a sequence will be called a (binary) *address*. The set of all binary addresses will be denoted by \mathbb{A} . The address of the root is the *empty* sequence, denoted \emptyset . So, for instance, the addresses of the three inner nodes of the tree of (1.1) are $\emptyset, 1, 10$, whereas the addresses of its four leaves are $0, 100, 101$, and 11 :



We deduce a natural indexation of subtrees by addresses. For T a tree and α a sufficiently short binary address, the α th subtree of T , denoted $T_{(\alpha)}$, is the subtree of T whose root is the vertex that has address α . Formally, $T_{(\alpha)}$ is recursively defined by the rules


$$(2.1) \quad T_{(\emptyset)} = T \text{ for every } T, \quad \text{and} \quad \begin{cases} T_{(0\alpha)} = T_{0(\alpha)} \\ T_{(1\alpha)} = T_{1(\alpha)} \end{cases} \quad \text{for } T = T_0 \wedge T_1.$$

Note that $T_{(\alpha)}$ exists and has positive size if and only if α is the address of an inner node of T , and it exists and has size 0 if and only if α is the address of a leaf of T .

With such an indexation, we naturally attach an address with each rotation.

Definition 2.1. We say that a positive base pair (T, T') has *address* α^+ if α is the address in T (and in T') of the root of the subtree that is rotated between T and T' . The address of the symmetric pair (T', T) is declared to be α^- .

For instance, in the positive base pair



the rotation involves the subtrees of T and T' whose roots have address 10, hence the address of (T, T') is declared to be 10^+ . See Figure 5 for more examples.

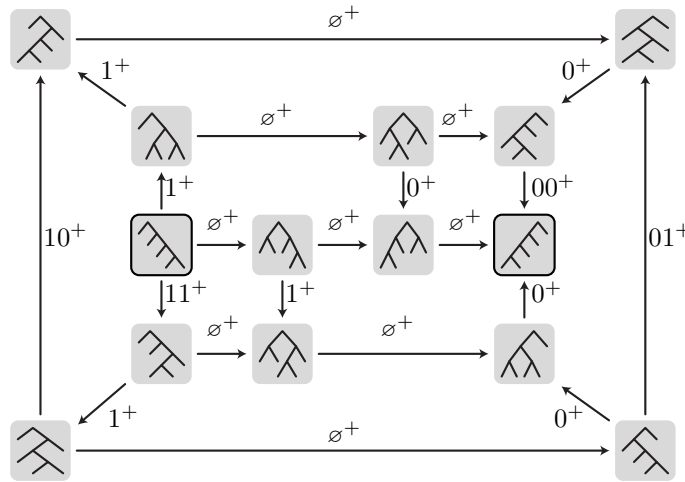


FIGURE 5. Orienting the edges of the associahedron K_4 yields the Tamari lattice, in which the right comb $\langle 1111 \rangle$ is minimal, and the left comb $\langle 0000 \rangle$ is maximal. Then taking positions into account provides for each edge a label that is a binary address.

The idea we shall develop in the sequel is to obtain lower bounds $\text{dist}(T, T') \geq \ell$ by proving that at least ℓ addresses of some prescribed type necessarily occur in any sequence of base pairs connecting T to T' , i.e., in any path from T to T' through the corresponding associahedron.

2.2. Connection with Thompson’s group F . To implement the above idea, it is convenient to view rotations as a partial action of Thompson’s group F on trees.

We recall from [3] that Thompson’s group F consists of all increasing piecewise linear self-homeomorphisms of $[0, 1]$ with dyadic slopes and discontinuities of the derivative at dyadic points. There is a simple correspondence between the elements of F and pairs of trees of equal size.

Definition 2.2. For $p \leq q$ in \mathbb{R} and T a size n tree, we recursively define a partition $\llbracket p, q \rrbracket_T$ of the real interval $[p, q]$ into $n + 1$ adjacent intervals by

$$(2.3) \quad \llbracket p, q \rrbracket_\bullet = \{[p, q]\}, \quad \text{and} \quad \llbracket p, q \rrbracket_T = \llbracket p, \frac{p+q}{2} \rrbracket_{T_0} \cup \llbracket \frac{p+q}{2}, q \rrbracket_{T_1}.$$

Then, for T, T' trees with equal sizes, we define $\Phi(T, T')$ to be the element of F that homothetically maps the i th interval of $\llbracket 0, 1 \rrbracket_T$ to the i th interval of $\llbracket 0, 1 \rrbracket_{T'}$ for each i .

Example 2.3. Assume $T = \langle 11 \rangle$ and $T' = \langle 00 \rangle$. The partitions of $[0, 1]$ respectively associated with T and T' are

$$\llbracket 0, 1 \rrbracket_T = \{[0, \frac{1}{2}], [\frac{1}{2}, \frac{3}{4}], [\frac{3}{4}, 1]\}, \quad \text{and} \quad \llbracket 0, 1 \rrbracket_{T'} = \{[0, \frac{1}{4}], [\frac{1}{4}, \frac{1}{2}], [\frac{1}{2}, 1]\}.$$

Therefore, $\Phi(T, T')$ is the element of Thompson’s group F whose graph is



i.e., the element denoted A in [3] (or x_0 in most recent references).

Infinitely many different pairs of trees represent a given element of F : if T and T' are obtained from T_0 and T'_0 respectively by replacing the k th leaf with a caret, then the subdivisions $\llbracket 0, 1 \rrbracket_T$ and $\llbracket 0, 1 \rrbracket_{T'}$ are obtained from $\llbracket 0, 1 \rrbracket_{T_0}$ and $\llbracket 0, 1 \rrbracket_{T'_0}$ by subdividing the k th intervals, and we have $\Phi(T, T') = \Phi(T_0, T'_0)$. However, for each element g of F , there exists a unique *reduced* pair of trees $(\tau(g), \tau'(g))$ satisfying $\Phi(\tau(g), \tau'(g)) = g$, where (T, T') is called reduced if it is obtained from no other pair (T_0, T'_0) as above [3, Section 2].

We can then introduce a partial action of the group F on trees as follows.

Definition 2.4. For T a tree and g an element of the group F , we define $T \bullet g$ to be the unique tree T' satisfying $\Phi(T, T') = g$, if it exists.

This is a partial action: $T \bullet g$ need not be defined for all T and g . By construction, $T \bullet g$ exists if and only if the partition $\llbracket 0, 1 \rrbracket_T$ refines the partition $\llbracket 0, 1 \rrbracket_{\tau(g)}$, *i.e.*, equivalently, if and only if the tree $\tau(g)$ is included in T . However, for each pair of trees of equal size (T, T') , there exists an element g satisfying $T \bullet g = T'$, namely $\Phi(T, T')$, and the rules for an action on the right are obeyed: if $T \bullet g$ and $(T \bullet g) \bullet h$ are defined, then $T \bullet (gh)$ is defined and it is equal to $(T \bullet g) \bullet h$ —we assume that the product in F corresponds to reverse composition: gh means “ g , then h ”. For our purpose, the main point is the following direct consequence of Definition 2.4.

Lemma 2.5. *The partial action of F on trees is free: if $T \bullet g$ and $T \bullet g'$ are defined and equal, then $g = g'$ holds.*

2.3. The generators A_α of the group F . It is now easy to describe the rotations with address α^\pm in terms of the action of an element of F .

Definition 2.6. For each address α in \mathbb{A} , we put $A_\alpha = \Phi(\langle \alpha 11 \rangle, \langle \alpha 00 \rangle)$. The family of all elements A_α is denoted \mathbf{A} .

Example 2.7. The element A_\emptyset is $\Phi(\langle 11 \rangle, \langle 00 \rangle)$, hence it is the element A (or x_0) considered in Example 2.3. More generally, if α is 1^i , *i.e.*, α consists of 1 repeated i times, A_α is the element usually denoted x_i , which corresponds to applying associativity at the i th position on the right branch of the tree. Viewed as a function of $[0, 1]$ to itself, A_α is the identity on $[0, 1 - 2^{-i}]$ and its graph on $[1 - 2^{-i}, 1]$ is that of x_0 contracted by a factor 2^i .

By definition, A_α corresponds to applying associativity at position α , and, therefore, we have the following equivalence, whose verification is straightforward.

Lemma 2.8. *For all trees T, T' with the same size, the following are equivalent:*

- (i) (T, T') is a base pair with address α^\pm ;
- (ii) $T' = T \bullet A_\alpha^{\pm 1}$ holds.

It is well-known that the group F is generated by the elements A_\emptyset and A_1 , hence, a fortiori, by the whole family \mathbf{A} . For g in F , we denote by $\ell_{\mathbf{A}}(g)$ the length of g with respect to \mathbf{A} , *i.e.*, the length of the shortest expression of g as a product of letters A_α and A_α^{-1} . Then the connection between the rotation distance and the length function $\ell_{\mathbf{A}}$ is very simple.

Proposition 2.9. *For all trees T, T' with the same size, we have*

$$(2.5) \quad \text{dist}(T, T') = \ell_{\mathbf{A}}(\Phi(T, T')).$$

Proof. Assume that (T_0, \dots, T_ℓ) is path from T to T' in $K_{|T|}$, *i.e.*, we have $T_0 = T$, $T_\ell = T'$, and (T_r, T_{r+1}) is a base pair for each r . Let $\alpha_r^{e_r}$ be the address of (T_r, T_{r+1}) . By construction, we have

$$T \bullet (A_{\alpha_1}^{e_1} \dots A_{\alpha_\ell}^{e_\ell}) = T',$$

hence $\Phi(T, T') = A_{\alpha_1}^{e_1} \dots A_{\alpha_\ell}^{e_\ell}$ by Lemma 2.5, and, therefore, $\ell_{\mathbf{A}}(\Phi(T, T')) \leq \ell$.

Conversely, assume that $A_{\alpha_1}^{e_1} \dots A_{\alpha_\ell}^{e_\ell}$ is an expression of the element $\Phi(T, T')$ of F in terms of the generators A_α . It need not be true that $T \bullet (A_{\alpha_1}^{e_1} \dots A_{\alpha_\ell}^{e_\ell})$ is defined, but we can always find an extension \widehat{T} of T (a tree obtained from T by adding more carets) such that $\widehat{T} \bullet (A_{\alpha_1}^{e_1} \dots A_{\alpha_\ell}^{e_\ell})$ is defined and equal to some extension \widehat{T}' of T' . Hence we have $\text{dist}(\widehat{T}, \widehat{T}') \leq \ell$. Now, anticipating on Section 4, there exists a set I such that T and T' are obtained from \widehat{T} and \widehat{T}' respectively by collapsing the labels of I . As will follow from Lemma 4.5, this implies $\text{dist}(T, T') \leq \text{dist}(\widehat{T}, \widehat{T}') \leq \ell$. \square

2.4. Presentation of F in terms of the generators A_α . We are thus left with the question of determining the length of an element of the group F in terms of the generators A_α . Formally, this problem is closed to similar length problems for which solutions are known. In [7] and [1], explicit combinatorial methods for computing the length of an element of F with respect to the generating family $\{x_0, x_1\}$, *i.e.*, $\{A_\emptyset, A_1\}$ with the current notation, are given. Similarly, the unique normal form of [3, Theorem 2.5] is geodesic with respect to the generating family $\{x_i \mid i \geq 0\}$ and,

therefore, there exists an explicit combinatorial method for computing the length with respect to that generating family, *i.e.*, with respect to $\{A_\emptyset, A_1, A_{11}, \dots\}$. Thus, one might expect a similar method for computing the length with respect to the family \mathbf{A} of all A_α 's. Unfortunately, no such method is known at the moment, and we can only obtain coarse inequalities.

The first step is to determine a presentation of the group F in terms of the (redundant) family \mathbf{A} . As we know a presentation of F from the x_i 's, it is sufficient to add the definitions of the elements A_α for α not a power of 1. Actually, a much more symmetric presentation exists.

If α, β are binary addresses, we say that α is a *prefix* of β , denoted $\alpha \sqsubseteq \beta$, if $\beta = \alpha\gamma$ holds for some address γ .

Lemma 2.10 ([4, Prop. 4] or [5, Prop. 2.13]). *In terms of \mathbf{A} , the group F is presented by the following relations, where α, β range over \mathbb{A} ,*

$$(2.6) \quad A_\alpha^2 = A_{\alpha 1} \cdot A_\alpha \cdot A_{\alpha 0},$$

$$(2.7) \quad A_{\alpha 0 \beta} \cdot A_\alpha = A_\alpha \cdot A_{\alpha 0 \beta},$$

$$(2.8) \quad A_{\alpha 1 0 \beta} \cdot A_\alpha = A_\alpha \cdot A_{\alpha 0 1 \beta},$$

$$(2.9) \quad A_{\alpha 1 1 \beta} \cdot A_\alpha = A_\alpha \cdot A_{\alpha 1 \beta},$$

$$(2.10) \quad A_\beta \cdot A_\alpha = A_\alpha \cdot A_\beta \quad \text{if neither } \alpha \sqsubseteq \beta \text{ nor } \beta \sqsubseteq \alpha \text{ holds.}$$

Relations (2.6) are MacLane–Stasheff pentagon relations, whereas (2.7)–(2.10) are quasi-commutation relations with an easy geometric meaning.

For a given pair of trees (T, T') , it is not difficult to find an expression of $\Phi(T, T')$ by an A_α -word, *i.e.*, a word in the letters $A_\alpha^{\pm 1}$ —see Remark 2.16 below. In general, nothing guarantees that this expression w is geodesic, and we can only deduce an upper bound $\text{dist}(T, T') \leq |w|$. However, we can also obtain lower bounds by using invariants of the relations of Lemma 2.10.

Lemma 2.11. *For w an A_α -word, let $I(w)$ be the algebraic sum of the exponents of the letters A_α with α containing no 0. Then I is invariant under the relations of Lemma 2.10.*

Proof. A simple inspection. For instance, if α contains no 0, there are two letters of the form A_{1^i} on both sides of (2.6), whereas, if α contains at least one 0, there are no such letter on either side. Also notice that the invariance property holds for the implicit free group relations $A_\alpha \cdot A_\alpha^{-1} = 1$ and $A_\alpha^{-1} \cdot A_\alpha = 1$. \square

Proposition 2.12. *An A_α -word that only contains letters A_{1^i} with positive exponents is geodesic.*

Proof. Using $|w|$ for the length (number of letters) of a word w , we have $I(w) \leq |w|$ for every A_α -word w , by definition of I . Now, if w contains only letters A_{1^i} with positive exponents, then we have $I(w) = |w|$. By Lemma 2.11, we deduce $|w| = I(w) = I(w') \leq |w'|$ for each A_α -word that is equivalent to w under the relations of Lemma 2.10. \square

Remark 2.13. By contrast, it is *not* true that a positive A_α -word, *i.e.*, an A_α -word in which all letters have positive exponents (no A_α^{-1}) need to be geodesic, or even quasi-geodesic: for $p \geq 1$, let us write $A_\alpha^{(p)}$ for $A_{\alpha 1^{p-1}} A_{\alpha 1^{p-2}} \dots A_{\alpha 1} A_\alpha$. Then

one can check that, for each p , the positive A_α -word

$$A_{\emptyset}^{(p)} A_{01}^{(p-1)} A_{0101}^{(p-2)} \dots A_{(01)^{p-2}}^{(2)} A_{(01)^{p-1}}^{(1)},$$

which has length $p(p+1)/2$, is equivalent to the A_α -word

$$A_{\emptyset}^{(p-1)} A_{\emptyset} A_0^{-1} A_{01} A_{010}^{-1} \dots A_{(01)^{p-2}} A_{(01)^{p-2}0}^{-1} A_{(01)^{p-1}},$$

which has length $3p - 2$. In other words, the submonoid of F generated by the elements A_α is not quasi-isometrically embedded in the group F .

As an application, we can determine the distance from a right comb to any tree.

Proposition 2.14. *For each tree T of size n , we have*

$$(2.11) \quad \text{dist}(\langle 1^n \rangle, T) = n - h_R(T),$$

where h_R denotes the length of the rightmost branch.

Proof. Using sh_1 for the word homomorphism that maps A_α to $A_{1\alpha}$ for each α , we recursively define an A_α -word w_T by

$$(2.12) \quad w_T = \begin{cases} \varepsilon \text{ (the empty word)} & \text{for } T = \bullet, \\ w_{T_0} \cdot A_{1^{h_R(T_0)-1}} \dots A_1 A_{\emptyset} \cdot \text{sh}_1(w_{T_1}) & \text{for } T = T_0 \wedge T_1. \end{cases}$$

An easy induction shows that, for each size n tree T , the length of w_T is $n - h_R(T)$, and that we have $\langle 1^n \rangle \bullet w_T = T$, i.e., w_T provides a distinguished way to go from the right comb $\langle 1^n \rangle$ to T .

Now, we see on (2.12) that w_T exclusively consists of letters A_{1^i} with a positive exponent. Hence, by Proposition 2.12, w_T is geodesic, and we deduce

$$\text{dist}(\langle 1^n \rangle, T) = |w_T| = n - h_R(T). \quad \square$$

Corollary 2.15. *For each n , we have $d(n) \geq n - 1$.*

Proof. Applying (2.11) when T is the left comb $\langle 0^n \rangle$ —or any size n term with right height 1—gives $\text{dist}(\langle 1^n \rangle, T) = n - 1$. \square


As can be expected, other proofs of the previous modest result can be given, for instance by counting left-oriented edges in the trees: our purpose in stating Proposition 2.14 is just to illustrate the general principle of using Thompson’s group F and its presentation from the A_α ’s.

Remark 2.16. The proof of Proposition 2.14 implies that, for each pair (T, T') of size n trees, the A_α -word $w_T^{-1} w_{T'}$ is an explicit expression of $\Phi(T, T')$ of length at most $2n - 2$. It corresponds to a distinguished path from T to T' via the right comb $\langle 1^n \rangle$ in the associahedron K_n .

2.5. Addresses of leaves. We turn to another way of using addresses to prove lower bounds on the rotation distance, namely analyzing the way the addresses of the leaves are modified in rotations.

Definition 2.17. For T a tree and $1 \leq i \leq |T| + 1$, we denote by $\text{add}_T(i)$ the address of the i th leaf of T in the left-to-right enumeration of leaves.

Equivalently, we can attribute labels 1 to $n + 1$ to the leaves of each size n tree T enumerated from left to right and, then, $\text{add}_T(i)$ is the address where the label i occurs in T . For instance, if T is the (thin) tree of (1.1), then the labelling of the leaves of T is



and we find $\text{add}_T(1) = 0$, $\text{add}_T(2) = 100$, $\text{add}_T(3) = 101$, and $\text{add}_T(4) = 11$.

The idea now is that, for each i , we can follow the parameter $\text{add}_T(i)$ when rotations are applied.

Lemma 2.18. *Assume that (T, T') is a base pair and $1 \leq i \leq |T| + 1$ holds. Then $\text{add}_{T'}(i)$ is equal to $\text{add}_T(i)$ or it is obtained from $\text{add}_T(i)$ by one of the following transformations, hereafter called special: adding or removing one 0, adding or removing one 1, replacing some subword 10 with 01 or vice versa.*

Proof. Assume that the address of (T, T') is α^+ and $\text{add}_T(i) = \gamma$ holds. Then four cases may occur. If α is not a prefix of γ , then we have $\text{add}_{T'}(i) = \gamma$. Otherwise, exactly one of $\alpha 0 \sqsubseteq \gamma$, $\alpha 10 \sqsubseteq \gamma$, or $\alpha 11 \sqsubseteq \gamma$ holds. If $\alpha 0 \sqsubseteq \gamma$ holds, say $\gamma = \alpha 0\beta$, then we have $\text{add}_{T'}(i) = \alpha 00\beta$, obtained from γ by adding one 0. If $\alpha 10 \sqsubseteq \gamma$ holds, say $\gamma = \alpha 0\beta$, then we have $\text{add}_{T'}(i) = \alpha 01\beta$, obtained from γ by replacing one 10 with 01. Finally, if $\alpha 11 \sqsubseteq \gamma$ holds, say $\gamma = \alpha 11\beta$, then we have $\text{add}_{T'}(i) = \alpha 1\beta$, obtained from γ by removing one 1. The results are symmetric for a negative base pair. \square

So Lemma 2.18 shows that the parameters $\text{add}_T(i)$ cannot change too fast, which can be easily translated into lower bounds on the rotation distance. For α, γ two binary addresses, we denote by $\#_\alpha \gamma$ the number of occurrences of α in γ .

Lemma 2.19. *For γ, γ' in \mathbb{A} , define*

$$\delta(\gamma, \gamma') = |\#_0 \gamma' - \#_0 \gamma| + |\#_1 \gamma' - \#_1 \gamma| + |\#_{10} \gamma' - \#_{10} \gamma|.$$

Then, for all trees T, T' of equal size, we have

$$(2.13) \quad \text{dist}(T, T') \geq \max_{1 \leq i \leq |T|} \delta(\text{add}_T(i), \text{add}_{T'}(i)).$$

Proof. Lemma 2.18 shows that, for each i , one rotation changes by at most one the value of $\delta(\text{add}_T(i), \text{add}_{T'}(i))$. Indeed, adding or removing one 0 can change the value of $|\#_0 \text{add}_{T'}(i) - \#_0 \text{add}_T(i)|$ by one, but it changes neither $|\#_1 \text{add}_{T'}(i) - \#_1 \text{add}_T(i)|$ nor $|\#_{10} \text{add}_{T'}(i) - \#_{10} \text{add}_T(i)|$. Similarly, exchanging 10 and 01 once can change the value of $|\#_{10} \text{add}_{T'}(i) - \#_{10} \text{add}_T(i)|$ by one, but it changes neither $|\#_0 \text{add}_{T'}(i) - \#_0 \text{add}_T(i)|$ nor $|\#_1 \text{add}_{T'}(i) - \#_1 \text{add}_T(i)|$. Thus, if $\delta(\text{add}_T(i), \text{add}_{T'}(i))$ is ℓ , then at least ℓ rotations are needed to transform T into T' . \square

Proposition 2.20. *For $T = (T_1 \wedge \bullet) \wedge T_2$, with $|T_1| = |T_2| = p - 1$, and $T' = \langle (10)^p \rangle$, we have*

$$(2.14) \quad \text{dist}(T, T') \geq 3p - 2.$$

Proof. Both T and T' have size $2p$. By construction, we have

$$\text{add}_T(p + 1) = 01 \quad \text{and} \quad \text{add}_{T'}(p + 1) = (10)^p.$$

We have $\#_0 01 = \#_1 01 = 1$, $\#_{10} 01 = 0$, and $\#_0 (10)^p = \#_1 (10)^p = \#_{10} (10)^p = p$, whence $\delta(01, (10)^p) = 3p - 2$, and (2.14) follows from (2.13). \square

Corollary 2.21. *For each n , we have $d(n) \geq \frac{3}{2}n - \frac{5}{2}$.*

Proof. Proposition 2.20 gives $d(n) \geq \frac{3}{2}n - 2$ for n even. A similar argument gives $d(n) \geq \frac{3}{2}n - \frac{5}{2}$ for n odd. \square

2.6. Addresses of leaves (continued). The method can be refined to obtain a more precise evaluation of the number of rotations needed to transform an address into a special given address, typically one of the form $1^p 0^q$. For γ a binary address, we denote by $\pi(\gamma)$ the unique path in \mathbb{N}^2 that starts on $\mathbb{N} \times \{0\}$, finishes on $\{0\} \times \mathbb{N}$ and contains one edge $(0, -1)$ for each 1 in γ and one edge $(1, 0)$ for each 0 in γ , following the order of letters in γ ; see Figure 6. Then we have the following result.

Lemma 2.22. *(See Figure 6.) For γ satisfying $\#_1\gamma \leq p$ and $\#_0\gamma \leq q$, put*

$$(2.15) \quad f(p, q, \gamma) = (p - \#_1\gamma) + (q - \#_0\gamma) + N(\gamma) + D(\gamma),$$

where $N(\gamma)$ denotes the number of squares lying below $\pi(\gamma)$ and adjacent to a coordinate axis, and $D(\gamma)$ is the distance in the \mathbb{N}^2 -grid from $(1, 1)$ to the region above $\pi(\gamma)$. Then, for all trees T, T' satisfying $\text{add}_T(i) = 1^p 0^q$, $\#_1 \text{add}_{T'}(i) \leq p$, and $\#_0 \text{add}_{T'}(i) \leq q$, we have

$$(2.16) \quad \text{dist}(T, T') \geq f(p, q, \text{add}_{T'}(i))$$




FIGURE 6. Application of Lemma 2.22 to $\gamma = 101001001$ with $p = 4$ and $q = 7$: we find $p - \#_1\gamma = 4 - 4 = 0$, $q - \#_0\gamma = 7 - 5 = 2$, the number of squares below $\pi(\gamma)$ touching one of the axes is 7, and the distance from $(1, 1)$ to the region above $\pi(\gamma)$ is 1, leading to $f(4, 7, \gamma) = 0 + 2 + 7 + 1 = 10$. Hence, if T and T' are trees in which, for some i , the i th leaf has address $1^4 0^7$ in T and γ in T' , we have $\text{dist}(T, T') \geq 10$.

Proof (sketch). As in Lemma 2.19, one checks that $f(p, q, \gamma)$ decreases by at most one when a special transformation is applied to γ . By Lemma 2.18, this implies that $f(p, q, \text{add}_{T'}(i))$ decreases by at most one when a rotation is applied to T' . \square

Proposition 2.23. *For $T = \langle 1^p 0^q \rangle$, $T' = \langle 0^q 1^p \rangle$ with $p, q \geq 1$, we have*

$$(2.17) \quad \text{dist}(T, T') \geq p + q + \min(p, q) - 2.$$


Proof. As Figure 11 will show, we have $\text{add}_T(p+1) = 1^p 0^q$ and $\text{add}_{T'}(p+1) = 0^q 1^p$, so applying Lemma 2.22 gives the lower bound $\text{dist}(T, T') \geq p + q + \min(p, q) - 2$. \square

(It is then easy to check that (2.17) is an equality by finding an explicit path from T to T' involving the expected number of rotations.) So we reobtain for $d(n)$ a lower bound $\frac{3}{2}n - 2$ for even n , and $\frac{3}{2}n - \frac{5}{2}$ for odd n . More refinements are possible, but approaches that only take one address at a time into account are unlikely to go beyond $\frac{3}{2}n + O(1)$.

3. USING NAMES

As the previous approach based on positions leads to limited results only, we now develop an alternative approach based on *names*. This is exactly similar to investigating a permutation not in terms of the positions of the elements that are permuted, but in terms of the names of the elements that have been permuted. Here we shall associate with every base pair a name $(a, b, c, d)^\pm$ consisting of four numbers and a sign. The principle for using names will be the same as with positions: we prove $\text{dist}(T, T') \geq \ell$ by showing that any path from T to T' through the associahedron $K_{|T|}$ must visit at least ℓ edges whose names satisfy some specific constraints. Here the constraints will involve the so-called covering relation, a binary relation that connects the leaves of the tree, taking the form “the i th leaf is covered by the j th leaf in T ”, denoted $i \triangleleft_T j$. The basic observation is that, if the i th leaf is not covered by the j th leaf in T , but is covered in T' , then any sequence of rotations from T to T' must contain a base pair whose name has a certain form, namely $(a, b, c, d)^+$ with $c = j$ and $a \leq i < b$. The rest of the paper consists in exploiting this principle in more and more sophisticated ways.

3.1. The name of a base pair. As in Section 2.5, we attach labels to the leaves of a tree. For each label i occurring in a tree T , we denote by $\text{add}_T(i)$ the address where i occurs in T . The only difference is that, in view of Section 4, we shall not necessarily assume that the labels used for a size n tree are 1 to $n + 1$. However we always assume that the labels increase from left to right. So, for instance,



is considered to be a legal labeling and, if T is the above tree, we would write $\text{add}_T(2) = 0$ and $\text{add}_T(6) = 101$. Formally, this amounts to hereafter considering *labeled* trees. However, every unlabeled tree is identified with the labeled tree where labels are $1, \dots, |T| + 1$.

Switching to labeled trees does not change anything to rotations and the derived notions. If T and T' are labeled trees, we say that (T, T') is a base pair if (\bar{T}, \bar{T}') is a base pair, where \bar{T} and \bar{T}' the unlabeled trees underlying T and T' , and, in addition, the same labels occur in T and T' (necessarily in the same order by our convention that labels increase from left to right). As the associativity law does not change the order of variables, this definition preserves the connection between rotation and associativity.

We now attach to each base pair (of labeled trees) a name that specifies the rotation in terms of the labels of leaves.

Definition 3.1. (See Figures 7 and 10 below.) Assume that (T, T') is a positive base pair. Let α^+ be the address of (T, T') . Then, the *name* of (T, T') , denoted $\nu(T, T')$, is defined to be $(a, b, c, d)^+$, where

- a is the unique label in T that satisfies $\text{add}_T(a) = \alpha 0^p$ for some p ,
- b is the unique label in T that satisfies $\text{add}_T(b) = \alpha 10^p$ for some p ,
- c is the unique label in T that satisfies $\text{add}_T(c) = \alpha 101^p$ for some p ,
- d is the unique label in T that satisfies $\text{add}_T(d) = \alpha 1^p$ for some p ,

In this case, the name of (T', T) is defined to be $(a, b, c, d)^-$. For $k = 1, \dots, 4$, the k th entry in $\nu(T, T')$ is denoted $\nu_k(T, T')$.





FIGURE 7. Name of a base pair of trees: $\nu_1(T, T')$ and $\nu_4(T, T')$, i.e., a and d , are the names of the extremal leaves in the subtree involved in the rotation, whereas $\nu_2(T, T')$ and $\nu_3(T, T')$, i.e., b and c , are the names of the extremal leaves in the (nested) subtree that is actually moved in the rotation.

Example 3.2. Let us consider the pair of (2.2) again, namely



With the default labels, the name of (T, T') is $(2, 3, 5, 6)^+$: indeed, 2 and 6 are the extreme labels of the subtree $T_{(10)}$ involved in the rotation, whereas 3 and 5 are the extreme labels in the subtree $T_{(1010)}$ that is actually moved.

By definition, when T is given, the address of a base pair (T, T') determines its name. The converse is true as well: if (T, T') has name $(a, b, c, d)^\pm$, then the address of (T, T') is α^\pm , where α is the longest common prefix of $\text{add}_T(a)$ and $\text{add}_T(d)$, hereafter denoted $\text{add}_T(a) \wedge \text{add}_T(d)$. This address is also $\text{add}_{T'}(a) \wedge \text{add}_{T'}(d)$.

3.2. The covering relation. The main tool that will enable us to use names to establish lower bounds for the rotation distance is a binary relation that provides a description of the shape of a tree in terms of the names of its leaves.

Definition 3.3. (See Figure 8.) If $i < j$ are labels in T , we say that i is *covered* by j in T , denoted $i \triangleleft_T j$, if there exists a subtree T' of T such that i is a non-final label in T' and j is the final label in T' .

The covering relation \triangleleft_T provides a complete description of the tree T .

Proposition 3.4. Every labeled tree T is determined by its covering relation \triangleleft_T .

Proof. Assume first that T has the default labeling, i.e., the labels are 1 to $|T| + 1$. By construction, $|T| + 1$ is the largest label in T , and every label $\leq |T|$ is covered by $|T| + 1$, so \triangleleft_T determines $|T|$. For an induction, it suffices now to show that, for $T = T_0 \wedge T_1$, the relation \triangleleft_T determines the size $|T_0|$ of T_0 . Now 1 is covered in T by $|T_0| + 1$, the final label in T_0 , but it is not covered by $|T_0| + 2$, nor is it either covered by any integer $\leq |T|$. So we have $|T_0| = \max\{j \leq |T| \mid 1 \triangleleft_T j\}$.

The argument is similar for an arbitrary labeled tree, which is determined by its shape plus the family of its labels. \square




FIGURE 8. The covering and co-covering relations: i is covered by j in T if there exists a subtree T' such that i is a non-final label in T' , whereas is the last (rightmost) label in T' . Symmetrically, i co-covers j in T if there exists a subtree T' such that i is the initial (leftmost) label in T' , and j is a non-initial label in T' .

According to the notation of (2.1), each subtree of a tree is specified by a binary address. Introducing the address of the subtree T' involved in Definition 3.3 gives the following rewording of the definition. We recall that, if α and β are binary addresses, $\alpha \sqsubseteq \beta$ means that α is a prefix of β , i.e., β consists of α possibly followed by additional 0's and 1's.

Lemma 3.5. *If i, j are labels in T , then i is covered by j in T if and only if there exists a binary address γ satisfying*

$$(3.2) \quad \gamma 0 \sqsubseteq \text{add}_T(i) \quad \text{and} \quad \text{add}_T(j) = \gamma 1^p \text{ for some } p \geq 1.$$

For future reference, we mention some simple properties of the covering relation.

Lemma 3.6. (i) *The set of labels covered by j in T is a (possibly empty) interval ending in $j - 1$: if $i \triangleleft_T j$ and $i \leq i' < j$ hold, then so does $i' \triangleleft_T j$.*

(ii) *The relation \triangleleft_T is transitive.*

Proof. Point (i) is clear from the definition: the elements covered by j in T are the non-final labels in the maximal subtree T' of T that admits j as its final label: this means that T' is either T itself, or it is some subtree $T_{(\gamma)}$ such that γ ends with 0.

For (ii), we observe that, if T' is a subtree of T witnessing for $j \triangleleft_T k$, and T'' is a subtree witnessing for $i \triangleleft_T j$, then the subtrees T' and T'' cannot be disjoint since j is the label of a leaf that belongs to both of them, and, then, T'' must be included in T' , so each non-final label in T'' is a non-final label in T' as well. \square

3.3. The co-covering relation. The definition of covering gives a distinguished role to the right side. Of course, there is a symmetric version involving the left side.

Definition 3.7. (See Figure 8.) If $i < j$ are labels in T , we say that i co-covers j in T , denoted $i \triangleright_T^* j$, if, for some subtree T' of T , the integer i is the initial label in T' , and j is a non-initial label in T' .

To avoid confusion, we shall always state the covering and co-covering properties for increasing pairs of labels $i < j$, thus saying “ i is covered by j ” rather than “ j covers i ”, and “ i co-covers j ” rather than “ j is co-covered by i ”. The counterpart of (3.2) is that i co-covers j in T if and only if there exists γ satisfying

$$(3.3) \quad \text{add}_T(i) = \gamma 0^p \text{ for some } p \geq 1 \quad \text{and} \quad \gamma 1 \sqsubseteq \text{add}_T(j).$$

The counterparts of Lemmas 3.5 and 3.6 are obviously true.

More interesting are the relations that connect covering and co-covering.

Lemma 3.8. *Assume $i < k \leq j < \ell$. Then the relations $i \triangleleft_T j$ and $k \triangleright_T^* \ell$ exclude each other.*

Proof. Assume $i < k \leq j$, $i \triangleleft_T j$ and $k \triangleright_T^* \ell$. We claim that $\ell \leq j$ holds. Indeed, let $T_{(\alpha)}$ be a subtree of T witnessing for $i \triangleleft_T j$, and $T_{(\beta)}$ be a subtree witnessing for $k \triangleright_T^* \ell$. According to (3.2) and (3.3), we have

$$\alpha 0 \sqsubseteq \text{add}_T(i), \quad \text{add}_T(j) = \alpha 1^p, \quad \text{add}_T(k) = \beta 0^q, \quad \text{and} \quad \beta 1 \sqsubseteq \text{add}_T(\ell)$$

for some $p, q \geq 1$. If α lies on the left of β , then $\alpha 1^p$ lies on the left of $\beta 0^q$ as well, contradicting the hypothesis $k \leq j$. If α lies on the right of β , then $\text{add}_T(i)$ lies on the right of $\beta 0^q$ as well, contradicting the hypothesis $i < k$. So α and β must be comparable, *i.e.*, one of the subtrees $T_{(\alpha)}, T_{(\beta)}$ is included in the other. The only case that makes $k \leq j$ possible is $T_{(\beta)}$ being strictly included in $T_{(\alpha)}$, in which case we have $i < k \leq \ell \leq j$. \square

Lemma 3.9. *Assume $a \triangleleft_T b$. Then the following are equivalent:*

- (i) *we have $\text{add}_T(a) = \gamma 0 1^p$ for some $p \geq 0$, where γ is $\text{add}_T(a) \wedge \text{add}_T(b)$;*
- (ii) *$a + 1 \triangleright_T^* b$ holds;*
- (iii) *$a \triangleleft_T i$ fails for each i in $[a + 1, b - 1]$.*

Proof. (See Figure 9.) By Lemma 3.5, the hypothesis implies that $\text{add}_T(b)$ is $\gamma 1^q$ for some positive q . Assume (i). Then $\text{add}_T(a + 1)$ must be of the form $\gamma 10^r$ for some positive r and, therefore, $a + 1$ co-covers b in T , *i.e.*, (ii) holds. On the other hand, a can be covered by no label between $a + 1$ and $b - 1$, so (iii) holds.

Conversely, we have $\gamma \sqsubseteq \text{add}_T(a + 1) \wedge \text{add}_T(b)$ under the hypotheses. Hence, if (ii) holds, then $\text{add}_T(a + 1)$ must be of the form $\gamma 10^r$, in which case $\text{add}_T(a)$ must be of the form $\gamma 0 1^p$ for some positive p . Hence, by the counterpart of Lemma 3.5, (i) holds.

Finally, if (i) fails, the address of a in T must be of the form $\gamma 0 1^p 0 \gamma'$ for some nonnegative p and some γ' . But, then, there exists a label $c > a$ whose address has the form $\gamma 0 1^r$, and we have $a \triangleleft_T c$, so (iii) fails. \square




FIGURE 9. Proof of Lemma 3.9: on the left, the case when (i), (ii), (iii) hold; on the right, the case when they fail.

3.4. The Key Lemma. We arrive at the main point, namely analyzing the influence of rotations on the covering and co-covering relations. The result is simple: a positive rotation creates some covering and deletes some co-covering, a negative rotation does the contrary. The nice point is that there exists a close relation between the name of a rotation and the covering or co-covering pairs it creates or deletes. This will directly lead to lower bounds on the rotation distance: as one rotation only changes the covering and co-covering relations by a small, well

controlled amount, if two trees T, T' have very different covering and co-covering relations, many rotations are needed to transform T into T' .

First, we note for further reference the following straightforward facts, whose verification should be obvious from Figure 7. We naturally use $i \triangleleft_T j$ for “ $i \triangleleft_T j$ or $i = j$ ”, and similarly for $i \triangleright_T^* j$.

Lemma 3.10. *If (T, T') is a base pair with name $(a, b, c, d)^\pm$, we have*

- (i) $a \triangleright_T^* i$ and $i \triangleleft_T d$ for each label i of T lying in $[a, d]$;
- (ii) $a \triangleright_T^* i$ and $i \triangleleft_T b - 1$ for each label i of T lying in $[a, b - 1]$;
- (iii) $b \triangleright_T^* i$ and $i \triangleleft_T c$ for each label i of T lying in $[b, c]$;
- (iv) $c + 1 \triangleright_T^* i$ and $i \triangleleft_T d$ for each label i of T lying in $[c + 1, d]$.

The roles of T and T' in Lemma 3.10 are symmetric, so all covering results stated for T also hold for T' .

Here comes the main point, namely the way covering and anticovering change in a base pair.

Lemma 3.11. *If (T, T') is a base pair with name $(a, b, c, d)^+$, then, for all labels i, j occurring in T ,*

- (i) $i \triangleleft_{T'} j$ holds if and only if we have either $i \triangleleft_T j$, or $a \leq i < b$ and $j = c$;
- (ii) $i \triangleright_{T'}^* j$ holds if and only if we have either $i \triangleright_T^* j$, or $i = b$ and $c < j \leq d$.

Proof. (i) For $i < a$, we have $i \not\triangleleft_T j$ and $i \not\triangleleft_{T'} j$ for $j < d$, and $i \triangleleft_T j \iff i \triangleleft_{T'} j$ for $j \geq d$. Similarly, we have $i \triangleleft_T j \iff i \triangleleft_{T'} j$ for $d < i < j$. So the point is to consider the pairs (i, j) with $a \leq i < j \leq d$. As is clear from Figure 7, nothing changes from T to T' when i and j both are in $[a, b - 1]$, or in $[b, c]$, or in $[c + 1, d]$. For the cases when i lies in one of the intervals and j in another, one sees on the figure that the only case when \triangleleft_T and $\triangleleft_{T'}$ disagree is when i lies in $[a, b - 1]$, thus corresponding to a leaf in the left subtree (the one with address $\alpha 0$), and j is corresponds to the rightmost leaf in the central tree, which has address $\alpha 10$ in T and $\alpha 01$ in T' . In this case, we have $i \not\triangleleft_T j$ and $i \triangleleft_{T'} j$.

The case of (ii) is symmetric, exchanging the roles of left and right and the inequalities. □

So, if (T, T') is a positive base pair, the only difference between the relations \triangleleft_T and $\triangleleft_{T'}$ are that $\nu_3(T, T')$ covers more elements in T' than in T . Symmetrically, the relations \triangleright_T^* and $\triangleright_{T'}^*$ are equally close, the only difference being that $\nu_2(T, T')$ co-covers less elements in T' than in T . This leads us to a criterion for recognizing that certain types of base pairs inevitably occur on any path connecting two trees.

Notation 3.12. In the sequel, we say “pair $(\dots, \dots, \dots, \dots)^\pm$ ” for “base pair with name $(\dots, \dots, \dots, \dots)^\pm$ ”. Also, we use abbreviated notation such as $(\leq i, > i, j, \dots)^+$ to refer to any pair $(a, b, c, d)^+$ satisfying $a \leq i$, $b > i$, and $c = b$.

Lemma 3.13 (Key Lemma). *Assume that the trees T, T' satisfy*

$$(3.4) \quad i \not\triangleleft_T j \quad \text{and} \quad i \triangleleft_{T'} j.$$

Then each sequence of rotations from T to T' contains a pair $(\leq i, > i, j, \dots)^+$.

Proof. Let (T_0, \dots, T_ℓ) be a path from T to T' in $K_{|T|}$, i.e., assume $T_0 = T$, $T_\ell = T'$, and (T_r, T_{r+1}) is a base pair for each r . By (3.4), we have $i \not\triangleleft_{T_0} j$ and $i \triangleleft_{T_\ell} j$, so there must exist an integer r satisfying

$$(3.5) \quad i \not\triangleleft_{T_r} j \quad \text{and} \quad i \triangleleft_{T_{r+1}} j.$$

By Lemma 3.11(i), this can occur only if we have

$$\nu_1(T_r, T_{r+1}) \leq i < \nu_2(T_r, T_{r+1}) \quad \text{and} \quad \nu_3(T_r, T_{r+1}) = j.$$

In other words, the pair (T_r, T_{r+1}) has the form $(\leq i, > i, j, \dots)^+$. □

Thus the covering relations partition the associahedra into regions. What Lemma 3.13 says is that one cannot go from one region to another one without crossing the border, which corresponds to base pairs with a certain type of name, see Figure 10.




FIGURE 10. Names of edges and covering relation in the associahedron K_4 . The regions correspond to the various possibilities for covering by 4. As $i < 4$ implies $j < 4$ for $i < j < 4$, there are four regions corresponding to $1 < 4$, to $1 \not< 4$ and $2 < 4$, to $2 \not< 4$ and $3 < 4$, and to $3 \not< 4$. Lemma 3.13 says for instance that the only way to leave the region $3 \not< 4$ is to cross a pair named $(\dots, 4, 4, \dots)^+$.

Of course, we have a symmetric statement involving the relation \triangleright^* .

Lemma 3.14. Assume that the trees T, T' satisfy

$$(3.6) \quad i \triangleright_T^* j \quad \text{and} \quad i \not\triangleright_{T'}^* j,$$

$$(3.7) \quad \text{or} \quad i \triangleright_T^* j \quad \text{and} \quad i - 1 \triangleleft_{T'} j - 1.$$

Then each sequence of rotations from T to T' contains a pair $(\dots, i, < j, \geq j)^+$.

Proof. The statement for (3.6) is the exact counterpart of Lemma 3.4 when left and right are exchanged. As for (3.7), Lemma 3.8 says that $i - 1 \triangleleft_{T'} j - 1$ implies $i \not\triangleright_{T'}^* j$, and, therefore, the hypotheses (3.7) imply the hypotheses (3.6). □

3.5. Refinements. More precise criteria will be needed in the sequel, and we shall now establish some refinements of Lemmas 3.13 and 3.14. All are based on these basic results, but, in addition, they exploit the geometric properties of the relations \triangleleft and \triangleright^* . The crucial advantage of the criteria below is that they provide stronger constraints for the parameters of the involved pairs: for instance, Lemma 3.15 specifies two of the four parameters completely (" $i + 1$ and j "), whereas Lemma 3.13

only gives one exact value (“ j ”), and not more than an inequality for another one (“ $>i$ ”). The price to pay for the improvement is a strengthening of the hypotheses and, chiefly, a disjunction in the conclusion.

Lemma 3.15. *Assume that the trees T, T' satisfy*

$$(3.8) \quad i \not\triangleleft_T j, \quad i \triangleleft_{T'} j, \quad \text{and} \quad i+1 \triangleright_{T'}^* j.$$

Then each sequence of rotations from T to T' contains a pair $(\dots, i+1, j, \dots)^+$, or a pair $(\dots, i+1, \dots, j)^-$.

Here we shall prove Lemma 3.15 directly—this can be done as a good exercise—but rather derive it from a more elaborate statement. The new refinement consists in getting constraints on three parameters at a time.

Lemma 3.16. *Assume $i < k < j$ and the trees T, T' satisfy*

$$(3.9) \quad i \not\triangleleft_T j, \quad i \triangleleft_{T'} j, \quad \text{and} \quad i+1 \triangleright_{T'}^* j.$$

Then each sequence of rotations from T to T' contains a pair $(\leq i, i < \dots \leq k, j, \dots)^+$, or a pair $(\leq i, i < \dots \leq k, \geq k, j)^-$.

Proof of Lemma 3.15 from Lemma 3.16. Assume first $j > i+1$. Put $k = i+1$. Then $i < b \leq k$ implies $b = i+1$. So Lemma 3.16 guarantees that there is at least one pair named

$$(\leq i, i+1, j, \dots)^+ \quad \text{or} \quad (\leq i, i+1, \geq i+1, j)^-,$$

which is the expected conclusion.

Assume now $j = i+1$. By hypothesis, we have $i \not\triangleleft_T j$ and $i \triangleleft_{T'} j$, so Lemma 3.13 gives a pair $(\leq i, >i, j, \dots)^+$, hence necessarily $(\dots, i+1, i+1, \dots)^+$, again of the expected type. \square

Proof of Lemma 3.16. We begin as in the proof of Lemma 3.13. Let (T_0, \dots, T_ℓ) be a path from T to T' in $K_{|T|}$. By (3.9), we have $i \not\triangleleft_{T_0} j$ and $i \triangleleft_{T_\ell} j$, so there exists a largest integer r satisfying

$$(3.10) \quad i \not\triangleleft_{T_r} j \quad \text{and} \quad i \triangleleft_{T_{r+1}} j,$$

and, as above, this requires $\nu_1(T_r, T_{r+1}) \leq i < \nu_2(T_r, T_{r+1})$ and $\nu_3(T_r, T_{r+1}) = j$, so (T_r, T_{r+1}) is a pair $(\leq i, >i, j, \dots)^+$. Write $b = \nu_2(T_r, T_{r+1})$.

If $b \leq k$ holds, (T_r, T_{r+1}) is a pair $(\leq i, k \leq \dots < i, j, \dots)^+$, and we are done.

So, from now on, we assume $k < b$. The hypotheses $i < k$ implies $i < b-1$, hence $i \triangleleft_{T_r} b-1$ by Lemma 3.10(i). As we have $b-1 < b \leq \nu_3(T_r, T_{r+1}) = j$, we deduce

$$(3.11) \quad \exists x \in [k, j-1] (i \triangleleft_{T_r} x)$$

On the other hand, by Lemma 3.9, the hypothesis $i+1 \triangleright_{T'}^* j$ implies

$$(3.12) \quad \forall x \in [k, j-1] (i \not\triangleleft_{T_\ell} x).$$

Therefore, there must exist $s \geq r$ satisfying

$$(3.13) \quad \exists x \in [k, j-1] (i \triangleleft_{T_s} x) \quad \text{and} \quad \forall x \in [k, j-1] (i \not\triangleleft_{T_{s+1}} x).$$

Choose such a s . Then, for some e in $[k, j-1]$, we have

$$(3.14) \quad i \triangleleft_{T_s} e \quad \text{and} \quad i \not\triangleleft_{T_{s+1}} e.$$

By Lemma 3.11(i), the name of (T_s, T_{s+1}) has the form $(\leq i, > i, e, \dots)^-$, hence, a fortiori, $(\leq i, > i, \geq k, \dots)^-$.

Moreover, we have $s \geq r$ by construction, so the choice of r implies that i is covered by j in T_{s+1} , and, therefore, by Lemma 3.6, so is $e + 1$ since we have $i < k \leq e < j$. On the other hand, Lemma 3.10(iv) gives $e + 1 \triangleright_{T_{s+1}}^* \nu_4(T_s, T_{s+1})$, and Lemma 3.9 then implies that $e + 1$ is covered in T_{s+1} by no element smaller than $\nu_4(T_s, T_{s+1})$. We deduce $\nu_4(T_s, T_{s+1}) \leq j$, and, even, $\nu_4(T_s, T_{s+1}) \in [k + 1, j]$ as we have $\nu_4(T_s, T_{s+1}) > \nu_3(T_s, T_{s+1}) \geq k$. Now, by Lemma 3.10(i), i is covered by $\nu_4(T_s, T_{s+1})$ in T_{s+1} , whereas, by (3.13), it is covered by no element of $[k, j - 1]$. It follows that the only possibility is $\nu_4(T_s, T_{s+1}) = j$.

Finally, let $f = \nu_2(T_s, T_{s+1})$. We already know that $f > i$ holds, and we claim that $f \leq k$ holds as well. Indeed, two cases are possible. For $f = i + 1$, the hypothesis $i < k$ directly implies $f \leq k$. For $f > i + 1$, Lemma 3.10(ii) implies that i is covered by $f - 1$ in T_{s+1} , and, therefore, (3.13) implies that $f - 1$ cannot belong to $[k, j - 1]$. As $f \leq e \leq j - 1$ is true by construction, the only possibility is $f - 1 < k$, i.e., $f \leq k$.

So (T_s, T_{s+1}) is a pair $(\leq i, i < \dots \leq k, \geq k, j)^-$, as expected. □

For future reference, we finally mention the right counterpart of Lemma 3.15. It can of course be proved by a direct argument, or deduced from the right counterpart of Lemma 3.16 (that we shall not need here).

Lemma 3.17. *Assume that the trees T, T' satisfy*

$$(3.15) \quad i \triangleright_T^* j, \quad i \triangleleft_T j - 1, \quad \text{and} \quad i \not\triangleright_{T'}^* j.$$

Then each sequence of rotations from T to T' contains a pair $(\dots, i, j - 1, \dots)^+$, or a pair $(i, \dots, j - 1, \dots)^-$.

3.6. Application: reproving a lower bound in $3n/2 + O(1)$. As a first application of the previous results and a warm-up for the sequel, we shall now reprove Proposition 2.23 about the distance of “bicombs”.

Proposition 3.18. *For $T = \langle 1^p 0^q \rangle$, $T' = \langle 0^q 1^p \rangle$ with $p, q \geq 1$, we have*

$$(3.16) \quad \text{dist}(T, T') \geq p + q + \min(p, q) - 2.$$




FIGURE 11. The bicombs of Propositions 2.23 and 3.18—here with $p = 4$ and $q = 6$.

The method of the proof consists in identifying various families of base pairs, and to prove, using the results of Sections 3.4 and 3.5, that every sequence of rotations from T to T' contains at least a certain number of pairs of these specific types.

Proof. Up to a symmetry, we may assume $p \leq q$. Put $n = p + q$. Let us say that a base pair is *special*

- of type I_a if it is $(\dots, a, p+1, \dots)^+$ or $(\dots, a, \dots, p+1)^-$ with $2 \leq a \leq p+1$,
- of type II_a if it is $(\dots, p+1, a, \dots)^+$ or $(p+1, \dots, a, \dots)^-$ with $p+2 \leq a \leq n$,
- of type III_a if it is $(1, \neq p+1, a, \dots)^+$ with $p+2 \leq a \leq n$,
- of type IV_a if it is $(\dots, a, \leq p, \dots)^+$ with $2 \leq a \leq p$.

First we observe that the various types of special pairs are disjoint, *i.e.*, a special pair has one type exactly: the four families are disjoint and, inside each family, the parameter a is uniquely determined.

Let (T_0, \dots, T_ℓ) be a path from T to T' in K_n . First choose a in $[2, p+1]$. We see on Figure 11 that $p+1$ covers $a-1$ in T' , but not in T , and, moreover, that a co-covers $p+1$ in T . Applying Lemma 3.15 with $i = a-1$ and $j = p+1$ guarantees that (T_0, \dots, T_ℓ) contains a pair $(\dots, a, p+1, \dots)^+$ or $(\dots, a, \dots, p+1)^-$, *i.e.*, a special pair of type I_a . Letting a vary from 2 to $p+1$ guarantees that (T_0, \dots, T_ℓ) contains at least p special pairs of type I.

Similarly, choose b in $[p+2, n]$. We see now that $p+1$ co-covers b in T , but not in T' , and, that, moreover, $p+1$ is not covered by $b+1$ in T' . Applying Lemma 3.17 with $i = p+1$ and $j = b+1$ guarantees that (T_0, \dots, T_ℓ) contains a pair $(\dots, p+1, b, \dots)^+$ or $(p+1, \dots, b, \dots)^-$, *i.e.*, a special pair of type II_b . Letting b vary from $p+2$ to n guarantees that (T_0, \dots, T_ℓ) contains at least $q-1$ special pairs of type II.

Consider b in $[p+2, n]$ again. Then 1 is covered by $p+1$ in T , and not covered by $p+1$ in T' . Applying Lemma 3.13 with $i = 1$ and $j = b$ guarantees that (T_0, \dots, T_ℓ) contains at least one pair $(1, a, b, \dots)^+$ for some a satisfying $2 \leq a \leq b$. Here two cases are possible.

Case 1. For each b in $[p+2, n]$, there is a pair $(1, a, b, \dots)^+$ with $a \neq p+1$, or type III_a . In this case, letting b vary from $p+2$ to n guarantees that (T_0, \dots, T_ℓ) contains at least $q-1$ special pairs of type III.

Case 2. There exists b such that (T_0, \dots, T_ℓ) contains no special pair of type III_b . Owing to the above observation, this implies that there exists r such that (T_r, T_{r+1}) is $(1, p+1, b, \dots)^+$. By Lemma 3.10(ii), this implies that 1 is covered by p in T_r . In this case, we claim that there must exist in (T_0, \dots, T_r) a pair of type IV_a for each a in $[2, p]$. Indeed, consider such an a . The hypothesis that 1 is covered by p in T_r implies that $a-1$ too is covered by p in T_r . On the other hand, we see that a co-covers $p+1$ in T . Applying Lemma 3.14 with $i = a$ and $j = p+1$ guarantees that (T_0, \dots, T_r) contains a pair $(\dots, a, \leq p, > p)^+$, hence a special pair of type IV_a . So, in this case, (T_0, \dots, T_ℓ) contains at least $p-1$ special pairs of type IV.

Summarizing, we conclude that (T_0, \dots, T_ℓ) contains at least

- p special pairs of type I,
 - $q-1$ special pairs of type II,
- and $q-1$ special pairs of type III or $p-1$ special pairs of type IV,

hence at least $3p-2$ special pairs. As these pairs are pairwise distinct, the distance between T_0 and T_ℓ , *i.e.*, between T and T' , is at least $3p-2$. \square

The previous argument is illustrated in the case of size 4 trees in Figure 12.




FIGURE 12. Proof of Proposition 3.18. Introducing special pairs of various types amounts to colouring the edges of the associahedron K_4 . In the current case, we use four colours, namely I_2 , I_3 , II_4 , and a common colour for III_4 and IV_2 , plus a neutral colour for the edges that receive none of the previous colours, *i.e.*, for non-special pairs. The proof shows that each path from $\langle 1100 \rangle$ to $\langle 0011 \rangle$ must contain at least one edge of each of the four colours—as can be checked on the picture—hence the distance between the trees $\langle 1100 \rangle$ and $\langle 0011 \rangle$ (framed squares) is four.

4. COLLAPSING

The previous method is powerful, but the results obtained so far remain limited. In order to establish stronger results, we now add one more ingredient called collapsing, which is a certain way of projecting an associahedron K_n onto smaller associahedra $K_{n'}$ with $n' < n$. The idea is very simple: collapsing a set of labels I in a tree T means erasing all leaves whose labels belong to I , and contracting the remaining edges to obtain a well-formed tree $\text{coll}_I(T)$. Collapsing a set of labels I in the two entries of a base pair (T, T') yields either a base pair, or twice the same tree, in which case we naturally say that the pair (T, T') is I -collapsing. This implies that the rotation distance between T and T' is at least the distance between the collapsed trees $\text{coll}_I(T)$ and $\text{coll}_I(T')$, plus the minimal number of inevitable I -collapsing pairs between T and T' . This principle will enable us to inductively determine the distances $\text{dist}(T_p, T'_p)$ for trees T_p, T'_p such that T_{p-1} is obtained from T_p and T'_{p-1} is obtained from T'_p by collapsing some set of labels (the same for both).

4.1. Collapsing. Assume that T is a finite binary tree. For $I \subseteq \mathbb{N}$, we consider the tree obtained from T by removing all leaves whose labels lie in I (if any). In order to include the case when all leaves are removed, we introduce an empty tree denoted \emptyset , together with the rules

$$(4.1) \quad \emptyset \wedge T = T, \quad T \wedge \emptyset = T.$$

It is then coherent to declare $|\emptyset| = -1$. Objects that are either a finite labeled tree or the empty tree will be called *extended trees*. We denote by $\text{Lab}(T)$ the family of labels occurring in T .

Definition 4.1. For $I \subseteq \mathbb{N}$ and T an extended tree, the I -collapse of T , denoted $\text{coll}_I(T)$, is recursively defined by

$$(4.2) \quad \text{coll}_I(T) = \begin{cases} \emptyset & \text{for } T = \emptyset, \\ \emptyset & \text{for } |T| = 0 \text{ and } \text{Lab}(T) \subseteq I, \\ T & \text{for } |T| = 0 \text{ and } \text{Lab}(T) \not\subseteq I, \\ \text{coll}_I(T_1) \wedge \text{coll}_I(T_2) & \text{for } T = T_1 \wedge T_2. \end{cases}$$

Example 4.2. Assume $T = ((1 \wedge 2) \wedge (3 \wedge 4))$. Then we find for instance

$$\text{coll}_{\{1\}}(T) = 2 \wedge (3 \wedge 4), \quad \text{coll}_{\{2,3\}}(T) = 1 \wedge 4, \quad \text{coll}_{\{1,2,3,4\}}(T) = \emptyset.$$

Properties of collapsing are mostly obvious. In particular, it should be clear that we always have $|\text{coll}_I(T)| = |T| - \#(\text{Lab}(T) \cap I)$. Also, we have the following compatibility with the covering and co-covering relations.

Lemma 4.3. *If i is covered by (resp. co-covers) j in T , and i, j do not belong to I , then i is covered by (resp. co-covers) j in $\text{coll}_I(T)$.*

Remark 4.4. We took the option not to change the remaining labels when some labels are collapsed. So, even we start with a tree T in which the labels are the default ones, namely 1 to $|T| + 1$, after collapsing we are likely to obtain a tree \overline{T} in which the labels are not 1 to $|\overline{T}| + 1$. That is why it seems preferable to consider general labeled trees, *i.e.*, to allow jumps in the sequence of labels.

4.2. Collapsing a base pair. For our current purpose, the question is to connect the rotation distance between two trees and the rotation distance between their images under collapsing. The point is that the collapsing of a base pair is either a diagonal pair, *i.e.*, a pair consisting of twice the same tree, or it is still a base pair, whose name is easily connected with that of the initial pair.

Lemma 4.5. *For $I \subseteq \mathbb{N}$ and (T, T') a base pair of name $(a, b, c, d)^+$,*

(i) *either we have*

$$(4.3) \quad [a, b - 1] \subseteq I \quad \text{or} \quad [b, c] \subseteq I \quad \text{or} \quad [c + 1, d] \subseteq I,$$

and then $\text{coll}_I(T)$ and $\text{coll}_I(T')$ coincide,

(ii) *or $(\text{coll}_I(T), \text{coll}_I(T'))$ is a base pair of name $(\bar{a}, \bar{b}, \bar{c}, \bar{d})^+$ with $\bar{a} = \min([a, b - 1] \setminus I)$, $\bar{b} = \min([b, c] \setminus I)$, $\bar{c} = \max([b, c] \setminus I)$, $\bar{d} = \max([c + 1, d] \setminus I)$.*

Proof. See Figure 7. If one of the three subtrees $T_{(\alpha 0)}$, $T_{(\alpha 10)}$, $T_{(\alpha 11)}$ completely vanishes, which happens when at least one of the three inclusions of (4.3) is true, then $\text{coll}_I(T)$ and $\text{coll}_I(T')$ coincide. Otherwise, *i.e.*, if at least one leaf of each of the above three subtrees remains, we have

$$\begin{aligned} \text{coll}_I(T_{(\alpha)}) &= \text{coll}_I(T_{(\alpha 0)}) \wedge (\text{coll}_I(T_{(\alpha 10)}) \wedge \text{coll}_I(T_{(\alpha 11)})), \\ \text{coll}_I(T'_{(\alpha)}) &= (\text{coll}_I(T_{(\alpha 0)}) \wedge \text{coll}_I(T_{(\alpha 10)})) \wedge \text{coll}_I(T_{(\alpha 11)}), \end{aligned}$$

Hence $(\text{coll}_I(T_{(\alpha)}), \text{coll}_I(T'_{(\alpha)}))$ is a positive base pair, and so is $(\text{coll}_I(T), \text{coll}_I(T'))$. The name should then be clear from the picture \square

Definition 4.6. A base pair is called I -collapsing if at least one of the three conditions of (4.3) is satisfied. If T, T' are trees, the I -distance between T and T' , denoted $I\text{-dist}(T, T')$, is defined to be the minimal number of I -collapsing steps occurring in a sequence of rotations from T to T' .


A direct application of Lemma 4.5 is the following useful relation:

Lemma 4.7. *For all trees T, T' and all sets I , we have*

$$(4.4) \quad \text{dist}(T, T') \geq \text{dist}(\text{coll}_I(T), \text{coll}_I(T')) + I\text{-dist}(T, T').$$

Proof. Let (T_0, \dots, T_ℓ) be a path from T to T' in $K_{|T|}$. By Lemma 4.5, the sequence $(\text{coll}_I(T_0), \dots, \text{coll}_I(T_\ell))$ is a path from $\text{coll}_I(T)$ to $\text{coll}_I(T')$ in $K_{\overline{n}}$, and the number of nontrivial pairs in this path is the number of non- I -collapsing pairs in (T_0, \dots, T_ℓ) . Therefore, we have $\ell \geq \text{dist}(\text{coll}_I(T), \text{coll}_I(T')) + I\text{-dist}(T, T')$. \square

Remark 4.8. By Lemma 4.5, the inequality (4.4) is an equality for $\text{dist}(T, T') \leq 1$. This need not be true in general. For instance, let $T = \langle 1100 \rangle$ and $T' = \langle 0011 \rangle$. We saw in Figure 12 that the distance between T and T' is 4. Now, we have $\text{coll}_{\{4,5\}}(T) = \text{coll}_{\{4,5\}}(T') = \langle 11 \rangle$, hence $\text{dist}(\text{coll}_{\{4,5\}}(T), \text{coll}_{\{4,5\}}(T')) = 0$. On the other hand, it can be checked on Figure 12 that there exists a path from T to T' in K_4 that contains only two $\{4, 5\}$ -collapsing pairs, namely



so we have $\{4, 5\}\text{-dist}(T, T') \leq 2$ (actually $= 2$).

4.3. Double collapsing. Technically, it will be convenient to use two collapsings at a time, with respect to sets that are strongly disjoint in the following sense.

Lemma 4.9. *Assume that I, J satisfy the condition*

$$(4.5) \quad \forall i \in I \forall j \in J \ ([i, j] \not\subseteq I \cup J).$$

Then, for all trees T, T' , we have

$$(4.6) \quad \text{dist}(T, T') \geq \text{dist}(\text{coll}_I(T), \text{coll}_I(T')) + I\text{-dist}(\text{coll}_J(T), \text{coll}_J(T')).$$

Proof. We claim that the inequality

$$(4.7) \quad I\text{-dist}(T, T') \geq I\text{-dist}(\text{coll}_J(T), \text{coll}_J(T'))$$

holds for all trees T, T' . Indeed, let (T_0, \dots, T_ℓ) be a path from T to T' . Then $(\text{coll}_J(T_0), \dots, \text{coll}_J(T_\ell))$ is a (possibly redundant) path from $\text{coll}_J(T)$ to $\text{coll}_J(T')$. Put $\ell = I\text{-dist}(\text{coll}_J(T), \text{coll}_J(T'))$. By definition of $I\text{-dist}$, there must be at least ℓ pairs (T_r, T_{r+1}) satisfying

$$(4.8) \quad \text{coll}_J(T_r) \neq \text{coll}_J(T_{r+1}) \quad \text{and} \quad \text{coll}_I(\text{coll}_J(T_r)) = \text{coll}_I(\text{coll}_J(T_{r+1})).$$

As $\text{coll}_I(\text{coll}_J(-)) = \text{coll}_{I \cup J}(-)$ always holds, (4.8) means that (T_r, T_{r+1}) is not J -collapsing, and is $(I \cup J)$ -collapsing. Now Condition (4.9) implies that every interval that is included in $I \cup J$ is included in I , or is included in J . Owing to the criterion of Lemma 4.5, we deduce that (T_r, T_{r+1}) is I -collapsing, and, therefore, we have $I\text{-dist}(T, T') \geq \ell$.

Then, using (4.7), we obtain

$$\begin{aligned} \text{dist}(T, T') &\geq \text{dist}(\text{coll}_I(T), \text{coll}_I(T')) + I\text{-dist}(T, T') \\ &\geq \text{dist}(\text{coll}_I(T), \text{coll}_I(T')) + I\text{-dist}(\text{coll}_J(T), \text{coll}_J(T')), \end{aligned}$$

which is the expected inequality (4.6). \square

4.4. **Application: a lower bound in $5n/3 + O(1)$.** Lemma 4.7 provides a natural method for establishing a lower bound on the distance $\text{dist}(T, T')$ in an inductive way: if $\text{dist}(\bar{T}, \bar{T}') \geq \bar{\ell}$ is known, and we can find a set I satisfying $\text{coll}_I(T) = \bar{T}$ and $\text{coll}_I(T') = \bar{T}'$, then it suffices to show that the minimal number of I -collapsing steps from T to T' is at least k to deduce $\text{dist}(T, T') \geq \bar{\ell} + k$. We shall now apply this principle to deduce from the results of Section 3.6, which provide a family with distance $3n/2 + O(1)$, a new family achieving distance $5n/3 + O(1)$.

Proposition 4.10. *For $T = \langle 1^p 0^p 1^p \rangle$, $T' = \langle 0^p (10)^p \rangle$ with $p \geq 1$, we have*

$$(4.9) \quad \text{dist}(T, T') \geq 5p - 4.$$

As T and T' above have size $3p$, we deduce

Corollary 4.11. *For $n = 3 \pmod 3$, we have $d(n) \geq \frac{5}{3}n - 4$.*

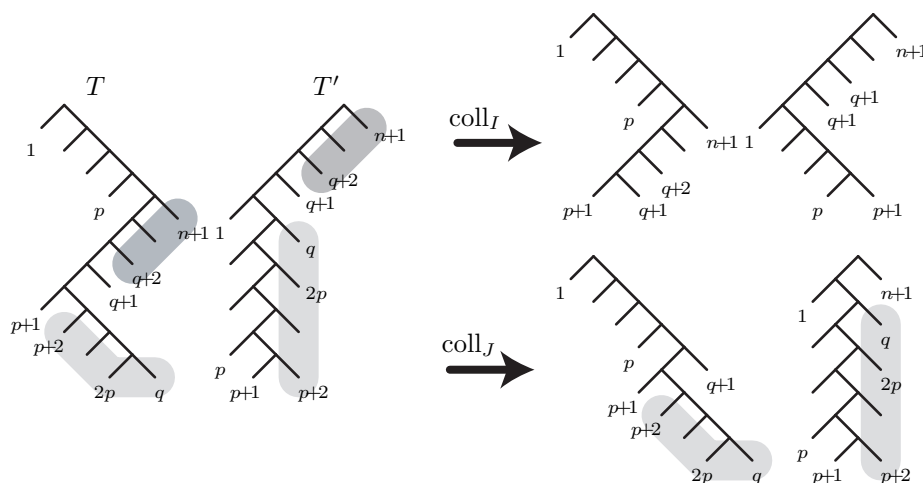


FIGURE 13. The trees of Proposition 4.10—here with $p = 4$ —and the proof of the latter from Lemma 4.12: collapsing I (light grey labels) leads to 2-combs, whose distance is known; collapsing J (dark grey labels) leads to trees that are, up to shifting the labels, those of Lemma 4.12.

To prove Proposition 4.10, we shall use Proposition 2.23 (or 3.18) and a convenient collapsing. Fix some p , and let \bar{T} and \bar{T}' be the corresponding trees of Proposition 2.23, namely $\langle 1^p 0^p \rangle$ and $\langle 0^p 1^p \rangle$. Collapsing T into \bar{T} is easy: T is a zigzag of three alternating length p combs (“tricomb”), whereas \bar{T} is a zigzag of two alternating combs (“bicomb”), so that we can project T to \bar{T} by collapsing all labels from $p + 2$ to $2p + 1$. It then turns out that collapsing the same labels in T' leads to \bar{T}' . By Proposition 2.23, the distance of \bar{T} and \bar{T}' is $3p - 2$, so, owing to Lemma 4.7, in order to establish Proposition 4.10, it is enough to prove

$$(4.10) \quad [p + 2, 2p + 1]\text{-dist}(T, T') \geq 2p - 2,$$

i.e., to prove that each sequence of rotations from T to T' contains at least $p - 2$ pairs that are $[p + 2, 2p + 1]$ -collapsing.

Instead of working with the trees of Proposition 4.10 themselves, it will be more convenient to use a second, auxiliary collapsing, and to use Lemma 4.9. We shall prove:

Lemma 4.12. For $\bar{T} = \langle 1^p 0 1^p \rangle$, $\bar{T}' = \langle (01)^p 0 \rangle$ with $p \geq 1$, we have

$$(4.11) \quad [p+2, 2p+1]\text{-dist}(\bar{T}, \bar{T}') \geq 2p-2.$$

Proof of Proposition 4.10 from Lemma 4.12. Let $I = [p+2, 2p+1]$ and $J = [2p+3, 3p+1]$. Then the sets I and J satisfy Condition (4.6), and we have $\bar{T} = \text{coll}_J(T)$ and $\bar{T}' = \text{coll}_J(T')$. Lemma 4.9 then gives

$$\text{dist}(T, T') \geq \text{dist}(\text{coll}_I(T), \text{coll}_I(T')) + I\text{-dist}(\bar{T}, \bar{T}').$$

As can be checked on Figure 13, we have $\text{coll}_I(T) = \langle 1^p 0^p \rangle$ and $\langle 0^p 1^p \rangle$. Using Proposition 2.23 and (5.3), we deduce

$$\text{dist}(T, T') \geq (3p-2) + (2p-2) = 5p-4,$$

as expected. \square

Proof of Lemma 4.12. The argument is similar to the one used for Proposition 3.18, with the additional difficulty that we need pairs that are $[p+2, 2p+1]$ -collapsing. Put $\bar{n} = 2p+1$, $I = [p+2, 2p]$, and, for $p+2 \leq a \leq 2p$, say that a base pair is

- special of type I_a if it is $(\dots, a, a, \dots)^+$
- special of type II_a^+ if it is $(\dots, \leq p+1, a, \leq \bar{n})^+$
- special of type II_a^- if it is $(\leq p, \dots, \geq p+1, a)^-$
- special of type III_a if it is $(\dots, a+1, \dots, \bar{n}+1)^-$.

It is straightforward that a special pair has a unique type, and that all special pairs are I -collapsing: with obvious notation, we have $[\nu_2, \nu_3] \subseteq I$ for type I_a , and $[\nu_3+1, \nu_4] \subseteq I$ for types II_a^+ , II_a^- , and III_a , which is enough to conclude using the criterion of Lemma 4.5.

Let (T_0, \dots, T_ℓ) be a path from \bar{T} to \bar{T}' in $K_{\bar{n}}$. Choose a in $[p+2, 2p]$. First, as can be read on Figure 13, we have

$$a-1 \not\prec_{\bar{T}} a \quad \text{and} \quad a-1 \prec_{\bar{T}'} a.$$

Lemma 3.13 guarantees that (T_0, \dots, T_ℓ) contains a pair $(\leq a-1, > a-1, a, \dots)^+$, *i.e.*, a pair of type I_a .

Next, let $a' = \bar{n}+1-a$. Then we have $2 \leq a' \leq p$, and we read the relations

$$a' \not\prec_{\bar{T}} a, \quad a' \prec_{\bar{T}} a, \quad \text{and} \quad a'+1 \triangleright_{\bar{T}'}^* a.$$

Applying Lemma 3.16 with $i = a'$, $j = a$, and $k = p+1$ guarantees the existence of r such that (T_r, T_{r+1}) is

$$(\leq a', a' < \dots < \leq p+1, a, \dots)^+ \quad \text{or} \quad (\leq a', a' < \dots < \leq p+1, \geq p+1, a)^-.$$

In the latter case, we have a special pair of type II_a^- . In the former case, we have a special pair of type II_a^+ provided the last parameter, namely $\nu_4(T_r, T_{r+1})$, is at most \bar{n} . Now assume this is not the case, *i.e.*, we have $\nu_4(T_r, T_{r+1}) = \bar{n}+1$. By Lemma 3.10(*iv*), we have $a+1 \sqsubseteq_{T_r}^* \bar{n}+1$, hence $a+1 \triangleright_{T_r}^* \bar{n}+1$ as $a+1 \leq \bar{n}$ holds by hypothesis. But, by hypothesis, we have $a+1 \not\triangleright_{T_r}^* \bar{n}+1$. By Lemma 3.14, there must exist $s \leq r$ such that (T_s, T_{s+1}) is $(\dots, a+1, \dots, \bar{n}+1)^-$, hence special of type III_a .

Hence, for each of the $p-1$ values $p+2, \dots, 2p$, the path (T_0, \dots, T_ℓ) contains a pair of type I_a , and a pair of type II_a^+ , II_a^- , or III_a . Hence (T_0, \dots, T_ℓ) contains

$p - 1$ pairs of type I,
 $p - 1$ pairs of type II $^\pm$ or III,

hence at least $2p - 2$ pairs that are I -collapsing. □

Remark 4.13. The previous result is optimal. It is not difficult to construct an explicit path of length $5p - 4$ from T to T' , and to check that its projection is a path of length $2p - 2$ from \bar{T} to \bar{T}' , all of which steps are I -collapsing. Also, one can observe that, in the above proof, the final argument showing the existence of a special pair of type III $_a$ also shows the existence of a pair $(\dots, a + 1, \dots, \bar{n} + 1)^+$. It follows that each path visiting a vertex satisfying $a \triangleright^* \bar{n} + 1$ for some $a \geq p + 3$ contains at least $2p - 1$ pairs that are I -collapsing. Hence such a path has length at least $5p - 3$ and it is not geodesic: moving $\bar{n} + 1$ to the right is never optimal.

4.5. Alternative families. To conclude this section, we mention still another family witnessing a growth rate of the form $3n/2 + O(1)$. The analysis of this family, which heavily uses collapsing, is easier than that of Proposition 3.18, but, contrary to the arguments developed above, it does not seem to extend for proving stronger results. The trees we consider are zigzags, with a slight change near the ends—though seemingly minor, that change modifies distances completely.

Proposition 4.14. *For $m \geq 0$, we have*

$$(4.12) \quad \text{dist}(\langle 1(10)^m 0 \rangle, \langle 0(01)^m 1 \rangle) = 3m + 1.$$




FIGURE 14. The trees of Proposition 4.14—here with $m = 5$: collapsing I (light grey labels) from T_m and T'_m leads (up to a relabeling) to T_{m-2} and T'_{m-2} , while collapsing J (dark grey labels) leads (up to a relabeling) to T_2 and T'_2 .

Proof. Let $T_m = \langle 1(10)^m 0 \rangle$ and $T'_m = \langle 0(01)^m 1 \rangle$. We use induction on m . For $m \leq 1$, the result is easily checked by a direct computation. Assume $m \geq 2$. Put

$$I_m = \{m, m + 1, m + 3, m + 4\} \quad \text{and} \quad J_m = [1, m - 2] \cup [m + 6, 2m + 3].$$

As can be read on Figure 14, we have

$$\text{coll}_{I_m}(T_m) = T_{m-2} \quad \text{and} \quad \text{coll}_{I_m}(T'_m) = T'_{m-2}.$$

as well as

$$\text{coll}_{J_m}(T_m) = T_2 \quad \text{and} \quad \text{coll}_{J_m}(T'_m) = T'_2.$$

The sets I_m and J_m satisfy the disjointness condition (4.5). Applying Lemma 4.7, we deduce

$$\text{dist}(T_m, T'_m) \geq \text{dist}(T_{m-2}, T'_{m-2}) + I_2\text{-dist}(T_2, T'_2).$$

A brute force verification—or a proof using the techniques of Section 3.4—gives

$$I_2\text{-dist}(T_2, T'_2) = 6,$$

and $\text{dist}(T_m, T'_m) \geq 3m + 1$ follows inductively. The other inequality, whence (4.12), are easily checked by a direct computation. \square

The remarkably simple proof of Proposition 4.14 relies on the conjunction of two properties. First, the two families of trees we consider are stable under two types of collapsing simultaneously: whether we collapse from the top or from the bottom, we can manage to remain in the same family. Second, the involved collapsing are perfect, in the sense that the inequality (4.4) turns out to be an equality, a necessary condition if we are to obtain an exact value.

It is not difficult to obtain other families satisfying one of the above two properties. In particular, for each $p \geq 1$, the thin trees $\langle 1^p(10)^m 0^p \rangle$ and $\langle 0^p(01)^m 1^p \rangle$ are eligible. Proposition 4.14 then corresponds to $p = 1$. The choice $p = 0$ is uninteresting, but $p \geq 2$ gives seemingly large distances. Applying the scheme above leads to looking for the unique parameter

$$([p + 1, 3p + 3] \setminus \{2p + 2\})\text{-dist}(\langle 1^p(10)^{p+1} 0^p \rangle, \langle 0^p(01)^{p+1} 1^p \rangle),$$

which we have seen is 6 for $p = 1$. For $p = 2$, one obtains the value 10, but it is then easy to see that such a value cannot give an equality in (4.4). In this way, one obtains lower bounds for the distances of the involved trees, but these bounds are not sharp, and it seems hard to obtain very strong results. For instance, for $p = 2$, collapsing six leaves guarantees ten collapsing steps, and one cannot obtain more than $5n/3 + O(1)$.

5. A LOWER BOUND IN $2n + O(\sqrt{n})$

It is not hard to repeat the argument of Section 4.4 so as to construct explicit trees of size n at distance $7n/4 + O(1)$. Unfortunately, a further iteration seems difficult and, in order to go farther, we shall have to develop a more intricate argument—yet the principle always remains the same. The main difference is that, now, we will collapse a size $2p$ bicomb rather than a size p comb.

5.1. Starting from an m -comb. The main result of this section is analogous to Proposition 4.10, but the source tree is a $2m$ -comb rather than a tricomb, *i.e.*, it is a tree obtained by stacking $2m$ left and right combs, alternately. The target tree is again a zigzag preceded by a short left comb—see Figure 15.

Proposition 5.1. *For $m, p \geq 1$, we have*

$$(5.1) \quad \text{dist}(\langle (1^p 0^p)^m \rangle, \langle 0^p (10)^{(m-1)p} 1^p \rangle) \geq 4mp - 3m - p + 1.$$

By letting m stay fixed and p vary, we obtain size n trees whose distance grows at least—actually, one can check that (5.1) is an equality—as $(4m - 1)n/(2m) + O(1)$. By letting m and p vary simultaneously, we obtain a lower bound in $2n + O(\sqrt{n})$:




FIGURE 15. The trees of Proposition 5.1—here with $m = 3$ and $p = 4$ —and the proof of the latter from Lemma 5.3: collapsing I (light grey labels) leads from $(T_{m,p}, T'_{m,p})$ to $(T_{m-1,p}, T'_{m-1,p})$; collapsing J (dark grey labels) leads from $(T_{m,p}, T'_{m,p})$ to the pair (\bar{T}_p, \bar{T}'_p) , which does not depend on m and is, up to shifting the labels, the pair of Lemma 5.3.

Corollary 5.2. For n of the form $2m^2$, we have

$$(5.2) \quad d(n) \geq 2n + 1 - 2\sqrt{2n}.$$

Moreover, $d(n) \geq 2n - C\sqrt{n}$ holds for each n with $C = \sqrt{70}$.

Proof (of Corollary 5.2 from Proposition 5.1). Assume first $n = 2m^2$. Then, choosing $p = m$ and using (5.1), we find

$$d(n) = d(2mp) \geq \text{dist}(\langle (1^p 0^p)^m \rangle, \langle 0^p (10)^{(m-1)p} 1^p \rangle) = 4m^2 - 4m + 1,$$

whence (5.2). For the general case, choose $m = \lfloor \sqrt{5n/14} \rfloor$, $p = \lfloor \sqrt{7n/10} \rfloor$. Then we have $2mp \leq n$, whence

$$d(n) \geq d(2mp) \geq \text{dist}(\langle (1^p 0^p)^m \rangle, \langle 0^p (10)^{(m-1)p} 1^p \rangle) = 4mp - 3m - p + 1.$$

by (5.1). Using $m \leq \sqrt{5n/14} < m + 1$ and $p \leq \sqrt{7n/10} < p + 1$, one obtains

$$4mp - 3m - p + 1 > 2n - (7\sqrt{5n/14} - 5\sqrt{7n/10}) + 5 = 2n - \sqrt{70n} + 5. \quad \square$$

The proof of Proposition 5.1 uses an induction on the parameter m , *i.e.*, on the number of alternations in the zigzag-tree $\langle (1^p 0^p)^m \rangle$. It is not hard to find how to collapse $T_{m,p}$ to $T_{m-1,p}$ and $T'_{m,p}$ to $T'_{m-1,p}$. Then the key point consists in identifying sufficiently many collapsing pairs, and this is done in the following result.

Lemma 5.3. *Let $T = \langle 1^p 0^1 p^0 \rangle$, $T' = \langle 0(10)^p 1^p \rangle$, and $I = [p + 2, 2p + 1]$. Then, for each $p \geq 1$, we have*

$$(5.3) \quad I\text{-dist}(T, T') \geq 4p - 3.$$

Proof of Proposition 5.1 from Lemma 5.3. The argument is exactly similar to the one used to deduce Proposition 4.10 from Lemma 4.12. We use induction on $m \geq 1$. The trees $T_{1,p}$ and $T'_{1,p}$ are the 2-combs $\langle 1^p 0^p \rangle$ and $\langle 0^p 1^p \rangle$, and Proposition 2.23 (or 3.18) gives $\text{dist}(T_{1,p}, T'_{1,p}) \geq 3p - 2$, in agreement with (5.1).

Assume now $m \geq 2$, and let

$$I = [(m-1)p + 2, (m+1)p + 1] \quad \text{and} \quad J = [1, (m-2)p] \cup [(m+1)p + 3, 2mp + 1].$$

Then I and J satisfy the disjointness condition of Lemma 4.9, and we read on Figure 15 that I -collapsing maps $T_{m,p}$ to $T_{m-1,p}$, and $T'_{m,p}$ to $T'_{m-1,p}$. On the other hand, J -collapsing maps $T_{m,p}$ to a tree \overline{T}_p that is, up to shifting the labels, the tree T of Lemma 5.3, and $T'_{m,p}$ to a tree \overline{T}'_p that is, up to shifting the labels, the tree T' of Lemma 5.3. Applying Lemma 4.9, the induction hypothesis for $m - 1$, and Lemma 5.3, we obtain

$$\begin{aligned} & \text{dist}(T_{m,p}, T'_{m,p}) \\ & \geq \text{dist}(\text{coll}_I(T_{m,p}), \text{coll}_I(T'_{m,p})) + I\text{-dist}(\text{coll}_J(T_{m,p}), \text{coll}_J(T'_{m,p})) \\ & = \text{dist}(T_{m-1,p}, T'_{m-1,p}) + I\text{-dist}(\overline{T}_p, \overline{T}'_p) \\ & \geq (4(m-1)p - 3(m-1) - p + 1) + (4p - 3) = 4mp - 3m - p + 1, \end{aligned}$$

which is (5.1). □

5.2. Special pairs. The proof of Lemma 5.3 will occupy the next sections. It is parallel to the proof of Lemma 4.12, but it is more involved and requires some care, mainly because we are to introduce many different types of special pairs.

In the sequel, we put $n = 3p + 1$ and $q = 2p + 1$. Then T and T' are size n trees, displayed in Figure 16—compared with the trees \overline{T}_p and \overline{T}'_p of Figure 15, the only difference is that the labeling has been standardized.




FIGURE 16. The trees of Lemma 5.3, here with $p = 4$, and, in grey, the labels of I , which are those for which we need to count the I -collapsing pairs.

Then we consider the following eleven (!) families of base pairs. We say that a base pair is

- special of type I_a^+ if it is $(\dots, a, q, \dots)^+$ with $p + 2 \leq a \leq q$,
- special of type I_a^- if it is $(\dots, a, \dots, q)^-$ with $p + 2 \leq a < q$,
- special of type II_a^+ if it is $(\dots, q, a, \dots)^+$ with $q + 1 \leq a < n$,
- special of type II_a^- if it is $(q, \dots, a, \dots)^-$ with $q + 1 \leq a < n$,
- special of type III_a^+ if it is $(\dots, a, < q, \dots)^+$ with $p + 2 \leq a < q$,
- special of type III_a^- if it is $(\dots, a, \dots, n + 1)^-$ with $p + 2 \leq a \leq q$,
- special of type IV_a^+ if it is $(\dots, \geq p + 2 \ \& \neq q, a, \dots)^+$ with $q + 1 < a \leq n$,
- special of type V_a^+ if it is $(\dots, \leq p + 1, a, \leq n)^+$ with $q + 1 \leq a < n$,
- special of type V_a^- if it is $(\dots, \leq p + 1, > p, a)^-$ with $q + 1 \leq a < n$,
- special of type VI_a^+ if it is $(\dots, a + 1, \dots, n + 1)^+$ with $q + 1 \leq a < n$,
- special of type VI_a^- if it is $(\dots, a + 1, \dots, n + 1)^-$ with $q + 1 \leq a < n$.

Claim 1. Every special pair is I -collapsing.

Proof. As we consider size n trees, every base pair satisfies $\nu_4 \leq n + 1$, hence $\nu_3 \leq n$ and, therefore $\nu_3 \in I$. For all special pairs except those of type V, we have $\nu_2 \geq p + 2$, hence $[\nu_2, \nu_3] \subseteq I$. For the special pairs of type V, we have $\nu_3 \geq p + 1$ and $\nu_4 \leq n$, hence $[\nu_3 + 1, \nu_4] \subseteq I$. In all cases, the criterion of Lemma 4.5 implies that the pair is I -collapsing. \square

Claim 2. A special pair has a unique type, except, for $q + 1 \leq b < a < n$,

- the pairs $(\dots, b + 1, a, n + 1)^+$, which have type both IV_a^+ and VI_b^+ ,
- the pairs $(q, b + 1, a, n + 1)^-$, which have type both II_a^- and VI_b^- .

Proof. A positive pair cannot coincide with a negative one, so we can consider positive and negative pairs separately. Next, for each type τ_a , the value of a can be recovered from one of the parameters ν_k of the pair, hence it is impossible that a pair of type τ_a be of type τ_b for some $b \neq a$. So it remains to check that, for each pair (τ_a, τ'_b) of distinct types (with the same sign), it is impossible that a special pair be simultaneously of type τ_a and τ'_b —up to the exceptions mentioned in the claim. This is done in the two arrays below. As should be clear, “ $\nu_2 > q / \nu_2 \leq q$ ” means that

a pair cannot be of the considered two types τ, τ' simultaneously because, for being of type τ , we must have $\nu_2 > q$ whereas, for being of type τ' , we must have $\nu_2 \leq q$.

\cap	II_b^+	III_b^+	IV_b^+	V_b^+	VI_b^+
I_a^+	$\nu_3=q/\nu_3>q$	$\nu_3=q/\nu_3<q$	$\nu_3=q/\nu_3>q$	$\nu_3=q/\nu_3>q$	$\nu_3=q/\nu_3>q$
II_a^+		$\nu_2=q/\nu_2<q$	$\nu_2=q/\nu_2\neq q$	$\nu_2=q/\nu_2\leq p+1$	$\nu_2=q/\nu_2>q$
III_a^+			$\nu_3<q/\nu_3>q$	$\nu_3<q/\nu_3>q$	$\nu_2<q/\nu_2>q$
IV_a^+				$\nu_2\geq p+2/\nu_2\leq p+1$	possible
V_a^+					$\nu_4\leq n/\nu_4>n$

\cap	II_b^-	III_b^-	V_b^-	VI_b^-
I_a^-	$\nu_2\leq q/\nu_2>q$	$\nu_4\leq n/\nu_4=n+1$	$\nu_2\geq p+2/\nu_2\leq p+1$	$\nu_4=q/\nu_4=n+1$
II_a^-		$\nu_1=q/\nu_1<q$	$\nu_1=q/\nu_1\leq p$	possible
III_a^-			$\nu_2\geq p+2/\nu_2\leq p+1$	$\nu_2\leq q/\nu_2>q$
V_a^-				$\nu_4\leq n/\nu_4>n$

So all cases have been considered. \square

We shall now exploit the differences of covering and co-covering between T and T' . First, we use the fact that $p+1$ to $q-1$ are not covered by q in T , whereas they are in T' , and, symmetrically, the fact that q co-covers $q+2$ to n in T , whereas it does not in T' .

Claim 3. *Every path from T to T'*

- contains a pair of type I_a^\pm for each a in $[p+2, q]$,
- and - contains a pair of type II_a^\pm for each a in $[q+1, n-1]$.

Proof. Assume $a \in [p+2, q]$. Then we have $p+1 \leq a-1 \leq 2p$, and, as can be read on Figure 16, we have

$$a-1 \not\triangleleft_T q, \quad a-1 \triangleleft_{T'} q, \quad \text{and} \quad a \triangleright_{T'}^* q.$$

Applying Lemma 3.15 with $i = a-1$ and $j = q$ shows that every path from T to T' contains a pair $(\dots, a, q, \dots)^+$, of type I_a^+ , or a pair $(\dots, a, \dots, q)^-$, of type I_a^- .

Assume now $a \in [q+1, n-1]$. Then we have $q+2 \leq a+1 \leq n$, and, as can be read on Figure 16, we have

$$q \triangleright_T^* a+1, \quad q \triangleleft_T a, \quad \text{and} \quad q \not\triangleright_{T'}^* a+1.$$

Applying Lemma 3.17 with $i = q$ and $j = a+1$ shows that every path from T to T' contains a pair $(\dots, q, a, \dots)^+$, of type II_a^+ , or a pair $(q, \dots, a, \dots)^-$, of type II_a^- . \square

Remark 5.4. The second argument above also applies to $a = q$, but then it possibly leads to a pair $(\dots, q, q, \dots)^+$ that would be of type I_q^+ and may have been already considered in the first argument.

We shall now exploit the fact that $p+1$ is not covered by any label from $q+1$ to $n-1$ in T , and is covered by these labels in T' .

Claim 4. *Every path from T to T'*

- contains a pair of type III_a^+ for each a in $[p+2, q-1]$,
- or - contains a pair of type IV_b for each b in $[q+1, n-1]$.

Proof. For $p = 1$, the result is vacuously true, and we assume $p \geq 2$. Let (T_0, \dots, T_ℓ) be any path from T to T' . Assume that, for some b in $[q + 1, n - 1]$, there exists no pair of type IV_b in (T_0, \dots, T_ℓ) . As we read on Figure 16 the relations

$$p + 1 \not\triangleleft_T b \quad \text{and} \quad p + 1 \triangleleft_{T'} b,$$

applying Lemma 3.13 with $i = p + 1$ and $j = b$ shows that there must exist r such that (T_r, T_{r+1}) is $(h, g, b, \dots)^+$ for some g, h satisfying $h \leq p + 1$ and $g \geq p + 2$. The hypothesis that this pair is *not* of type IV_b implies $g = q$. As $p \geq 2$ holds, we have $h \leq p + 1 < q - 1$ and, therefore, Lemma 3.10(i) implies $h \triangleleft_{T_r} q - 1$. Let a be any element of $[p + 2, q - 1]$. By Lemma 3.6, $h \triangleleft_{T_r} q - 1$ implies $a - 1 \triangleleft_{T_r} q - 1$. On the other hand, we see on Figure 16 that a co-covers q in T . Applying Lemma 3.17 with $i = a$ and $j = q$ shows that (T_0, \dots, T_r) contains a pair $(\dots, a, < q, \geq q)^+$, of type III_a^+ . \square

The next result uses the fact that each label a between 2 and p is not covered by $n + 1 - a$ in T , whereas it is in T' . The possible interference of the label $n + 1$ makes the result slightly more complicated—as was already the case in the proof of Lemma 4.12. This step is the most delicate one, as it requires the full power of Lemma 3.16 and not only Lemma 3.13 or Lemma 3.15.

Claim 5. *For each a in $[q + 1, n - 1]$, every path from T to T'*

- contains a pair of type V_a^\pm ,
- or - contains a pair of type VI_a^+ and a pair of type VI_a^- .

Proof. Let (T_0, \dots, T_ℓ) be a path from T to T' , and let a be an element of $[q + 1, n - 1]$. Put $a' = n + 1 - a$. Then we have $2 \leq a' \leq p$, and we read on Figure 16 the relations

$$a' \not\triangleleft_T a, \quad a' \triangleleft_{T'} a, \quad \text{and} \quad a' + 1 \triangleright_{T'}^* a.$$

Applying Lemma 3.16 with $i = a'$, $j = a$, and $k = p + 1$ shows that there exists r such that (T_r, T_{r+1}) is $(\leq a', a' < \dots \leq p + 1, a, \dots)^+$ or $(\leq a', a' < \dots \leq p + 1, > p, a)^-$. The latter pair is of type V_a^- . As for the former one, two cases are possible: if $\nu_4(T_r, T_{r+1}) \leq n$ holds, then (T_r, T_{r+1}) is of type V_a^+ , else we necessarily have $\nu_4(T_r, T_{r+1}) = n$. Then (T_r, T_{r+1}) is a pair $(\dots, \dots, a, n + 1)^+$, in which case, by Lemma 3.10(iv), we have $a + 1 \triangleright_{T_r}^* n + 1$, and even $a + 1 \triangleright_{T_r}^* n + 1$ as $a < n$ is assumed. Now, we read on Figure 16 that $a + 1$ co-covers $n + 1$ neither in T nor in T' . Applying Lemma 3.14 with $i = a + 1$ and $j = n + 1$ guarantees that (T_0, \dots, T_r) contains at least a pair $(\dots, a + 1, \dots, n + 1)^-$, of type VI_a^- , and that (T_r, \dots, T_ℓ) contains at least a pair $(\dots, a + 1, \dots, n + 1)^+$, of type VI_a^+ . \square

The last claim of the series will be used to cope with the possible interference between types II^- and VI^- .

Claim 6. *Assume that T_* is a size n tree and q co-covers $n + 1$ in T_* . Then every path from T to T_* contains*

- a pair of type III_a^\pm for each a in $[p + 2, q]$.

Proof. Let $\overline{T} = \text{coll}_{[q+1, n]}(T)$ and $\overline{T}_* = \text{coll}_{[q+1, n]}(T_*)$. Then we have $\overline{T} = \langle 01^p \rangle$, and, by Lemma 4.3, the hypothesis that q co-covers $n + 1$ in T_* implies that q co-covers $n + 1$ in \overline{T}_* . Let a be any element of $[p + 2, q]$. By Lemma 3.8, $a - 1$ cannot be covered by n in \overline{T}_* . On the other hand, $a - 1$ is covered by n in T , and, moreover, a co-covers q in T . Applying Lemma 3.15 to $(\overline{T}_*, \overline{T})$ with $i = a$ and $j = q$ shows that every path from \overline{T} to \overline{T}_* contains a pair $(\dots, a, q, n + 1)^-$

or $(\dots, a, \dots, q)^+$. It follows that every path from T to T_* contains a pair that projects to a pair of the previous form when $[q+1, n]$ is collapsed.

Now, Lemma 4.5 shows that $(a', b', c', d')^-$ projects to $(\dots, a, q, n+1)^-$ if and only if we have $b' = a$, $\sup([b', c'] \setminus [q+1, n]) = q$, hence $c' \geq q$, and $\sup([c'+1, d'] \setminus [q+1, n]) = n+1$, hence $d' = n+1$, implying that $(a', b', c', d')^-$ is special of type III_a^- .

Similarly, $(a', b', c', d')^+$ projects to $(\dots, a, \dots, q)^+$ for $b' = a$, $c' = \sup([b', c'] \setminus [q+1, n]) < q$, hence $c' < q$, and $\sup([c'+1, d'] \setminus [q+1, n]) = q$, hence $d' \leq n$, so that $(a', b', c', d')^+$ is special of type III_a^+ . \square

5.3. Proof of Lemma 5.3. We are now ready to prove Lemma 5.3. The argument is similar to the one used for proving Lemma 4.12, but we have to be more careful because of the possible interferences between special pairs of type VI^\pm and special pairs of other types. It may be noted that such problems never occur when $\text{add}(n)$ remains $1^p 01^{p+1}$ throughout the considered path from T to T' : proving the result for the trees $\text{coll}_{\{n+1\}}(T)$ and $\text{coll}_{\{n+1\}}(T')$ would be easier.

Proof of Lemma 5.3. Let (T_0, \dots, T_ℓ) be any path from T to T' . We have to show that this path contains at least $4p-3$ special pairs. The latter can correspond to several combinations of types, and we consider three cases.

Case 1: (T_0, \dots, T_ℓ) contains a pair of type III_a^+ for each a in $[p+2, q-1]$.

Then (T_0, \dots, T_ℓ) contains at least

p	pairs of type I^\pm	by Claim 3,
$p-1$	pairs of type II^\pm	by Claim 3,
$p-1$	pairs of type III^+	by hypothesis,
$p-1$	pairs of type V^\pm or VI^+	by Claim 5,

hence at least $4p-3$ special pairs, which are I -collapsing by Claim 1, and are pairwise distinct by Claim 2, since we appeal to no pair of type IV^+ (which could interfere with VI^+) or VI^- (which could interfere with II^-).

Case 2: There exists a in $[p+2, q-1]$ such that (T_0, \dots, T_ℓ) contains no pair of type III_a^+ , and there is no r such that q co-covers $n+1$ in T_r .

Then Claim 2 implies that (T_0, \dots, T_ℓ) contains no special pair that is simultaneously of types II^- and VI^- . Indeed, if a pair is both of types II_a^- and VI_b^- , it has the form $(q, \dots, \dots, n+1)^-$, and, therefore, by Lemma 3.10(i), q co-covers $n+1$ in the two trees of that pair. Then (T_0, \dots, T_ℓ) contains

p	pairs of type I^\pm	by Claim 3,
$p-1$	pairs of type II^\pm	by Claim 3,
$p-1$	pairs of type IV^+	by Claim 4 and the hypothesis,
$p-1$	pairs of type V^\pm or VI^-	by Claim 5,

and we have again $4p-3$ special pairs, which are I -collapsing by Claim 1, and are pairwise distinct, since we appeal to no pair of type VI^+ (which could interfere with IV^+), and interferences between type II^- and VI^- are discarded by our hypotheses.

Case 3: There exists a in $[p+2, q-1]$ such that (T_0, \dots, T_ℓ) contains no pair of type III_a^+ , and there exists r such that q co-covers $n+1$ in T_r .

Then (T_0, \dots, T_ℓ) contains

p	pairs of type I^\pm	by Claim 3,
$p - 1$	pairs of type II^\pm	by Claim 3,
$p - 1$	pairs of type III^\pm	by Claim 6,
$p - 1$	pairs of type IV^+	by Claim 5 and the hypothesis.

We still have found $4p - 3$ special pairs, all I -collapsing by Claim 1, and pairwise distinct by Claim 2, since we appeal now to no pair of type VI^\pm . \square

So the proofs of Lemma 5.3 and, therefore, of Proposition 5.1, are complete, yielding the expected lower bound $d(n) \geq 2n - O(\sqrt{n})$ on the diameter of the n th associahedron.

5.4. Going further. Proving the conjectured value $d(n) = 2n - 6$ for $n > 10$ using the above methods seems feasible, but is likely to require a more intricate argument. Experiments easily suggest families of trees that should achieve the maximal distance, namely symmetric zigzag-trees with small combs attached at each end, on the shape of the trees $T'_{m,3}$ of Proposition 5.1.

Conjecture 5.5. *Define*

$$T_n = \begin{cases} \langle 111(01)^{p-3}00 \rangle \\ \langle 111(01)^{p-3}000 \rangle \end{cases} \quad T'_n = \begin{cases} \langle 000(10)^{p-3}11 \rangle & \text{for } n = 2p - 1, \\ \langle 000(10)^{p-3}111 \rangle & \text{for } n = 2p, \end{cases}$$

Then one has $\text{dist}(T_n, T'_n) = 2n - 6$ for $n \geq 11$.

The problem for establishing Conjecture 5.5 is that counting I -collapsing pairs cannot suffice here: various solutions exist for projecting (T_n, T'_n) onto (T_{n-2}, T'_{n-2}) by collapsing two labels, but the minimal number of collapsing pairs is then 3, and not 4, as would be needed to conclude. On the other hand, the highly symmetric shape of the trees T_n and T'_n allows for other arguments that will not be developed here.

ACKNOWLEDGMENT

The author thanks Shalom Eliahou for introducing him to the fascinating problem addressed in this paper.

REFERENCES

- [1] J. Belk & K. Brown, *Forest diagrams for elements of Thompson's group F* , Internat. J. Algebra Comput. **15** (2005) 815–850.
- [2] P. Bose & F. Hurtado, *Flips in planar graphs*, Computational Geometry **42-1** (2009) 60–80.
- [3] J.W. Cannon, W.J. Floyd, & W.R. Parry, *Introductory notes on Richard Thompson's groups*, Ens. Math. **42** (1996) 215–257.
- [4] P. Dehornoy, *The structure group for the associativity identity*, J. Pure Appl. Algebra **111** (1996) 59–82.
- [5] P. Dehornoy, *Geometric presentations of Thompson's groups*, J. Pure Appl. Algebra **203** (2005) 1–44.
- [6] D. Epstein, J. Cannon, D. Holt, S. Levy, M. Paterson & W. Thurston, *Word Processing in Groups*, Jones & Bartlett Publ. (1992).
- [7] B. Fordham, *Minimal length elements of Thompson's group F* , Geom. Dedicata **99** (2003) 179–220.
- [8] W. Geyer, *On Tamari lattices*, Disc. Math. **133** (1994) 99–122.
- [9] S. Hanke, T. Ottmann, & S. Schuierer, *The edge-flipping distance of triangulations*, J. Universal Comput. Sci. **2-8** (1996) 570–579.

- [10] S. Huang & D. Tamari, *Problems of associativity: A simple proof for the lattice property of systems ordered by a semi-associative law*, J. Combinat. Th., Series A **13** (1972) 7–13.
- [11] F. Hurtado & M. Noy, *Graph of triangulations of a convex polygon and tree of triangulations*, Computational Geometry **13** (1999) 179–188.
- [12] D. Krammer, *A class of Garside groupoid structures on the pure braid group*, Trans. Amer. Math. Soc. **360** (2008) 4029–4061.
- [13] J. Lucas, *The rotation graph of binary trees is Hamiltonian*, J. Algorithms **8-4** (1988) 503–535.
- [14] R. McKenzie & R.J. Thompson, *An elementary construction of unsolvable word problems in group theory*, in Word Problems, Boone & al. eds., Studies in Logic vol. 71, North Holland (1973).
- [15] N. Reading, *Cambrian lattices*, Adv. Math. **205-2** (2006) 313–353.
- [16] R.O. Rogers, *On finding shortest paths in the rotation graph of binary trees*, Congressus Numerantium **197** (1999) 77–95.
- [17] D. Sleator, R. Tarjan, W. Thurston, *Rotation distance, triangulations, and hyperbolic geometry*, J. Amer. Math. Soc. **1** (1988) 647–681.
- [18] Z. Šunić, *Tamari lattices, forests, and Thompson monoids*, Europ. J. Combinatorics **28-4** (2007) 1216–1238.
- [19] D. Tamari, *The algebra of bracketings and their enumeration*, Nieuw Archief voor Wiskunde **3-10** (1962) 131–146.
- [20] R.J. Thompson, *Embeddings into finitely generated simple groups which preserve the word problem*, in Word problems II: The Oxford book, S.I. Adjan & al. eds., Studies in Logic vol. 95, North Holland (1980) 401–440.

LABORATOIRE DE MATHÉMATIQUES NICOLAS ORESME, UNIVERSITÉ DE CAEN, 14032 CAEN, FRANCE

E-mail address: dehornoy@math.unicaen.fr

URL: [//www.math.unicaen.fr/~dehornoy](http://www.math.unicaen.fr/~dehornoy)