

Self-avoiding walk is ballistic on graphs with more than one end

Florian Lehner^{*1}, Christian Lindorfer^{†2}, and
Christoforos Panagiotis^{‡3}

¹University of Auckland, New Zealand

²Technische Universität Graz, Austria

³University of Bath, UK

March 18, 2024

Abstract

1 Introduction

A *self-avoiding walk* on a graph G is a walk that visits each vertex at most once. This notion was originally introduced in the work of the chemist Flory [12] to model long polymer chains, and it soon attracted the interest of the mathematical community. The primary focus has been on studying the asymptotic behaviour of a self-avoiding of a given length sampled uniformly at random, giving rise to questions that, while simple to pose, frequently prove challenging to resolve.

A significant amount of research on self-avoiding walks has concentrated on answering these questions in the case of lattices in \mathbb{R}^d , where the model is now well-understood in dimensions $d \geq 5$ by the seminal work of Hara and Slade [30, 31]. The low-dimensional cases continue to present serious challenges. See [29, 33, 34, 8, 7, 6, 9, 35] and the references therein for some of the most important results. For a comprehensive introduction to the model in this context interested readers can refer to [2, 38].

Over the years, the study of self-avoiding walk beyond lattices has received increasing attention. The systematic study of self-avoiding walk on general transitive graphs was

^{*}F. Lehner was supported by FWF (Austrian Science Fund) project P31889-N35.

[†]C. Lindorfer was partially supported by FWF (Austrian Science Fund) projects P31889-N35 and DK W1230.

[‡]C. Panagiotis was supported by the Swiss National Science Foundation and the NCCR SwissMAP.

initiated in a series of papers by Grimmett and Li [15, 20, 16, 17, 21, 18, 19, 23], whose work is primarily concerned with properties of the connective constant. Other works on self-avoiding walk in this context include [39, 40, 22, 14, 3, 37, 41, 4, 13]. See [23] for a survey of these results.

In this paper, we consider the self-avoiding walk on quasi-transitive graphs with more than one end. Such graphs admit a group invariant *tree decomposition* $\mathcal{T} = (T, \mathcal{V})$; we refer the reader to Section 3 for the precise definition, but note that any such tree decomposition gives rise a quasi-transitive action of $\text{AUT}(G)$ on T . The concept of tree decomposition was initially introduced by Halin in 1976 [26] and later rediscovered by Robertson and Seymour [42], playing a pivotal role in proving the Graph Minor Theorem. Tree decompositions that are invariant under some group action were perhaps first studied by Dunwoody and Krön [11], drawing inspiration from a method involving edge cuts introduced by Dunwoody in [10]. In our context, it is crucial that these tree decompositions are not only invariant under a quasi-transitive group of automorphisms, but also satisfy some additional properties – see Corollary 3.2 for the precise statement.

One of the first instances in the study of self-avoiding walks on graphs with more than one end can be found in the work of Alm and Janson [1], where they specifically study the case of 2-ended graphs. This particular case turns out to be more tractable due to the fact that the large-scale structure of the graph is similar to that of a line, that is, the tree T of the tree decomposition is isomorphic to \mathbb{Z} . The remaining case of graphs with infinitely many ends proves to be more challenging, with the currently known results limited to either the properties of the SAW-generating functions away from its critical point [36] or to graphs that satisfy additional geometric assumptions [14, 37]. Our main results answer two fundamental questions in the study of self-avoiding walks in the case of transitive graphs with infinitely many ends, and more generally, quasi-transitive graphs G that satisfy the additional technical property that a quasi-transitive group of automorphisms Γ of G does not fix an end of the tree decomposition T .

Denote by c_n the number of self-avoiding walks of length n on G started from some fixed vertex o . A fundamental quantity in the study of self-avoiding walk is the *connective constant*

$$\mu_w = \lim_{n \rightarrow \infty} c_n^{1/n},$$

where the fact that the limit exists and does not depend on the starting point follows from a standard subadditivity argument [28]. The connective constant is not typically known or expected to take an interesting value, with a notable exception being the hexagonal lattice, where Duminil-Copin and Smirnov proved in a celebrated paper [9] that $\mu_w = \sqrt{2 + \sqrt{2}}$. For this reason, it is often more interesting to estimate the subexponential correction to μ_w^n .

Our first result states that c_n grows asymptotically like μ_w^n .

Theorem 1.1. *If a quasi transitive graphs G has infinitely many ends, and $\text{AUT}(G)$ does not fix an end of T , then there exist $k \geq 1$, $a_1, a_2, \dots, a_k > 0$ and $c > 0$ such that for every $q \geq 1$ and $r = 0, 1, \dots, k - 1$ we have*

$$c_{qk+r} = a_r \mu_w^{qk+r} (1 + O(e^{-cq})). \quad (1)$$

In particular, there exist $C_1, C_2 > 0$ such that

$$C_1 \mu_w^n \leq c_n \leq C_2 \mu_w^n \tag{2}$$

for every $n \geq 1$, and we can choose $C_1 = 1$ if G is a transitive graph.

Our next theorem concerns the displacement of a typical self-avoiding walk of length n . We define \mathbb{P}_n to be the uniform measure on self-avoiding walks of length n in G starting from o , and we write $w = (w_0, w_1, \dots, w_n)$ for a random self-avoiding walk sampled from \mathbb{P}_n . We write $d_G(\cdot, \cdot)$ for the graph distance in G .

Theorem 1.2. *If G has infinitely many ends, and $\text{AUT}(G)$ does not fix an end of T , then there exist constants $c, t > 0$ such that $\mathbb{P}_n(d_G(w_0, w_n) \geq cn) \geq 1 - e^{-tn}$ for every $n \geq 1$.*

It turns out that if $\text{AUT}(G)$ fixes an end of T then it is non-unimodular. To see this, let x and y be two vertices where y lies in some small separator separating x from the fixed end ω , but far away from x . Then the size of the orbit of y under $\text{AUT}(G)_x$ is bounded by some absolute constant, but the size of the orbit of x under $\text{AUT}(G)_y$ grows as we increase the distance between x and y . Transitive graphs whose automorphism group admits a non-unimodular transitive subgroup were treated by Hutchcroft [32], who proved that (2) and the conclusion of Theorem 1.2 hold in this context. Thus we obtain the following result.

Theorem 1.3. *Let G be a transitive graph with infinitely many ends. Then there exists constant $C > 0$ such that*

$$\mu_w^n \leq c_n \leq C \mu_w^n.$$

Moreover, there exist constants $c, t > 0$ such that $\mathbb{P}_n(d_G(w_0, w_n) \geq cn) \geq 1 - e^{-tn}$ for every $n \geq 1$.

1.1 Tools and proof methods

The proofs of Theorems 1.1 and 1.2 follow an approach similar to that of Alm and Janson in [1]. Leveraging the tree decomposition, we decompose the graph into finite or infinite parts, and analyse the restrictions of self-avoiding walks to these parts. These restrictions give rise to the notions of *configurations*, *shapes*, and *arrangements*. Similar notions were also introduced in [36], but as we are working in a greater generality (in particular, parts can be infinite in our setting), some modifications in the definitions are necessary. Intuitively, each shape describes the restriction of a self-avoiding walk to a part of the tree decomposition, and each configuration describes how a self-avoiding walk crosses between two neighbouring parts. Arrangements are collections of compatible shapes and configurations, that is shapes and configurations whose partial descriptions of a self-avoiding walk fit together; see Figure 1 for an example.

Precise definitions are introduced in Section 4; for now we point out that there is a natural bijection between arrangements and self-avoiding walks and hence we may study arrangements in order to understand self-avoiding walks. We note that even though




Figure 1: A self avoiding walk and its decomposition into configurations and shapes. Note that every edge of the graph appears in exactly one part. Dashed edges correspond to detours whose edges do not lie in this part; configurations also keep track which side the detour lies on. The little arrows pointing in and out of every configuration indicate the side on which the first and last edge of the self-avoiding walk lie, respectively.

self-avoiding walks are highly non-Markovian, arrangements satisfy a spatial Markov property in the sense that compatibility of shapes and configurations is a local property; this means that in order to study arrangements, we do not have to consider their global structures, but may focus on the individual configurations and shapes. This observation is only useful for self-avoiding walks which intersect more than one part, as otherwise the self-avoiding walk is the same as its restriction to a part and we have not gained any insight. However, it turns out that the probability that a self-avoiding walk stays within one part is exponentially small – see Lemma 6.6.

The spatial Markov property of configurations gives rise to a natural recursive structure which enables us to obtain a system of equations linking the partition functions of arrangements ‘rooted at’ given configurations. This system can be encoded using a Jacobian matrix $\mathfrak{J}(z)$, thereby reducing the problem of determining the asymptotic behaviour of self-avoiding walk to identifying which irreducible components of $\mathfrak{J}(z)$ have spectral radius equal to 1 at the critical point $z = 1/\mu_w$.

We will distinguish between irreducible components rooted at *I-configurations* and *U-configurations*. Precise definitions will be given later, but intuitively, an I-configuration means that the corresponding self-avoiding walk starts and ends on different sides, and a U-configuration between two parts means that the corresponding self-avoiding walk starts and ends on the same side. For instance, configurations c_2 and c_3 in the example in Figure 1 are I-configurations, and configurations c_1 , c_4 , and c_5 are U-configurations.

As a key step in our approach, we establish that the spectral radius of each component of U-configurations is strictly smaller than 1 at the critical point. After developing a

geometric construction that allows us to transform walks, we employ an argument in the spirit of Kesten’s pattern lemma [33] to show that self-avoiding walks consisting only of U-configurations are exponentially rare. In particular, this holds for *self-avoiding polygons*, walks that start and end at the same vertex and are otherwise self-avoiding. The latter is known to imply that self-avoiding walk is ballistic [41], and Theorem 1.2 follows.

To deduce Theorem 1.1, we establish that only one irreducible component possesses a spectral radius equal to 1 at $z = 1/\mu_w$, namely the *persistent* I-configurations. Leveraging the Perron-Frobenius theorem, we conclude that the spectral radius is a simple eigenvalue, allowing us to deduce that the self-avoiding walk-generating function has only simple poles. This, in turn, implies that c_n grows asymptotically like μ_w^n .

1.2 Paper organisation

In Section 2 we define the model and gather some relevant definitions. In Sections 3 and 4 we introduce the notions of tree decompositions, configurations, and arrangements and prove some basic facts about them. In particular, we give a bijection between self-avoiding walks and certain configurations which will allow us to translate between the two objects. Next, in Section 5 we introduce the Jacobian matrix $\mathfrak{J}(z)$ and we establish the connection with the spectral radius of its irreducible components. Subsequently, in Section 6 we show the analyticity of $\mathfrak{J}(z)$ and the partition functions of U-configurations at the critical point. In Section 7 we prove that the spectral radius of several irreducible components is strictly smaller than 1 at the critical point. Finally, in Section 8 we finish the proof of Theorems 1.1 and 1.2.

2 Definitions and basic background

In this section, we gather some definitions that will be used throughout the paper. Most of our notation is standard, but there is some new notation as well; we hence encourage even readers familiar with graph theory to skim this section.

2.1 Graph-theoretic definitions

A *digraph* G consists of a set of *vertices* $V(G)$ and a set of *arcs* $E(G)$. Arcs are considered to be *oriented*, so every arc e is assigned an *initial vertex* $e^- \in V(G)$ and a *terminal vertex* $e^+ \in V(G)$, which are different vertices of G ; note that we allow different edges to have the same initial and terminal vertex. A *graph* is a digraph G together with a bijection $\bar{\cdot}: E(G) \rightarrow E(G)$ such that $\bar{e}^+ = e^-$, $\bar{e}^- = e^+$ and $\bar{\bar{e}} = e$. This means that arcs appear in pairs e, \bar{e} having the same endpoints but different direction. We call such an unordered pair $\{e, \bar{e}\}$ an *edge*. A digraph is called *simple* if it contains no arcs which are different but have the same initial and terminal vertex. In this case we sometimes abuse notation and write $e = e^-e^+$. For a vertex v of G , we denote by $E(v)$ the set of all arcs with initial vertex v and by $\deg(v) = |E(v)|$ the *degree* of v , that is, $\deg(v)$ is

the number of outgoing arcs of v . A digraph is called *locally finite*, if all vertices have finite degree.

2.2 Definition of the model

A *walk* in a digraph is an alternating sequence $p = (v_0, e_1, v_1, \dots, e_n, v_n)$ of vertices $v_i \in V(G)$ and arcs $e_i \in E(G)$ such that $e_i^- = v_{i-1}$ and $e_i^+ = v_i$ for every $i \in [n]$. We point out that we cannot define a walk purely as a sequence of vertices because there may be more than one edge connecting the same pair of vertices. The *length* of the walk $p = (v_0, e_1, v_1, \dots, e_n, v_n)$, denoted by $|p|$, is the number n of arcs in p . Its *initial vertex* is $p^- = v_0$ and its *terminal vertex* is $p^+ = v_n$. For convenience, we include in our definition the *trivial walk* (v) of length 0, starting (and ending) at a vertex v and also the *empty walk* \emptyset consisting of no vertices and no arcs.

A *self-avoiding walk* (or *SAW*) is a walk p whose vertices are pairwise different. We write $c_n(x)$ for the number of self-avoiding walks of length n with initial vertex x , and c_n for $c_n(o)$, where o is some fixed root vertex in G . We define

$$\mu_w(x) = \limsup_{n \rightarrow \infty} c_n(x)^{1/n}.$$

We recall that on a quasi-transitive graph, the limit exists and does not depend on the base point [28]. In this case, we will simply write μ_w .

A *self-avoiding polygon* is a walk $p = (v_0, e_1, v_1, \dots, e_n, v_n)$ with $v_0 = v_n$ and with $v_i \neq v_j$ for distinct pairs i, j other than the pair $0, n$. As in [41], we identify two self-avoiding polygons which share the same set of edges. We write $p_n(x)$ for the number of self-avoiding polygons of length n with initial vertex x . We also write p_n for the number of self-avoiding polygons of length n starting at the fixed root vertex o . We define

$$\mu_p(x) = \limsup_{n \rightarrow \infty} p_n(x)^{1/n},$$

and

$$\mu_p = \limsup_{n \rightarrow \infty} \left(\sup_{x \in V} p_n(x) \right)^{1/n}.$$

In certain cases, it is more convenient to work with self-avoiding walks that start at a vertex x and end at a neighbour vertex of x . We call these walks *self-avoiding returns*.

2.3 Graph distance and connectedness

Let G be a digraph and let u and v be vertices of G . The distance $d_G(u, v)$ from u to v is the length of the shortest walk with initial vertex u and terminal vertex v . If no such walk exists, $d_G(u, v)$ is infinite. We extend this notation to subsets of $V(G)$ in the obvious way: the distance of two sets is the minimal distance between elements of these sets. We also extend this definition to arcs of G . For technical reasons we prefer a pair of arcs to have distance 0 if and only if the two arcs have the same set of endpoints.

Thus for two arcs $e, f \in E(G)$ we define

$$d_G(e, f) = \begin{cases} 0 & \text{if } \{e^-, e^+\} = \{f^-, f^+\} \text{ (as sets),} \\ d_G(\{e^-, e^+\}, \{f^-, f^+\}) + 1 & \text{otherwise.} \end{cases}$$

Observe that if G is a graph, the distance is symmetric on both $V(G)$ and $E(G)$, however it is only a metric on $V(G)$ because an arc and its inverse have distance 0.

A digraph G is *strongly connected* if for any two vertices u, v there is a walk p in G starting at u and ending at v . Note that this implies that any two vertices in the digraph are connected by walks in both directions. A *strong component* of G is a maximal strongly connected subgraph. If G is a graph, we usually omit the word ‘‘strong’’ in both notations.

For $K \subseteq V(G)$ we denote by $G - K$ the subgraph obtained from G by removing K and all arcs incident to K . If removing K disconnects G , then K is called a *separating set*. Furthermore, we denote by $G[K]$ the subgraph of G *induced* by K , that is the graph $G - (V(G) \setminus K)$.

A walk is *closed* if its initial and terminal vertex coincide. A closed walk p of length at least 3 is called a *cycle* if all vertices except the initial and terminal vertex are pairwise different; note that closed walks of length 2 are not considered cycles in this paper even though they use different arcs for both directions. A *tree* is a connected and cycle free graph which does not contain any cycle. A tree consisting only of vertices of degree at most 2 is called a *path*. We point out that unlike walks, we consider paths to be graphs and therefore a path has no ‘direction’; a finite path can be seen as the support of a SAW. Given two disjoint subsets A and B of vertices of a graph G , an *A - B -path* on G is a finite path intersecting A and B only in its two endpoints.

2.4 Surgery on walks and multi-walks

Let p be a walk. For two vertices u and v of p we write upv for the maximal sub-walk of p starting at u and ending at v . If $u = v_0$ or $v = v_n$ we omit the corresponding vertex and denote the sub-walk by pv or up , respectively. We extend this notation even further. For walks p_1, \dots, p_n and vertices v_0, \dots, v_n in the respective walks, we denote the concatenation $(v_0p_1v_1)(v_1p_2v_2) \dots (v_{n-1}p_nv_n)$ of the sub-walks $v_{i-1}p_iv_i$ by $v_0p_1v_1p_2 \dots p_nv_n$. If the terminal vertex v of p_1 coincides with the initial vertex of p_2 , we write p_1p_2 instead of p_1vp_2 , and similarly for concatenations of multiple walks. If e is an arc connecting the terminal vertex v_1 of p_1 to the initial vertex v_2 of p_2 , then we write p_1ep_2 instead of $p_1v_1(v_1, e, v_2)v_2p_2$, and similarly for concatenations with more than two parts.

A *multi-walk* p is a sequence of vertices and arcs obtained by stringing together the sequences of vertices and arcs corresponding to walks p_1, \dots, p_k ; the p_i are called the *walk components* of p . In other words, a multi-walk is a sequence of vertices and arcs, such that every arc in the sequence is preceded by its initial vertex and succeeded by its terminal vertex. Note that each of the walks p_i is a sequence starting and ending with a vertex, so that the final vertex of p_i and the initial vertex of p_{i+1} will appear next to

each other in the sequence p . In fact every appearance of two consecutive vertices in a multi-walk always marks the start of a new walk component.

2.5 Rays and ends

A *ray* is a one-way infinite path and a *double ray* is a two-way infinite path.

Two rays in a graph G are called *equivalent*, if for every finite set $K \subseteq V(G)$ they end up in the same component of $G - K$, that is, all but finitely many of their vertices are contained in that component. An *end* of G is an equivalence class of rays with respect to this equivalence relation. Note that for every finite set $K \subseteq V(G)$ and every end ω , there is a component H of $G - K$ such that all but finitely many vertices of each ray in ω are contained in H ; in such a case we say that ω *lies in* H . Two ends ω_1 and ω_2 of a graph G are *separated* by K if they lie in different components of $G - K$. Halin [24] showed that an end containing arbitrarily many disjoint rays must contain an infinite family of disjoint rays, hence the maximum number of disjoint rays contained in an end ω is well defined and lies in $\mathbb{N} \cup \{\infty\}$. This number is called the *size* of the end ω . An end of finite size is called *thin*, and an end of infinite size is called *thick*.

2.6 Graph automorphisms

An *automorphism* γ of a graph G is a permutation of $V(G)$ and $E(G)$ preserving the neighbourhood relation in G . More precisely, for every arc $e \in E(G)$ it satisfies $(\gamma e)^- = \gamma e^-$ and $\gamma \bar{e} = \overline{\gamma e}$. The set of all automorphisms of G forms a group which is called the *automorphism group* of G and denoted by $\text{AUT}(G)$. For a subgroup $\Gamma \subseteq \text{AUT}(G)$ we can define an equivalence relation on $V(G)$ by $u \sim v \iff \exists \gamma \in \Gamma: u = \gamma v$. The equivalence classes with respect to this relation are called *orbits*, the orbit containing a vertex v is denoted by Γv . We say that Γ acts *vertex transitively*, or simply *transitively*, if there is exactly one orbit, and that it acts *quasi-transitively*, if there are only finitely many orbits. In this case the graph G is also called *(quasi-)transitive*. Similarly, we say that Γ acts *arc transitively*, if the action of Γ on $E(G)$ has a single orbit and *edge transitively*, if it admits a single orbit on the set of edges. A subgraph H of G is called γ -invariant for $\gamma \in \text{AUT}(G)$ if $\gamma(H) = H$.

It is well known, that if an infinite, locally finite, connected graph is quasi-transitive, then it has either one, two, or infinitely many ends. If it has one end, this end is thick. If it has two ends, both are thin and must have the same size. Finally, if it has infinitely many ends, then it must have thin ends. These and many more results were given by Halin in [25].

2.7 Open subgraphs

It will be convenient for us to consider subgraphs which are allowed to contain edges of which they do not contain both endpoints.

An *open subgraph* H of G consists of a vertex set $V(H) \subseteq V(G)$ and the set $E(H) = \{e \in E(G) \mid e^- \in V(H)\}$ of all outgoing arcs of $V(H)$ in G . Note that H is uniquely determined by its vertex set $V(H)$; we say that H is the open subgraph *induced* by $V(H)$.

We call H finite, if $V(H)$ is finite. Arcs $e \in E(H)$ such that $e^+ \in V(G) \setminus V(H)$ are called *boundary arcs* of H , and the set of all boundary arcs is denoted $\partial E(H)$. Observe that an open subgraph is a graph if and only if the set of boundary arcs is empty. The main idea behind the definition of open subgraphs is the following: for every partition of the vertex set $V(G)$, the graph G is the disjoint union of the open subgraphs of G induced by the sets in the partition.

An open subgraph H of G is called *open subtree* if $G[V(H)]$ is a tree and it is called *open path* if $G[V(H)]$ is a path. The *leaves* of an open subtree are the leaves of the underlying tree; in other words, when talking about the leaves of an open subtree, we ignore all boundary edges. For convenience, we let the length of an open path be the number of vertices it contains. Note that this is not the length of the underlying path—intuitively we may imagine that an open path contains an additional ‘half edge’ in the form of an outgoing arc at each of its endpoints. An *open star* S in G is an open path of length 1; we point out that the underlying tree of an open star is not a star, but a single vertex. It is induced by a single vertex s of G , called the *center* of S and we denote it by $\text{star}_G(s)$ or $\text{star}(s)$ if the graph is clear. Additionally, for an arc e we write $\text{star}(e)$ for the open subtree with vertex set $\{e^-, e^+\}$. If T is a tree, the (*open*) *cone* K_e in T rooted at the arc e consists of e and the component of $T - \{e, \bar{e}\}$ containing the initial vertex e^- of e .

Let T be a tree and e, f be two different arcs of T . We call e and f *linkable* if they point away from each other, that is, they are not contained in the unique path in T connecting e^- and f^- . If e and f are linkable, the (unique) *link* W connecting e and f is the open path induced by this path: it connects e^- and f^- , and e and f are boundary arcs at the two endpoints of W . Observe that whenever e and f are linkable, the length of their link W is $d_T(e, f) = d_T(e^-, f^-) + 1$.

3 Tree decompositions

Recall that our general proof strategy is to decompose SAWs into configurations and shapes. In order to be able to do this, we first need to decompose the graph appropriately. To this end, the goal of the current section is to show that every quasi-transitive graph with more than one end admits a tree decomposition with certain ‘nice properties’. We start by giving some definitions.

A *tree decomposition* of a graph G is a pair $\mathcal{T} = (T, \mathcal{V})$, consisting of a tree T and a function $\mathcal{V} : V(T) \rightarrow 2^{V(G)}$ assigning a subset of $V(G)$ to every vertex of T , such that the following three conditions are satisfied:

- (T1) $V(G) = \bigcup_{t \in V(T)} \mathcal{V}(t)$.
- (T2) For every arc $e \in E(G)$ there is a vertex $t \in V(T)$ such that $\mathcal{V}(t)$ contains both endpoints of e .
- (T3) $\mathcal{V}(s) \cap \mathcal{V}(t) \subseteq \mathcal{V}(r)$ for every vertex r on the unique s - t -path in T .

For every $t \in V(t)$, the set $\mathcal{V}(t)$ is called the *part* of t . For an arc $e = st$ of T , the intersection $\mathcal{V}(e) = \mathcal{V}(s, t) := \mathcal{V}(s) \cap \mathcal{V}(t)$ is called the *adhesion set* of e . Note that by definition $\mathcal{V}(e) = \mathcal{V}(\bar{e})$.

A *separation* of a graph G is a pair (A, B) of vertex sets such that $G[A] \cup G[B] = G$, in other words $A \cup B = V(G)$, and there are no edges between $A \setminus B$ and $B \setminus A$. The intersection $A \cap B$ is called the *separator*, and $|A \cap B|$ is called the *order* of (A, B) . A separation (A, B) of finite order is said to *separate* two ends ω_1 and ω_2 if one of them lies in a component of $B \setminus A$, and the other one lies in a component of $A \setminus B$. We say that (A, B) *minimally separates* ω_1 and ω_2 if (A, B) has the minimal order among all separations separating ω_1 and ω_2 . Note that every arc a of the decomposition tree T corresponds to a separation of G with separator $\mathcal{V}(a)$: we call

$$\left(\bigcup_{t \in V(K_a)} \mathcal{V}(t), \bigcup_{t \in V(K_{\bar{a}})} \mathcal{V}(t) \right)$$

the separation of G *induced by* a . Clearly, if the separation induced by a is (A, B) , then the separation induced by \bar{a} is (B, A) .

A tree decomposition (T, \mathcal{V}) of G is called Γ -*invariant* for a group $\Gamma \leq \text{AUT}(G)$, if every $\gamma \in \Gamma$ maps parts onto parts and thereby induces an automorphism of T (which we denote by γ as well); this clearly induces an action of Γ on T .

The following result which guarantees the existence of a Γ -invariant tree decomposition is at the heart of our construction. It can be seen as a special case of [11, Theorem 8.1]; simply note that the separations (A, B) with $|A \cap B| = k$ and infinite sides A and B correspond to a cut system in the sense of [11], and that the G -tree $T(\mathcal{C})$ can be interpreted as a tree decomposition whose induced separations are all contained in this cut system.

Theorem 3.1. *Let Γ be a group acting quasi-transitively on a locally finite graph G . Assume that G has more than one end, and let k be the minimal integer such that there are two ends which can be separated by removing k vertices. Then there is a Γ -invariant tree decomposition of G in which every induced separation has order k and separates two ends.*

Corollary 3.2 below describes the tree decompositions we will be working with for the rest of this paper. It is similar to [27, Corollary 4.3] and [36, Corollary 3.2]. We still provide a proof for the convenience of the reader since these two results are slightly less general than ours.

We call a tree decomposition *reduced* if every induced separation minimally separates some pair of ends of G , and no two parts corresponding to adjacent vertices in T coincide. We call it *strongly reduced*, if in addition every induced separation has order k , where k is the minimal integer such that there is a pair of ends which can be separated by removing k vertices.

In order to extract a tree decomposition with the desired properties from the above theorem, we will need the following construction. Let $\mathcal{T} = (T, \mathcal{V})$ be a tree decomposition

of G , and let $F \subseteq E(T)$ be a subset of arcs such that $\bar{e} \in F$ whenever $e \in F$. We define the *contraction* $\mathcal{T}/F = (T/F, \mathcal{V}/F)$ of the tree decomposition \mathcal{T} as follows. The vertices of T/F are components of $T - (E(T) \setminus F)$ with an arc between two of them if there is an arc in T (which is necessarily an element of $E(T) \setminus F$) connecting the corresponding components. For a component $S \in V(T/F)$ we set $\mathcal{V}(S) = \bigcup_{s \in V(S)} \mathcal{V}(s)$. It is not hard to see that this defines a tree decomposition, and that the arcs of the decomposition tree are precisely the arcs in $E(T) \setminus F$. Moreover, the separations induced by an arc $e \in E(T) \setminus F$ and its counterpart in $E(T/F)$ coincide.

Corollary 3.2. *Let Γ be a group acting quasi-transitively on a locally finite graph G . Then there is a strongly reduced, Γ -invariant tree decomposition (T, \mathcal{V}) such that the action of Γ on T is edge transitive, but not arc transitive.*

Proof. We start with the tree decomposition $\mathcal{T} = (T, \mathcal{V})$ provided by Theorem 3.1 and successively change it to satisfy the additional conditions.

First, let $e \in E(T)$ and let $F = E(T) \setminus (\Gamma e \cup \Gamma \bar{e})$. Then \mathcal{T}/F is easily seen to be Γ -invariant, and the action of Γ on T/F is transitive on $E(T/F)$.

We claim that \mathcal{T}/F is reduced. Assume that it is not. Then there are two parts corresponding to adjacent vertices of T/F which coincide. Edge transitivity implies that all parts of \mathcal{T}/F coincide, and thus all adhesion sets coincide with the parts as well. This cannot be the case since G is infinite, but the adhesion sets of \mathcal{T}/F are precisely the adhesion sets of \mathcal{T} corresponding to arcs in Γe and thus they are finite. Finally, \mathcal{T}/F is strongly reduced due to the size of the adhesion sets in the tree decomposition \mathcal{T} provided by Theorem 3.1.

Note that if the action on T/F in the above construction is not arc transitive, then we are done. Hence we may assume that we have a strongly reduced tree decomposition $\mathcal{T} = (T, \mathcal{V})$ such that the action of Γ on T is arc transitive. Let T' be the tree obtained by subdividing each edge in T precisely once; more formally, $V(T') = V(T) \cup \{\{e, \bar{e}\} \mid e \in E(T)\}$, and $t \in V(T)$ and $\{e, \bar{e}\}$ are adjacent in T' if and only if t and e are incident in T . Then it is straightforward to check that $\mathcal{T}' = (T', \mathcal{V})$ defines a tree decomposition; recall that $\mathcal{V}(st)$ is defined as $\mathcal{V}(s) \cap \mathcal{V}(t)$. This tree decomposition is still Γ -invariant, since the action on the tree T induces an action on the tree T' which takes parts to parts and adhesion sets to adhesion sets.

Next we show that \mathcal{T}' is strongly reduced. Let $st \in E(T)$. Then $|\mathcal{V}(s)| = |\mathcal{V}(t)|$ by arc transitivity, and consequently $|\mathcal{V}(s)| > |\mathcal{V}(st)|$ because otherwise the parts corresponding to s and t would coincide. Moreover, note that every adhesion set of \mathcal{T}' is also an adhesion set of \mathcal{T} since $\mathcal{V}(s) \cap \mathcal{V}(st) = \mathcal{V}(s) \cap \mathcal{V}(t)$. This shows that \mathcal{T}' is strongly reduced.

The action of Γ is not transitive on the arcs of T' , since it is impossible to map the set $\mathcal{V}(s)$ to the set $\mathcal{V}(st)$ because the sets have different cardinalities. To see that the action is transitive on edges, note that by arc transitivity of the action on T , for any two arcs $e_1 = s_1 t_1$ and $e_2 = s_2 t_2$ there is $\gamma \in \Gamma$ such that $\gamma e_1 = e_2$ and thus in particular $\gamma s_1 = s_2$. We conclude that γ maps the edge connecting s_1 and $\{e_1, \bar{e}_1\}$ to the edge connecting s_2 and $\{e_2, \bar{e}_2\}$ in T' . \square

We now collect some useful properties of the tree decompositions provided by Corol-

lary 3.2. For the remainder of this section let Γ be a group acting quasi-transitively on a locally finite graph G and let $\mathcal{T} = (T, \mathcal{V})$ be a tree decomposition as provided by Corollary 3.2.

Proposition 3.3. *For every vertex $v \in V(G)$ there are only finitely many $s \in V(T)$ such that $v \in \mathcal{V}(s)$.*

Proof. Assume for a contradiction that v is contained in infinitely many parts. Let $U \subseteq V(G)$ be a maximal subset such that $v \in U$ and there are infinitely many parts containing all of U . Note that such a maximal subset exists because if a set of vertices is contained in two parts then by (T3) it is contained in two neighbouring parts. Therefore any set of vertices which is contained in infinitely many parts must be contained in an adhesion set, and thus there is an upper bound on the size of such a set.

By (T3), the set $\{s \in V(T) : U \subseteq \mathcal{V}(s)\}$ induces a subtree T' of T . This subtree is infinite because there are infinitely many parts containing U . Define $F = E(T')$ and note that $U \subseteq \mathcal{V}(e)$ for every $e \in F$. The separation (A_e, B_e) induced by $e \in F$ separates two ends, hence v must have neighbours $a_e \in A_e \setminus B_e$ and $b_e \in B_e \setminus A_e$, otherwise it would be possible to separate the ends by fewer vertices.

By (T2), both endpoints of the arc $va_e \in E(G)$ must be contained in some part of the tree decomposition, and since $a_e \notin B_e$ this part is $\mathcal{V}(s)$ for some $s \in K_e$. Similarly, both endpoints of the arc $vb_e \in E(G)$ are contained in a part $\mathcal{V}(s)$ for some $s \in K_{\bar{e}}$.

Since F contains infinitely many arcs, there is an infinite subset $S \subseteq V(T)$ such that $\mathcal{V}(s)$ contains a neighbour of v for every $s \in S$. But v only has finitely many neighbours, so infinitely many parts contain the same neighbour u of v .

Taking an appropriate infinite subset of F (either the arcs of some infinite star, or arcs on a ray which are sufficiently far apart) we can make sure that for any two vertices in S , the path connecting them contains at least one vertex of T' . By (T3) this implies that u is contained in $\mathcal{V}(s)$ for infinitely many $s \in T'$; this contradicts the maximality of U . \square

Corollary 3.4. *There is a constant N such that $\mathcal{V}(e) \cap \mathcal{V}(f) = \emptyset$ whenever $e, f \in E(T)$ such that $d_T(e, f) \geq N$.*

Proof. By Proposition 3.3 every vertex $v \in V(G)$ is contained in only finitely many parts of the tree decomposition. Moreover G is quasi-transitive, so there is a constant N such that every $v \in V(G)$ is contained in at most N parts. Assume that e and f are arcs of T such that $\mathcal{V}(e) \cap \mathcal{V}(f)$ contains a vertex v and $d_T(e, f) \geq N$. Then the unique path in T connecting e^- and f^- has at least $N + 1$ vertices. It follows from (T3) that $v \in \mathcal{V}(t)$ for each vertex t of this path, contradicting the fact that every vertex is contained in at most N parts. \square

Proposition 3.5. *There is a map $\theta: E(G) \rightarrow V(T)$ such that for every arc $e \in E(G)$*

- (i) *both endpoints of e are contained in $\mathcal{V}(\theta(e))$,*
- (ii) *$\theta(e) = \theta(\bar{e})$, and*

(iii) $\gamma\theta(e) = \theta(\gamma e)$ for every $\gamma \in \Gamma$.

Proof. For every arc $e \in E(G)$, let $S(e) = \{s \in V(T) : e^-, e^+ \in \mathcal{V}(s)\}$. The set $S(e)$ is non-empty by (T2), the induced subgraph $T[S(e)]$ is a tree by (T3), and this tree is finite by Proposition 3.3. Hence it either has a central vertex, or a central edge. In the former case, we let $\theta(e)$ be the central vertex of $T[S(e)]$. In the latter case, we note that, because Γ acts edge transitively but not arc transitively, there is an orientation of the edges such that Γ preserves the chosen orientations. Let $\theta(e)$ be the initial vertex of the central edge of $T[S(e)]$ with respect to this orientation.

Clearly, every $\gamma \in \Gamma$ which maps e to f also maps $S(e)$ to $S(f)$, and thus maps the central vertex or edge of $T[S(e)]$ to the central vertex or edge of $T[S(f)]$. Since γ also preserves the orientation we picked above, we conclude that $\gamma\theta(e) = \theta(\gamma e)$ as desired. \square

Now let $\mathcal{E}(s) = \theta^{-1}(s) \subseteq E(G)$, where θ is the function given by Proposition 3.5. Then $\mathcal{E}(s)$ is a subset of the edge set of $G[\mathcal{V}(s)]$.

Additionally we introduce for every arc $e = st$ of T a new set of *virtual arcs* $\mathcal{E}(e) = \mathcal{E}(st)$, such that every pair of vertices of $\mathcal{V}(e)$ is connected by an arc in $\mathcal{E}(e)$. By definition these arcs come in pairs connecting the same vertices but having different direction, so that $\mathcal{V}(e)$ and $\mathcal{E}(e)$ form a complete graph. Note that the sets $\mathcal{E}(e)$ and $\mathcal{E}(\bar{e})$ are disjoint by definition.

We define the *adhesion graph* $\mathcal{G}(e) = (\mathcal{V}(e), \mathcal{E}(e) \uplus \mathcal{E}(\bar{e})) (= \mathcal{G}(\bar{e}))$ and note that every pair of vertices is connected by precisely two arcs, and thus $\mathcal{G}(e)$ is not simple unless it only consists of a single vertex. In order to enhance readability, we usually write $\mathcal{E}(s, t)$ instead of $\mathcal{E}(st)$ and $\mathcal{G}(s, t)$ instead of $\mathcal{G}(st)$.

Finally, we assign to every vertex s of T the *part graph*

$$\mathcal{G}(s) = \left(\mathcal{V}(s), \mathcal{E}(s) \uplus \biguplus_{e: e^- = s} \mathcal{E}(e) \right).$$

Again, $\mathcal{G}(s)$ generally is not a simple graph since $\mathcal{E}(s)$ and the various sets $\mathcal{E}(s, t)$ potentially contain arcs with the same endpoints.

Recall that an open subtree S of T consists of a vertex set $V(S)$ and all arcs of T starting at those vertices such that $T[V(S)]$ is a tree. The *part graph induced by S* is

$$\mathcal{G}(S) = \left(\bigcup_{t \in V(S)} \mathcal{V}(t), \bigcup_{t \in V(S)} \mathcal{E}(t) \uplus \biguplus_{e: \partial E(S)} \mathcal{E}(e) \right).$$

This is a natural extension of part graphs; it is easy to see that $\mathcal{G}(\text{star}(s)) = \mathcal{G}(s)$.

Proposition 3.6. *For every $s \in V(T)$, the setwise stabiliser $\Gamma_{\mathcal{V}(s)}$ acts quasi-transitively on $\mathcal{G}(s)$.*

Proof. If $u \in \mathcal{V}(t)$ does not lie in any adhesion set, then neither does any image of u under a graph automorphism. In particular, any $\gamma \in \Gamma$ mapping u to some vertex

$v \in \mathcal{V}(t)$ fixes $\mathcal{V}(t)$ and thus, under the action of the stabiliser of $\mathcal{V}(t)$ there are only finitely many orbits of vertices in $\mathcal{V}(t)$ not contained in any adhesion set.

Whenever $\gamma \in \Gamma$ fixes s and maps a neighbour t of s onto some other neighbour t' , γ lies in $\Gamma_{\mathcal{V}(t)}$ and maps the adhesion set $\mathcal{V}(s, t)$ onto $\mathcal{V}(s, t')$. By edge transitivity, there are at most two orbits of adhesion sets contained in $\mathcal{V}(s)$ under the action of $\Gamma_{\mathcal{V}(t)}$. As every adhesion set contains the same finite number of elements, $\Gamma_{\mathcal{V}(t)}$ acts with finitely many orbits on the vertices which lie in adhesion sets of the tree decomposition. \square

Proposition 3.7. *There is $M \in \mathbb{N}$ such that for every $s \in V(T)$ and every $v \in \mathcal{V}(s)$ there is a neighbour t of s such that $d_G(v, \mathcal{V}(s, t)) \leq M$.*

Proof. If u and v are vertices whose distance to the nearest adhesion set differs, then they lie in different orbits with respect to the action of $\Gamma_{\mathcal{V}(s)}$ because $\Gamma_{\mathcal{V}(s)}$ maps adhesion sets to adhesion sets. Hence, if the distance from a vertex to the nearest adhesion set was unbounded, then the action of $\Gamma_{\mathcal{V}(s)}$ would have infinitely many orbits on $\mathcal{V}(s)$, thus contradicting Proposition 3.6. \square

Proposition 3.8. *Let e and f be arcs of T . If Γ does not fix an end of T , then there is an automorphism $\gamma \in \Gamma$ such that e and $\gamma(f)$ are linkable.*

If e and f lie in the same orbit, then for every odd k we can choose γ such that $d_T(e, \gamma(f)) = k$.

Proof. Since the action is transitive on edges but not on arcs, we can choose an orientation of the edges of E which is preserved under the action of Γ . There are two possibilities: either every vertex has both incoming and outgoing arcs with respect to this orientation, or all arcs are oriented from one bipartite part to the other.

In the first case note that every vertex has at least two incoming and outgoing arcs (otherwise Γ would fix an end), and hence there is $\gamma \in \Gamma$ such that $\gamma(f)^- = e^-$ and the two arcs are linked by the open path of length 1 containing only the vertex e^- . In the second case we can find an open path containing one or two vertices, depending on the orientations of the arcs e and f .

If e and f lie in the same orbit, then we can choose γ such that $d_T(e, \gamma(f)) = 1$. Concatenation of these open paths shows that we can also choose γ such that $d_T(e, \gamma(f)) = k$ for any odd k . \square

Proposition 3.9. *Let $e, f \in E(T)$. There is a family of $|\mathcal{V}(e)|$ disjoint paths connecting $\mathcal{V}(e)$ to $\mathcal{V}(f)$ in G . Moreover, there is $N \in \mathbb{N}$ such that if $d_T(e, f) \geq N$, then for any pair of vertices $u \in \mathcal{V}(e)$ and $v \in \mathcal{V}(f)$ there is a u - v -path in G which meets $\mathcal{V}(e) \cup \mathcal{V}(f)$ only in u and v .*

Proof. Let (A, B) and (C, D) be the separations induced by e and f , respectively. It is clear from the definition of induced separations that we may without loss of generality assume that $A \cap D \subseteq \mathcal{V}(e) \cap \mathcal{V}(f)$. There are ends ω_A and ω_D which lie in components of $A \setminus B$ and $D \setminus C$, respectively. If there were fewer than $|\mathcal{V}(e)|$ disjoint paths connecting $\mathcal{V}(e)$ to $\mathcal{V}(f)$ in G , then by Menger's theorem it would be possible to separate these

two ends by removing fewer than $|\mathcal{V}(e)|$ vertices; this contradicts the fact that the tree decomposition is strongly reduced.

For the ‘moreover’ part, let $e' \in E(T)$, and let H be a finite connected subgraph of G containing all vertices in $\mathcal{V}(e')$. By Corollary 3.4 there is some $N' \in \mathbb{N}$ such that $\mathcal{V}(f') \cap V(H) = \emptyset$ whenever the distance between e' and f' in T is at least N' .

Now assume that the distance between e and f is at least $N = 2N'$, and let e' be an arc lying in the center of the path connecting e and f in T . The union of H and the collection of disjoint paths from the first part is a connected subgraph of G because each of the paths must pass through a vertex in $\mathcal{V}(e')$. Vertices in $\mathcal{V}(e) \cup \mathcal{V}(f)$ have degree 1 in this graph because of the choice of N . Hence they cannot appear as internal vertices of any path, and thus any path in this subgraph meets $\mathcal{V}(e) \cup \mathcal{V}(f)$ at most in its endpoints. Since the graph is connected, it contains the desired paths. \square

4 Configurations and arrangements

The goal of this section is to utilize tree decompositions to build SAWs from so-called shapes living on the graphs $\mathcal{G}(t)$ for $t \in V(T)$ and configurations living on $\mathcal{G}(e)$ for $e \in E(T)$ where $\mathcal{T} = (T, \mathcal{V})$ is a tree decomposition of a quasi-transitive graph G ; usually the one provided by Corollary 3.2. In the case where G is a one-dimensional lattice, this was already done by Alm and Janson [1]. Here we use a more general approach similar to the one introduced in [36] for tree decompositions whose parts are finite. As we also want to treat infinite parts, our definitions and notation are slightly different from the ones used there.

Let G be connected, locally finite, simple and quasi-transitive, and let $\mathcal{T} = (T, \mathcal{V})$ be a tree decomposition of G . Assume that we have a map θ mapping each edge to a part containing its endpoints as in Proposition 3.5. Define adhesion graphs and part graphs for this tree decomposition in the same way as we did after Proposition 3.5 for the tree decomposition provided by Corollary 3.2. A *configuration* on an arc e of T with respect to \mathcal{T} is a triple $c = (q, x, y)$, where $x, y \in \{e, \bar{e}\}$ are not necessarily different orientations of e , and q is either a SAW on the graph $\mathcal{G}(e)$ or equal to the empty set \emptyset . In the latter case c is called the *empty configuration*.

We call x the *entry direction* and y the *exit direction* of c . The inverse of the configuration $c = (q, x, y)$ is the configuration $\bar{c} = (\bar{q}, y, x)$ consisting of the reverse walk \bar{q} of q , and which has entry direction y and exit direction x .

A *shape* on a vertex s of T with respect to \mathcal{T} is a SAW p on the part graph $\mathcal{G}(s)$. We say that a shape p on s and a configuration (q, x, y) on an arc $e \in E(s)$ are compatible if the following three conditions hold; the intersection of a walk p with a subgraph is defined as the subsequence of p consisting of the vertices and edges contained in that subgraph and therefore this intersection is in general not a walk but a multi-walk.

$$(C1) \quad p \cap \mathcal{G}(e) = q \cap \mathcal{G}(s).$$

$$(C2) \quad \text{If } x = e, \text{ then } p \text{ starts in } \mathcal{V}(e).$$

$$(C3) \quad \text{If } y = e, \text{ then } p \text{ ends in } \mathcal{V}(e).$$

An *arrangement* on a finite open subtree S of T with respect to \mathcal{T} is a pair $A = (P, C)$ of maps assigning to every vertex $s \in V(S)$ a shape $P(s)$ and to every arc $e \in E(S)$ a configuration $C(e) = (Q(e), X(e), Y(e)) = C(\bar{e})$, where $X(e), Y(e)$ may be arcs in $E(S)$ or their inverses, such that for $s \in V(S)$ the following conditions hold:

- (D1) $P(s)$ and $C(e)$ are compatible for every $e \in E(s)$.
- (D2) There is at most one arc $e \in E(s)$ such that $X(e) = e$. If there is no such arc, then $P(s)$ starts with a non-virtual arc.
- (D3) There is at most one arc $e \in E(s)$ such that $Y(e) = e$. If there is no such arc, then $P(s)$ ends with a non-virtual arc.

The *weight* $\|A\|$ of the arrangement $A = (P, C)$ on the open subtree S is the total number of non-virtual arcs contained in all the walks $P(s)$ for $s \in V(S)$, so $\|A\| = \sum_{s \in V(S)} \|P(s)\|$, where $\|P(s)\|$ denotes the number of non-virtual arcs in $P(s)$.

The arrangement A is called *boring* on an arc $e \in E(S)$ if $X(e) = Y(e)$ and all arcs of $Q(e)$ are contained in $\mathcal{E}(X(e))$; in this case we also say that the configuration $C(e) = (Q(e), X(e), Y(e))$ is *boring*. Intuitively, if an arrangement is boring on an arc e , then this means that all non-trivial shapes lie on one side of this arc; in other words we won't lose any information by removing everything that lies on the other side of this arc from the open subtree. We can thus reduce an arrangement by iteratively pruning subtrees attached to edges on which the arrangement is boring until all such subtrees are trivial. We call an arrangement A *reduced* if it is non-boring on all arcs in $E(S) \setminus \partial E(S)$. Call a reduced arrangement *complete* if the configuration $C(e)$ is boring and $X(e) = \bar{e}$ for every $e \in \partial E(S)$.

Our goal is to establish a relation between (complete) arrangements on finite open subtrees S of T and self-avoiding walks of length at least 1 on $\mathcal{G}(S)$. While our definition of arrangements differs slightly from the definitions of configurations in [36], we will follow the same strategy, thus our proofs are quite similar.

The main result of this section is Theorem 4.15 which intuitively states that every self avoiding walk has an arrangement associated to it and (subject to some technical conditions) this is a bijection. The technical details are quite messy, hence we start by providing a very rough sketch of the main ideas for the convenience of the reader. Tree decompositions can be made coarser by contracting edges, and finer by decontracting these edges again, and we can define corresponding operations on arrangements; see Figure 2 for a sketch and Constructions 4.4 and 4.7 for details. Every self-avoiding walk p consists of finitely many edges. Thus after some finite number of contractions in the tree decomposition, all edges of p are contained in one part. Hence p is a shape on this part, and this shape can be extended to a configuration on the open star. Decontracting this arrangement gives the arrangement on the original tree decomposition corresponding to p . We may also perform these steps in reverse order to translate an arrangement into a self-avoiding walk.

In the remainder of this section we provide the details to the above proof sketch. As a first step, we show that a self avoiding walk whose edges belong to a single part can indeed be transformed into an arrangement on the open star.

Lemma 4.1. *Let $S = \text{star}(s)$ be an open star in T and p be a shape on s .*

- (i) *There is an arrangement $A = (P, C)$ on S such that $P(s) = p$.*
- (ii) *For $e \in E(S)$ the walk $Q(e)$ of the configuration $C(e)$ is uniquely defined by p .*
- (iii) *If p starts in $\mathcal{V}(f)$ for some $f \in E(S)$ and/or ends in $\mathcal{V}(f')$ for some $f' \in E(S)$, we may choose A such that $X(f) = f$ and/or $Y(f') = f'$ holds, respectively.*

Proof. Let $e \in E(S)$. Our first goal is to construct a walk q on $\mathcal{G}(e)$ such that (C1) is satisfied and show that this walk is unique. If $p \cap \mathcal{G}(e) = \emptyset$, we set $q = \emptyset$ the empty walk. Otherwise $p \cap \mathcal{G}(e)$ is a multi-walk, let p_1, \dots, p_n be its walk-components. Set $q = p_1 e_1 p_2 \dots e_{n-1} p_n$, where $e_i \in \mathcal{E}(\bar{e})$ connects p_i^+ and p_{i+1}^- . It is not hard to check that (C1) is satisfied. Furthermore, by definition $\mathcal{V}(e) \subseteq V(\mathcal{G}(s))$ and $\mathcal{E}(e) \subseteq E(\mathcal{G}(s))$ holds, thus there is no other way to complete q .

Let now $P(s) = p$ and for $e \in E(S)$ write $Q(e)$ for the walk q on $\mathcal{G}(e)$ constructed above. If $P(s)$ starts in some adhesion set $\mathcal{V}(f)$, we may choose $X(f) = f$ and $X(e) = \bar{e}$ for $e \in E(S) \setminus \{f\}$. Otherwise $P(s)$ starts with a non-virtual arc and we have to choose $X(e) = \bar{e}$ for every $e \in E(S)$. Similarly, if $P(s)$ ends in some adhesion set $\mathcal{V}(f)$ we may choose $Y(f) = f$ and $X(e) = \bar{e}$ for $e \in E(S) \setminus \{f\}$, otherwise we have to choose $X(e) = \bar{e}$ for every $e \in E(S)$. In any case this construction satisfies (D1) – (D3). \square

The following Lemma formalises the intuition behind the notion of boring configurations. In particular, it shows that it is usually sufficient to work with reduced arrangements, as any non-reduced arrangement can be restricted to a smaller open subtree without losing any information.

Lemma 4.2. *Let $S = \text{star}(s)$ be an open star in T , let $e \in E(S)$ be one of its arcs and let $c = (q, x, y)$ be a non-empty boring configuration on e such that $x = y = e$. Then there is a unique arrangement $A = (P, C)$ on S such that $C(e) = c$. Additionally, $P(s)$ contains no non-virtual arcs and $C(f)$ is boring for every $f \in E(S) \setminus \{e\}$.*

Proof. By definition of boring configurations the walk q contains only arcs in $\mathcal{E}(e)$ and no arcs in $\mathcal{E}(\bar{e})$. Thus q is a shape on s and Lemma 4.1 provides the existence of an arrangement $A = (P, C)$ on S such that $P(s) = q$. Additionally q starts and ends in $\mathcal{V}(e)$, so we may choose A such that $X(e) = Y(e) = e$. Note that this is the only way of choosing entry and exit directions which is consistent with the given boring configuration c , and it is easy to check that indeed $C(e) = c$.

Let $f \neq e$ be an arc of S . Then properties (D2) and (D3) imply $X(f) = Y(f) = \bar{f}$ because $X(e) = Y(e) = e$. Also $P(s) \cap \mathcal{G}(f)$ consists only of vertices, so (C1) implies that all arcs of $Q(f)$ are contained in $\mathcal{E}(\bar{f})$, so that $C(f)$ is boring. \square

Remark 4.3. Observe that for any arrangement $A = (P, C)$ on a finite open subtree S of T empty configurations can only occur on boundary arcs of S . This follows from the fact that if $X(e) = e$ for some interior arc e of S , then the walk $P(e^-)$ has to start in $\mathcal{V}(e)$ by (C2) and thus $Q(e)$ is not empty by (C1).




Figure 2: Sketch of contraction (see Construction 4.4) and projection (see Construction 4.7) of arrangements with respect to the arc f .

In the next step we will construct coarser tree decompositions from a given tree decomposition $\mathcal{T} = (T, \mathcal{V})$ of G by contracting arcs of T (similar to the construction used in the proof of Corollary 3.2).

We first recall the definition of contraction in the special case where only a single edge is contracted, and introduce some additional notation for this special case. Let f be an arc of T . The tree decomposition $\mathcal{T}/f = (T/f, \mathcal{V}/f)$ of G obtained from \mathcal{T} by contracting f is given as follows. The vertex set $V(T/f)$ of the decomposition tree T/f is obtained from $V(T)$ by replacing f^- and f^+ by a single new vertex s_f and its edge set $E(T/f)$ is obtained from $E(T)$ by deleting f and \bar{f} and for the remaining edges changing all endpoints in $\{f^-, f^+\}$ to the new vertex s_f . In other words, we contract f and leave the names of vertices and edges unchanged wherever possible. Furthermore, the part corresponding to s_f is $\mathcal{V}/f(s_f) = \mathcal{V}(f^-) \cup \mathcal{V}(f^+)$, for all other vertices $t \in V(T/f)$ we define $\mathcal{V}/f(t) = \mathcal{V}(t)$. It is not hard to check that the result satisfies properties (T1) – (T3), thus being a tree decomposition. The part graph of the new vertex is $\mathcal{G}(s_f) = \mathcal{G}(\text{star}(f))$, all other part graphs and adhesion graphs are inherited from \mathcal{T} and stay the same.

Given an arrangement with respect to the tree decomposition \mathcal{T} , we can define the contraction of the arrangement as described below. The construction and its inverse (Construction 4.7) are sketched in Figure 2.

Construction 4.4. Let $A = (P, C)$ be an arrangement on the open subtree $\text{star}(f)$ of T and let $\text{star}(s_f)$ be the open star in T/f centered in the new vertex s_f introduced by contracting f in T . We construct a pair $(P/f, C/f)$ and show that it is an arrangement on $\text{star}(s_f)$. Configurations stay the same; we set $C/f(e) = C(e)$ for every $e \in E(s_f)$. By Remark 4.3 the configuration $C(f)$ is non-empty, let $Q(f) = (v_1, e_1, \dots, e_{k-1}, v_k)$. By (C1) we have that $P(f^-)$ and $P(f^+)$ both visit v_1, \dots, v_k in this order and no other vertices of $\mathcal{V}(f)$. Let $a_0 = X(f)$, let $a_k = Y(f)$ and for $j \in [k-1]$ let $a_j \in \{f, \bar{f}\}$ such that $e_j \in \mathcal{E}(a_j)$. We define the walk $P/f(s_f)$ as the concatenation

$$P(a_0^+)v_1P(a_1^+)v_2 \dots v_kP(a_k^+).$$

In other words, $P/f(s_f)$ is obtained from $P(f^-)$ and $P(f^+)$ by deleting all arcs in

$\mathcal{E}(f)$ and then piecing the walk components of the resulting multi-walks together in a consistent manner.

Intuitively it should be clear that applying Construction 4.4 to an arrangement yields an arrangement on the contracted tree decomposition. We will now prove this formally.

Lemma 4.5. *The walk $P/f(s_f)$ is a self-avoiding walk on $\mathcal{G}(s_f)$ satisfying $P/f(s_f) \cap \mathcal{G}(f^-) = P(f^-) - \mathcal{E}(f)$ and $P/f(s_f) \cap \mathcal{G}(f^+) = P(f^+) - \mathcal{E}(\bar{f})$. In particular, the set of arcs contained in $P/f(s_f)$ consists of the arc sets of $P(f^-) - \mathcal{E}(f)$ and $P(f^+) - \mathcal{E}(\bar{f})$.*

Proof. Let $P/f(s_f) = P(a_0^+)v_1P(a_1^+)v_2 \dots v_kP(a_k^+)$ as defined above. If one of $P(f^-)$ and $P(f^+)$ is the empty walk, then all claimed properties are trivially satisfied, so we may assume that both walks are non-empty. By (C2), $P(X(f)^-)$ must start in v_1 and by (C3), $P(Y(f)^-)$ must end in v_k , so both $P(a_0^+)v_1$ and $v_kP(a_k^+)$ are trivial. By definition $Q(f)$ is a walk consisting of arcs e_1, \dots, e_{k-1} and by (C1), $v_jP(a_j^-)v_{j+1}$ consists only of the virtual arc e_j . Combining these observations with the fact that $P(f^-)$ can be decomposed as $P(f^-) = P(f^-)v_1P(f^-)v_2 \dots v_lP(f^-)$ we conclude that

$$P/f(s_f) \cap \mathcal{G}(f^-) = P(f^-) - \mathcal{E}(f),$$

and similarly for f^+ . This implies that $P/f(s_f)$ uses no vertex more than once: for vertices in $\mathcal{V}(f)$, this holds by definition, for vertices outside of $\mathcal{V}(f)$, this follows from the fact that $P(f^-)$ and $P(f^+)$ are self-avoiding. Hence $P/f(s_f)$ is self-avoiding. \square

Lemma 4.6. *The pair $A/f = (P/f, C/f)$ defined in Construction 4.4 is an arrangement on $\text{star}(s_f)$.*

Proof. We have already seen that $P/f(s_f)$ is a shape and by definition $C/f(e) = C(e)$ is a configuration, so we only need to verify (D1) – (D3). Let $e \in \text{star}(s_f)$ and assume without loss of generality that $e^- = f^-$ in T . Then by construction

$$P/f(s_f) \cap \mathcal{G}(e) = P(f^-) \cap \mathcal{G}(e) = Q(e) \cap \mathcal{G}(f^-) = Q(e) \cap \mathcal{G}(s_f),$$

so (C1) holds. Furthermore $X(e) = e$ if and only if $P(f^-)$ starts in $\mathcal{V}(e)$ and additionally $X(f) = \bar{f}$ by (D2) and thus $P(f^+)$ starts in $\mathcal{V}(f)$. Then by construction the starting vertices of $P/f(s_f)$ and $P(f^-)$ coincide, so $P/f(s_f)$ starts in $\mathcal{V}(e)$ and (C2) is satisfied. Finally either $X(f) = f$ or $X(f) = \bar{f}$, so again by (D2) at most one arc $e \in \text{star}(s_f)$ can satisfy $X(e) = e$. If there is no such arc, then each $e \in E(X(f)^-)$ satisfies $X(e) = \bar{e}$, so $P(X(f)^-)$ starts with a non-virtual arc, so $P/f(s_f)$ starts with a non-virtual arc and (D2) follows. Finally (C3) and (D3) follow analogously by considering exit directions. \square

We have managed to contract an arrangement on the open tree $\text{star}(f)$ of T to obtain an arrangement on the open star $\text{star}(s_f)$ of T/f . In the next step, we want to do the converse, projecting a given arrangement $A = (P, C)$ on $\text{star}(s_f)$ down to obtain an arrangement $\pi A = (\pi P, \pi C)$ on $\text{star}(f)$, see again Figure 2 for a sketch. For technical reasons we assume that $P(s_f)$ meets $\mathcal{V}(f)$.

Construction 4.7. As one might expect, we choose $\pi C(e) = C(e)$ for $e \in \partial E(f)$. By definition $P(s_f) \cap \mathcal{G}(f)$ is a multi-walk on the adhesion graph $\mathcal{G}(f)$. Each of its walk components consists only of a single vertex v_i . Therefore the walk $P(s_f)$ can be written as a concatenation $q_0 \dots q_k$ of sub-walks q_i , where $q_0 = P(s_f)v_1$, $q_k = v_k P(s_f)$ and $q_j = v_j P(s_f)v_{j+1}$ for $j \in [k-1]$. Note that while q_0 and q_k may be trivial, all other q_j must have length at least 1. By the basic properties of tree-decompositions, the arcs of each (non-trivial) q_j are contained in exactly one of $\mathcal{G}(f^-), \mathcal{G}(f^+)$. Let $a_j \in \{f, \bar{f}\}$ be such that $q_j \in \mathcal{G}(a_j^+)$; in other words, a_j marks on which side of f the part containing q_j lies. For $j \in [k-1]$ let e_j be the (virtual) arc in $\mathcal{E}(a_j)$ connecting v_j and v_{j+1} . We are able to construct shapes $\pi P(s)$ for $s \in \{f^-, f^+\}$ and a configuration $\pi C(f) = (\pi Q(f), \pi X(f), \pi Y(f)) = \pi C(\bar{f})$. The shape $\pi P(s)$ is given as the concatenation $p_0 p_1 \dots p_k$, where

$$p_j = \begin{cases} q_j & \text{if } a_j^+ = s \\ (v_j, e_j, v_{j+1}) & \text{if } j \in [k-1] \text{ and } a_j^+ \neq s \\ (v_j) & \text{if } j \in \{0, k\} \text{ and } a_j^+ \neq s. \end{cases}$$

The walk $\pi Q(f)$ is given as $\pi Q(f) = (v_1, e_1, v_2, \dots, e_{k-1}, v_k)$. Finally, by (D2) there is at most one $e \in E(s_f)$ such that $X(e) = e$. If such an arc exists, we choose $\pi X(f) \in \{f, \bar{f}\}$ such that $\pi X(f)^+ = e^-$. Otherwise $P(s_f)$ starts with a non-virtual arc $e_0 \in \mathcal{E}(s_f)$ and we choose $\pi X(f) \in \{f, \bar{f}\}$ such that $e_0 \in \mathcal{E}(\pi X(f)^+)$. Similarly if there is an arc $e \in E(s_f)$ such that $Y(e) = e$ we choose $\pi Y(f)$ such that $\pi Y(f)^+ = e^-$, otherwise we choose $Y(f)$ such that the last arc of $P(s_f)$ is contained in $\mathcal{E}(\pi Y(f)^+)$.

We now prove that applying Construction 4.7 to an arrangement gives an arrangement.

Lemma 4.8. *The walks $\pi P(f^-)$ and $\pi P(f^+)$ are self-avoiding walks on $\mathcal{G}(f^-)$ and $\mathcal{G}(f^+)$ satisfying $P(s_f) \cap \mathcal{G}(f^-) = \pi P(f^-) - \mathcal{E}(f)$ and $P(s_f) \cap \mathcal{G}(f^+) = \pi P(f^+) - \mathcal{E}(\bar{f})$.*

Proof. By construction $\pi P(f^-)$ consists of all walk-components q_j of $P(s_f)$ contained in $\mathcal{G}(f^-)$, while those q_j not contained in $\mathcal{G}(f^-)$ are replaced by virtual arcs $e_j \in \mathcal{E}(f)$ (shortcuts) connecting the same endpoints. In particular $\pi P(f^-)$ is a walk and $P(s_f) \cap \mathcal{G}(f^-) = \pi P(f^-) - \mathcal{E}(f)$. Furthermore it is self-avoiding because its sequence of vertices also occurs in the SAW $P(s_f)$. The statements for $\pi P(f^+)$ follow analogously. \square

Lemma 4.9. *The map $\pi A = (\pi P, \pi C)$ defined in Construction 4.7 is an arrangement on $\text{star}(f)$.*

Proof. We show that (D1) – (D3) holds for $s = f^-$, the other vertex f^+ is treated analogously. For $e \in E(f^-) \setminus \{f\}$ compatibility follows from the fact that $\pi P(f^-) \cap \mathcal{G}(e) = P(s_f) \cap \mathcal{G}(e)$ and

$$\pi X(e) = e \iff X(e) = e \iff P(s_f) \text{ starts in } \mathcal{V}(e) \iff \pi P(f^-) \text{ starts in } \mathcal{V}(e).$$

and the analogue statement for $\pi Y(e)$. We still need to consider the arc f . By construction $\pi Q(f) \cap \mathcal{G}(f^-)$ is a non-empty multi-walk consisting of all virtual arcs e_j in

$\mathcal{E}(f)$ such that $q_j \in \mathcal{G}(f^+)$, which are exactly the virtual arcs in $\pi P(f^-) \cap \mathcal{G}(f)$. Furthermore $\pi X(f) = f$ holds either if $X(e) = e$ holds for some arc $e \in \mathcal{E}(f^+) \setminus \{\bar{f}\}$ or if q_0 is non-trivial and contained in $\mathcal{G}(f^+)$. In both cases $\pi P(f^-)$ starts at $v_0 \in \mathcal{V}(f)$. Similarly, if $\pi Y(f) = f$, then $\pi P(f^-)$ ends at $v_k \in \mathcal{V}(f)$, so (D1) is satisfied. Again by construction $\pi X(f) = f$ can only hold if $\pi X(e) = \bar{e}$ for all $e \in \mathcal{E}(f^-) \setminus \{\bar{f}\}$, so at most one $e \in \mathcal{E}(f^-)$ satisfies $\pi X(e) \neq e$. If there is no such arc, then $P(s_f)$ starts with a non-virtual arc which is by construction contained in $\mathcal{G}(f^-)$. In particular $\pi P(f^-)$ starts with a non-virtual arc, yielding (D2). Similar arguments for the exit direction πY yield (D3). We conclude that πA is indeed an arrangement on $\text{star}(f)$. \square

Next we show that contraction and projection are inverses of one another. The technical condition that the shape meets $\mathcal{V}(f)$ ensures that all shapes are non-empty. This is crucial for the bijection between self-avoiding walks and arrangements. Without this technical condition, every self-avoiding walk would correspond to not just one, but infinitely many different arrangements, obtained by attaching arbitrarily many boring configurations and suitable (often empty) shapes.

Proposition 4.10. *Contraction of arrangements $A \mapsto A/f$ defines a bijection between arrangements on the open subtree $\text{star}(f)$ of T with respect to \mathcal{T} and arrangements on the open star $\text{star}(s_f)$ centered in the contracted vertex s_f of T/f with respect to \mathcal{T}/f whose shape on s_f meets $\mathcal{V}(f)$. Its inverse map is the projection $A \mapsto \pi A$ of arrangements. Moreover $\|A\| = \|A/f\|$ holds.*

Proof. We have already seen that $A \mapsto A/f$ maps arrangements on $\text{star}(f)$ to arrangements on $\text{star}(s_f)$ and that $A \mapsto \pi A$ maps arrangements on $\text{star}(s_f)$ to arrangements on $\text{star}(f)$. By Remark 4.3 the walk $Q(f)$ of the configuration $C(f)$ is non-empty, so by construction $P/f(s_f)$ intersects $\mathcal{V}(f)$. Furthermore by Lemma 4.5 the set of non-virtual arcs of $P/f(s_f)$ is the disjoint union of the sets of non-virtual arcs of $P(f^-)$ and $P(f^+)$, so $\|A\| = \|A/f\|$ holds.

Let A be an arrangement on $\text{star}(f)$. We claim that $\pi(A/f) = A$. By construction $\pi(A/f)$ is an arrangement on $\text{star}(f)$ such that $\pi(C/f)(e) = C/f(e) = C(e)$ for every $e \in E(f)$. Additionally, by Lemmas 4.5 and 4.8 we have that $\pi(P/f)(f^-)$ satisfies

$$\pi(P/f)(f^-) - \mathcal{E}(f) = P/f(s_f) \cap \mathcal{G}(f^-) = P(f^-) - \mathcal{E}(f).$$

Any virtual arcs of $\pi P/f(f^-)$ in $\mathcal{E}(f)$ must connect the walk-components of $\pi P/f(f^-) - \mathcal{E}(f)$, thus we conclude $\pi(P/f)(f^-) = P(f^-)$. In a similar way we obtain $\pi(P/f)(f^+) = P(f^+)$.

It remains to show that $\pi(C/f)(f) = C(f)$. Equality of the walks $\pi(Q/f)(f)$ and $Q(f)$ follows directly from (C1) of compatibility: Both walks $\pi(Q/f)(f)$ and $Q(f)$ consist of all virtual arcs of the shapes $\pi(P/f)(f^-) = P(f^-)$ and $\pi(P/f)(f^+) = P(f^+)$ contained in $\mathcal{G}(f)$ and also the order of these arcs coincide. Finally, if there is some arc $e \in \partial E(f)$ such that $X(e) = e$, then also $\pi(X/f)(e) = e$ and thus (D2) yields $X(f) = \pi(X/f)(f)$. Otherwise $X(e) = \bar{e}$ holds for all $e \in E(X(f)^+)$, so by (D2) the walk $P(X(f)^+)$ starts with a non-virtual arc e_0 and by (C2) the walk $P(X(f)^-)$ starts in

$\mathcal{V}(f)$. In particular, the first arc of $P/f(s_f)$ is e_0 and by construction $\pi(X/f)(f) = X(f)$. Analogous arguments yield $\pi(Y/f)(f) = Y(f)$.

The proof that $(\pi A)/f = A$ holds for every arrangement $A = (P, C)$ on $\text{star}(s_f)$ such that $P(s_f)$ intersects $\mathcal{V}(e)$ works similarly and is left to the reader. \square

Let S be a finite open sub-tree of T . Starting with a given tree decomposition $\mathcal{T} = (T, \mathcal{V})$ of G the process of edge contraction can be iteratively applied to contract each interior edge of S . In this way S is contracted into an open star centered at a single vertex s_S . We denote the obtained tree decomposition by $\mathcal{T}/S = (T/S, \mathcal{V}/S)$. As before, an arrangement A on S will be contracted to an arrangement A/S on the open star $\text{star}(s_S)$ of T/S .

Remark 4.11. It is not hard to see that neither the contracted tree decomposition $\mathcal{T}/S = (T/S, \mathcal{V}/S)$ nor the contracted arrangement $A/S = (P/S, C/S)$ on $\text{star}(s_S)$ depends on the order of edge contractions: Clearly the tree obtained by consecutive edge contractions does not depend on the order. Additionally $\mathcal{V}/S(s_S)$ is the union of all vertices in $\mathcal{V}(s)$ for $s \in V(S)$ and $\mathcal{V}/S(t) = \mathcal{V}(t)$ for $t \in V(T/S) \setminus \{s_S\}$. In particular \mathcal{T}/S does not depend on the order of contractions of internal edge in S .

The configurations on boundary arcs remain the same after each step of contraction of A , so in particular $C/S(e) = C(e)$ holds for any arc $e \in E(s_S) = \partial E(S)$. Additionally, by Lemma 4.5 the walk $P(s_S)$ consists of all non-virtual arcs of walks $P(s)$ for $s \in V(S)$ and its direction is uniquely defined. Again, this representation of A/S on S does not depend on the order of contractions of edges in S .

We immediately obtain the following corollary of Proposition 4.10.

Corollary 4.12. *Let S be an open subtree of T . The contraction of arrangements $A \mapsto A/S$ defines a bijection between arrangements on S with respect to \mathcal{T} and arrangements on the open star $\text{star}(s_S)$ centered in the contracted vertex s_S of T/S with respect to \mathcal{T}/S whose shape on s_S meets $\mathcal{V}(s)$ for every $s \in V(S)$. Moreover $\|A\| = \|A/S\|$ holds.*

Corollary 4.12 can be used to define a map φ translating any arrangement A on an open subtree S of T to the self-avoiding walk $\varphi(A) = P/S(s_S)$ on the graph $\mathcal{G}(S)$. Observe that by definition the length of $\varphi(A)$ is $\|A\|$. We call $\varphi(A)$ the walk *represented* by A and A a *representation* of $\varphi(A)$. Theorem 4.13 below shows that every SAW has some representation.

Let S be a not necessarily finite subtree of T and w be a SAW on $\mathcal{G}(S)$. The *support* of w in T is the smallest open subtree S' of S such that all arcs of w are contained in $\mathcal{G}(S')$.

Theorem 4.13. *Let T' be an open subtree of T and let w be a SAW on $\mathcal{G}(T')$. Then there is an arrangement $A = (P, C)$ on the support of w in T representing w . Moreover, if w starts in $\mathcal{V}(f)$ for some $f \in \partial E(T')$ and/or ends in $\mathcal{V}(f')$ for some $f' \in \partial E(T')$, we may choose A such that $X(f) = f$ and/or $Y(f') = f'$ holds, respectively.*

Proof. Let S be the support of w in T . By definition w is a shape on the contracted vertex s_S of T/S , thus Lemma 4.1 (i) provides us with an arrangement $A' = (P', C')$

on $\text{star}_{T/S}(s_S)$ such that $P'(s_S) = w$. By minimality of S the shape w meets $\mathcal{V}(s)$ for every $s \in V(S)$. Thus Corollary 4.12 provides us with an arrangement A on S such that $A/S = A'$ and this arrangement satisfies $\varphi(A) = w$. The moreover part follows from Lemma 4.1 (iii). \square

We have just seen that every arrangement A on an open subtree S of T represents a self-avoiding walk $\varphi(A)$ on $\mathcal{G}(S)$ and that every such walk is represented by some arrangement. However, in general there can be several different arrangements representing a single walk. To enumerate SAWs on G , we need each walk to be represented by exactly one arrangement. This can be done by restricting ourselves to complete arrangements, as the upcoming Theorem 4.15 shows. For its proof, we need the following auxiliary result.

Lemma 4.14. *Let $A = (P, C)$ be a complete arrangement on an open subtree S of T and let t be a leaf of S . Then $P(t)$ contains a non-virtual arc.*

Proof. Let t be a leaf of S . If $S = \text{star}(t)$, completeness implies that $X(e) = \bar{e}$ for every $e \in E(S)$. Thus by (D2) the shape $P(t)$ starts with a non-virtual arc.

Suppose now that S is not an open star and let f be the unique arc in $E(t) \setminus \partial E(S)$. By completeness of A we have $X(e) = Y(e) = \bar{e}$ for every $e \in \partial E(S)$. If $X(f) = \bar{f}$ or $Y(f) = \bar{f}$, then again $P(t)$ has to start or end with a non-virtual arc. Thus we may assume that $X(f) = Y(f) = f$. By reducedness $C(f)$ is not boring, so $Q(f)$ has to contain a virtual arc in $\mathcal{E}(\bar{f})$. Then by compatibility (C1) the shape $P(t)$ contains a sub-walk p connecting the two endpoints of this walk. Now $C(e)$ is boring for every $e \in \partial E(S)$ and thus p cannot contain any virtual arcs. We conclude that $P(t)$ contains a non-virtual arc. \square

Theorem 4.15. *Let w be a self-avoiding walk on G . Then there is a unique complete arrangement representing w .*

Proof. We start by constructing a representation of w as in the proof of Theorem 4.15. Let S be the support of w in T and let $S' = \text{star}_{T/S}(s_S)$. Then Lemma 4.1 (i) provides us with an arrangement $A' = (P', C')$ on S' such that $P'(s_S) = w$. Note that we can choose A' to be complete: The walk w starts and ends with a non-virtual arc and $X'(e) = Y'(e) = \bar{e}$ is a valid choice for every $e \in E(S')$. Additionally, all arcs of w are non-virtual arcs of $\mathcal{G}(S) = \mathcal{G}(s_S)$, thus $Q'(e)$ contains only arcs in $\mathcal{E}(\bar{e})$ and $C'(e)$ is boring. Corollary 4.12 provides us with an arrangement A on S such that $A/S = A'$ and thus $\varphi(A) = w$. By Lemma 4.2 whenever $C(e)$ is boring for some arc e of S , $P(s)$ cannot contain any non-virtual arcs for every $s \in V(K_{X(e)}) \cap V(S)$. By minimality of S we conclude $e \in \partial E(S)$. Thus A is complete.

For the uniqueness part, observe that whenever A is a complete arrangement on a subtree S of T representing w , by Lemma 4.14 this subtree must be the support of w . By Corollary 4.12 each A corresponds to a unique arrangement A/S on S' and clearly A/S must be complete. By completeness $X/S(e) = Y/S(e) = \bar{e}$ holds for every $e \in E(S')$ and by Lemma 4.1 (ii) the configuration $C/S(e)$ is uniquely defined by $w = P/S(s_S)$, so A/S is uniquely defined by w . \square

Remark 4.16. For any complete arrangement A on a finite open subtree S of T there is exactly one vertex s_0 of S such that $X(e) = \bar{e}$ for every $e \in E(s_0)$. This can be easily seen by double counting the set $\{(s, e) \mid e^- = s \text{ and } X(e) = \bar{e}\}$ and using (D2) and that configurations on boundary arcs e satisfy $X(e) = \bar{e}$. We call s_0 the *source* of the complete arrangement A . Similarly there is a unique vertex t_0 of S such that $Y(e) = \bar{e}$ for every $e \in E(t_0)$ and this vertex is called the *target* of A .

Lemma 4.17. *Let A be a complete arrangement representing a SAW w on G . The part graph $\mathcal{G}(s_0)$ of the source s_0 of A contains the first arc of w . Similarly, the graph $\mathcal{G}(t_0)$ of the target t_0 contains the final arc of w .*

Proof. Observe that for any $t \in V(T)$ by Construction 4.7 the walk $P(t)$ visits exactly the vertices in $w \cap \mathcal{V}(t)$ in the same order as w . Let t be a vertex of T such that $w^- \in \mathcal{V}(t)$ and let s be the neighbour of t on the unique s_0 - t -path and let $e = st$. Then $X(e) = \bar{e}$, so by (C2) the initial vertex w^- of $P(t)$ must be in $\mathcal{E}(e)$ and thus $w^- \in \mathcal{V}(s)$. Induction yields $w^- \in \mathcal{V}(s_0)$. By (D2) the first arc of $P(s_0)$ is non-virtual and thus coincides with the first arc of w . The statement about the target follows analogously. \square

5 A generating system of equations for configurations

Throughout this section let G be a locally finite connected graph and let Γ be a group acting quasi-transitively on G . We always assume that Γ does not fix an end of G . Corollary 3.2 provides us with a Γ -invariant tree decomposition $(\mathcal{T}, \mathcal{V})$ of G such that there are exactly two Γ -orbits on the arcs of T . More precisely, we can pick an arc e_0 such that every arc e of T either lies in the orbit Γe_0 or in the orbit $\Gamma \bar{e}_0$. We denote by $\rho(e) \in \{e_0, \bar{e}_0\}$ the respective representative of the arc e .

Two non-boring configurations $c = (q, x, y)$ and $c' = (q', x', y')$ on the arcs e and e' of T are Γ -equivalent if $(q', x', y') = (\gamma(q), \gamma(x), \gamma(y))$ holds for some $\gamma \in \Gamma$, where $\gamma(q)$ refers to the natural extension of γ to the virtual edges. Clearly this defines an equivalence relation on the set of configurations on $E(T)$. Let \mathcal{C} be a set of representatives chosen in a way such that each $c \in \mathcal{C}$ is a non-boring configuration on the arc e_0 . Observe that the set \mathcal{C} has finite cardinality because the adhesion graph $\mathcal{G}(e_0)$ is finite and thus carries only finitely many different configurations. For any non-boring configuration c we denote its representative in \mathcal{C} by $\rho(c)$.

Let $c = (q, x, y)$ be a non-boring configuration. A c -completion is an arrangement $A = (P, C)$ on a finite open sub-tree S of K_x containing x such that $C(x) = c$ and $C(e)$ is boring on $e \in E(S)$ if and only if $e \in \partial E(S) \setminus \{x\}$. We define the generating function of c -completions as

$$F_c(z) = \sum_{A \text{ } c\text{-completion}} z^{\|A\|},$$

Let r_c denote the radius of convergence of the generating function $F_c(z)$ and let $R = \min_{c \in \mathcal{C}} r_c$.

The main goal in this section is to find a system of equations for the generating functions of c -completions and study this system to obtain properties of these generating functions.

For an arrangement A on an open subtree S of T and a vertex s of S , we denote by $\text{nb}(A, s)$ the set of arcs $e \in E(s)$ such that A is non-boring on e . Observe that by definition $\text{nb}(A, s)$ is a finite set.

Let $c = (q, x, y)$ be a configuration and let $A = (P, C)$ be a c -completion on the open cone K_x . Then we can decompose A into its restriction A_S on the open star $S = \text{star}(x^-)$ and $C(e)$ -completions A_e on the cones $K_{\bar{e}}$ for all outgoing arcs e of S different from x . Observe that Lemma 4.2 implies that whenever $C(e)$ is boring, the arrangement A_e is uniquely given by $C(e)$ and does not contain any non-virtual arcs. Thus it is possible to recursively build all c -completions by completing arrangements A on $\text{star}(x^-)$ satisfying $C(x) = c$ with $C(e)$ -completions for all $e \in E(x^-) \setminus \{x\}$ carrying non-boring configurations $C(e)$. Thus we obtain

$$F_c(z) = \sum_{\substack{A \text{ arr. on } \text{star}(x^-): \\ C(x)=c}} z^{\|A\|} \prod_{e \in \text{nb}(A, x^-) \setminus \{x\}} F_{C(e)}(z). \quad (3)$$

It follows from the definitions that $F_c(z) = F_{c'}(z)$ whenever c and c' are Γ -equivalent configurations. Write $\mathbf{F}(z) = (F_c(z))_{c \in \mathcal{C}}$. Then the equations above can be rewritten by replacing every $F_c(z)$ by $F_{\rho(c)}(z)$ to obtain

$$F_c(z) = P_c(z, \mathbf{F}(z)), \quad c \in \mathcal{C}. \quad (4)$$

Here

$$P_c(z, \mathbf{y}) = \sum_{\mathbf{n} \in \mathbb{N}_0^{|\mathcal{C}|}} a_{c, \mathbf{n}}(z) \mathbf{y}^{\mathbf{n}}, \quad c \in \mathcal{C},$$

is a formal power series in the variables $\mathbf{y} = (y_c)_{c \in \mathcal{C}}$ whose coefficients are formal power series $a_{c, \mathbf{n}}(z)$ in the variable z given by the sums in equation (3), respectively, where

$$\mathbf{y}^{\mathbf{n}} = \prod_{c \in \mathcal{C}} y_c^{n_c}$$

for an exponent vector $\mathbf{n} = (n_c)_{c \in \mathcal{C}}$.

Let $c \in \mathcal{C}$ be a non-boring configuration. By Lemma 4.14 each c -completion contains at least one non-virtual arc, so that $F_c(0) = 0$. Furthermore we may assume without loss of generality that $F_c(z)$ is not identically zero, that is, that there exists at least one c -completion. If this is not the case for some configuration $c \in \mathcal{C}$, we remove c from \mathcal{C} and replace every occurrence of y_c in $P_c(z, \mathbf{y})$ by 0. Observe that after this process equations (4) remain valid.

Consider the Jacobian matrix

$$\mathfrak{J}(z, \mathbf{y}) = \left(\frac{\partial P_c}{\partial y_{c'}}(z, \mathbf{y}) \right)_{c, c' \in \mathcal{C}}.$$

Let $c = (q, x, y)$ and c' be two non-boring configurations. A c - c' -completion is an arrangement $A = (P, C)$ on a finite open subtree S of K_x containing x such that $C(x) = c$ and there is an arc $f \in \partial E(S)$ such that $\rho(C(f)) = c'$ and $C(e)$ is boring on $e \in E(S)$ if and only if $e \in \partial E(S) \setminus \{x, f\}$. The arc x is the *source arc* of the c - c' -completion A and f is the *target arc*. The length of A is the distance between the source arc x and the target arc f . We denote by $\mathcal{A}(c, c', n)$ the set of c - c' -completions of length n .

Let us now take a closer look at the entries of $\mathfrak{J}(z, \mathbf{F}(z))$. To simplify notation, for any expression $\phi(z, \mathbf{y})$ depending on the variable z and the vector \mathbf{y} , we denote its evaluation at $\mathbf{y} = \mathbf{F}(z)$ by

$$\phi(z) = \phi(z, \mathbf{F}(z)).$$

Lemma 5.1. *Let c, c' be non-boring configurations. Then*

$$(\mathfrak{J}(z))_{c, c'} = \sum_{A \in \mathcal{A}(c, c', 1)} z^{\|A\|}.$$

Proof. Let $c = (q, x, y)$. From equations (3) and (4) we obtain

$$P_c(z, \mathbf{y}) = \sum_{\substack{A \text{ arr. on } \text{star}(x^-): \\ C(x) = c}} z^{\|A\|} \prod_{e \in \text{nb}(A, x^-) \setminus \{x\}} y_{\rho(C(e))}.$$

An application of the Leibniz rule provides

$$\frac{\partial P_c}{\partial y_{c'}}(z, \mathbf{y}) = \sum_{\substack{A \text{ arr. on } \text{star}(x^-): \\ C(x) = c}} z^{\|A\|} \sum_{\substack{f \in \text{nb}(A, x^-) \setminus \{x\}: \\ \rho(C(f)) = c'}} \prod_{e \in \text{nb}(A, x^-) \setminus \{x, f\}} y_{\rho(C(e))}$$

By definition every c - c' -completion A of length 1 consists of an arrangement $A_s = (P_s, C_s)$ on $\text{star}(x^-)$ such that there is an arc $f \in E(x^-) \setminus \{x\}$ with $\rho(C_s(f)) = c'$ and $C_s(e)$ -completions A_e for all $e \in \text{nb}(A, x^-) \setminus \{x, f\}$. Conversely, every such A_s together with some $C_s(e)$ -completions A_e forms a c - c' -completion. Finally note that in this case

$$\|A\| = \|A_s\| + \sum_{e \in \text{nb}(A, x^-) \setminus \{x, f\}} \|A_e\|.$$

We conclude

$$(\mathfrak{J}(z))_{c, c'} = \sum_{A \in \mathcal{A}(c, c', 1)} z^{\|A_s\|} \prod_{e \in \text{nb}(A_s, x^-) \setminus \{x, f\}} F_{\rho(C_s(e))} = \sum_{A \in \mathcal{A}(c, c', 1)} z^{\|A\|}.$$

□

Observe that a c - c'' -completion A_1 and a c'' - c' -completion A_2 can be concatenated to obtain a c - c' -completion whose length is the sum of the lengths of A_1 and A_2 . In particular, the matrix counting c - c' -completions of length n can be obtained as the n -th power of the matrix of c - c' -completions of length 1.

Corollary 5.2. *Let c, c' be non-boring configurations and denote $\mathcal{A}(c, c', n)$ the set of c - c' -completions of length n . Then*

$$(\mathfrak{J}(z)^n)_{c, c'} = \sum_{A \in \mathcal{A}(c, c', n)} z^{\|A\|}.$$

The *dependency digraph* D of the system of equations (4) is given as follows. Its vertex set is the set \mathcal{C} of representatives of configurations and there is an arc from $c = (q, x, y)$ to c' if $\frac{\partial P_c}{\partial y_{c'}}(z)$ is non-zero. In this case we also write $c \rightarrow c'$. Recall that we removed all c such that $F_c(z)$ is identically zero from \mathcal{C} . Therefore we have that $c \rightarrow c'$ if and only if $y_{c'}$ occurs in the power series P_c , that is, if $a_{c', \mathbf{n}}(z) > 0$ for some $\mathbf{n} = (n_c)_{c \in \mathcal{C}}$ with $n_{c'} > 0$. By definition this is the case if and only if there is a c -completion $A = (P, C)$ such that $\rho(C(f)) = c'$ holds for some $f \in E(x^-) \setminus \{x\}$.

The following lemma follows readily from the definitions.

Lemma 5.3. *If $c \rightarrow c'$ for two configuration $c, c' \in \mathcal{C}$, then $r_c \leq r_{c'}$.*

We further classify configurations as follows. A configuration $c = (q, x, y)$ is called an *I-configuration* if $x \neq y$ and it is called a *U-configuration* if $x = y$. We denote the respective subsets of \mathcal{C} by \mathcal{I} and \mathcal{U} , so that $\mathcal{C} = \mathcal{I} \cup \mathcal{U}$ is a disjoint union. We call $c \in \mathcal{I}$ *simple*, if q consists only of a single vertex and denote by \mathcal{I}_s the set of simple configurations. Furthermore, we call $c \in \mathcal{I}$ *persistent*, if there are two simple configurations c_1 and c_2 such that c lies on a walk from c_1 to c_2 in D . Note that this is equivalent to the existence of a c_1 - c_2 -completion such that there is some arc e of the shortest walk from the source arc to the target arc for which $C(e)$ is equivalent to c . Non-persistent I-configurations are called *transient* and we denote by \mathcal{I}_p and \mathcal{I}_t the sets of persistent and transient configurations, respectively. Note in particular that all simple configurations are persistent.

Lemma 5.4. *Let $e, f \in E(T)$. Then there is an arc $f' \in \Gamma f$ such that for every pair of simple configurations $c = (q, x, y), c' = (q', x', y')$ with entry directions $x = e$ and $x' = \bar{f}$ there is a c - c' -completion with target arc f' .*

Proof. Let N be the constant defined in Proposition 3.9. Proposition 3.8 provides an automorphism $\gamma \in \Gamma$ such that e and $f' = \gamma(f)$ are linkable and have distance at least N . Let u be the single vertex of q and v be the single vertex of $\gamma(q')$. Then Proposition 3.9 provides us with a SAW w connecting u and v and meeting $\mathcal{V}(e)$ and $\mathcal{V}(f')$ only in these vertices.

Theorem 4.13 provides us with a representation $A = (P, C)$ of the SAW w on its support S such that $X(e) = e$ and $Y(f') = f'$. It is easy to check that this representation is a c - c' -completion with target arc f' . \square

Remark 5.5. The previous lemma immediately implies that $r_c \leq 1$ holds for the radius of convergence r_c of F_c for every simple configuration $c \in \mathcal{I}_s$. Indeed, the walk represented by a c - c' -completion is also represented by a c -completion, thus $F_c(z)$ contains infinitely many non-zero coefficients.

Lemma 5.6. *The dependency-digraph D satisfies the following conditions.*

- (i) *There are no arcs from U-configurations to I-configurations.*
- (ii) *Each strong component is contained in one of the sets \mathcal{I}_p , \mathcal{I}_t and \mathcal{U} . We call components persistent, transient or U-components depending on the type of configurations they contain.*
- (iii) *The set \mathcal{I}_p is a strong component of D .*

Proof. Let $c = (q, x, y) \in \mathcal{C}$ be a U-configuration and let $A = (P, C)$ be an arrangement on $\text{star}(x^-)$ such that $C(x) = c$. As c is a U-configuration, $x = y \in E(x^-)$, so by (D2) and (D3) in the definition of arrangements $X(f) = Y(f) = \bar{f}$ holds for every $f \in E(x^-) \setminus \{x\}$, thus $C(f)$ is a U-configuration. In particular D cannot contain arcs from U-configurations to I-configurations.

Lemma 5.4 shows that all simple configurations are contained in the same strong component of D . Furthermore, persistent configurations are defined to lie on paths between simple configurations, so all of them are contained in this component, showing that \mathcal{I}_p is contained in a single strong component.

Now let \mathcal{K} be an arbitrary strong component in D . If \mathcal{K} contains a U-configuration, then (i) implies that \mathcal{K} cannot contain any I-configuration, so $\mathcal{K} \subseteq \mathcal{U}$.

Next assume that \mathcal{K} contains a transient configuration. Recall that by definition a transient configuration c cannot lie on a c_1 - c_2 -path in D for two simple configurations c_1, c_2 . Hence, if \mathcal{K} contains a transient configuration, it cannot contain a simple configuration, so in this case $\mathcal{K} \subseteq \mathcal{I}_t$.

Finally, if \mathcal{K} contains neither U-configurations nor transient configurations, then it is contained in \mathcal{I}_p . Since all of \mathcal{I}_p is contained in a single strong component we conclude that in this case $\mathcal{K} = \mathcal{I}_p$, in particular \mathcal{I}_p is a strong component. \square

The submatrix obtained from $\mathfrak{J}(z)$ by restricting the index set to some subset $\mathcal{K} \subset \mathcal{C}$ is denoted by $\mathfrak{J}_{\mathcal{K}}(z)$. It follows directly from Lemma 5.3 that if \mathcal{K} is a strong component of D , the radii of convergence r_c coincide for all $c \in \mathcal{K}$. We denote the common value by $r_{\mathcal{K}}$. Let $c \in \mathcal{K}$ and let $c' \in \mathcal{C}$ such that $c \rightarrow c'$. Then Lemma 5.3 implies that $F_{c'}(z)$ is well-defined and finite for $0 \leq z < r_{\mathcal{K}}$. Hence $\mathfrak{J}_{\mathcal{K}}(z)$ is well defined for $0 \leq z < r_{\mathcal{K}}$. Denote the spectral radius of the matrix $\mathfrak{J}_{\mathcal{K}}(z)$ by $\lambda_{\mathcal{K}}(z)$.

Let us now state some properties of $\mathfrak{J}_{\mathcal{K}}(z)$ and $\lambda_{\mathcal{K}}(z)$ which will be important later on.

Lemma 5.7. *For each strong component \mathcal{K} of D the following statements hold.*

- (i) *All entries of $\mathfrak{J}_{\mathcal{K}}(z)$ are continuous, non-negative, and non-decreasing in $[0, r_{\mathcal{K}})$.*
- (ii) *The spectral radius $\lambda_{\mathcal{K}}(z)$ is continuous and non-decreasing in $[0, r_{\mathcal{K}})$. Additionally, $\lambda_{\mathcal{I}_p}(z)$ is strictly increasing in $[0, r_{\mathcal{I}_p})$.*
- (iii) *$\mathfrak{J}_{\mathcal{K}}(0)$ is nilpotent, in particular $\lambda_{\mathcal{K}}(0) = 0$.*

Proof. By Lemma 5.1, the entries of the matrix $\mathfrak{J}_{\mathcal{K}}(z)$ are generating functions counting c - c' -completions of length 1 for $c, c' \in \mathcal{K}$. In particular they are continuous, non-negative and non-decreasing within their radius of convergence and the same holds for the spectral radius $\lambda_{\mathcal{K}}(z)$.

By Proposition 3.9, there are two persistent configurations c, c' such that $(\mathfrak{J}_{\mathcal{I}_p}(z))_{c, c'}$ is non-constant and thus strictly increasing in z . Irreducibility of $\mathfrak{J}_{\mathcal{I}_p}(z)$ implies that also $\lambda_{\mathcal{I}_p}(z)$ is strictly increasing in z .

To see that (iii) holds, note that by Corollary 3.4 for n large enough, any arcs $e, f \in E(T)$ of distance at least n satisfy $\mathcal{V}(e) \cap \mathcal{V}(f) = \emptyset$. By Corollary 5.2 the entries of $\mathfrak{J}_{\mathcal{K}}(z)^n$ are generating functions counting c - c' -completions A of length n . But the walk represented by such an arrangement A connects $\mathcal{V}(e)$ and $\mathcal{V}(f)$ and thus must be non-trivial. In particular the respective entry of $\mathfrak{J}_{\mathcal{K}}(z)^n$ has no constant term, so $\mathfrak{J}_{\mathcal{K}}(0)^n = 0$. \square

Our aim now is to show that $\mathfrak{J}_{\mathcal{K}}(z)$ has finite entries at $z = r_{\mathcal{K}}$ for every strong component \mathcal{K} .

Lemma 5.8. *For each strong component \mathcal{K} of D , the matrix $\mathfrak{J}_{\mathcal{K}}(r_{\mathcal{K}})$ is well-defined and has finite entries. Moreover, $\lambda_{\mathcal{K}}(z) < 1$ for every $0 \leq z < r_{\mathcal{K}}$ and $\lambda_{\mathcal{K}}(r_{\mathcal{K}}) \leq 1$.*

In order to prove this lemma, we first need the following result.

Lemma 5.9. *Let $A = (a_{ij}) \in \mathbb{R}^{n \times n}$ be an irreducible matrix such that $a_{ij} \geq 0$ for every $i, j \in \{1, 2, \dots, n\}$ and $\lambda(A) \leq 1$. Let also $M = \max\{a_{ij}\}$ and $m = \min\{a_{ij} \mid a_{ij} > 0\}$. Then $M \leq m^{-n}$.*

Proof. Consider some i_0, j_0 such that $a_{i_0 j_0} = M$. By the irreducibility of A , there are $l \leq n$ and indices $z_1 = j_0, z_2, \dots, z_l = i_0$ with $a_{z_i z_{i+1}} > 0$ for every $i = 1, 2, \dots, l-1$. The entry $a_{i_0 i_0}^{(l)}$ of the matrix A^l satisfies

$$a_{i_0 i_0}^{(l)} \geq a_{i_0 j_0} \prod_{i=1}^{l-1} a_{z_i z_{i+1}} \geq M m^{l-1}.$$

Since A^l has non-negative entries and $\lambda(A^l) = \lambda(A)^l \leq 1$, we have by the monotonicity of the spectral radius that

$$1 \geq \lambda(A^l) \geq \lambda(B) = M m^{l-1},$$

where B is the matrix with $b_{i_0 i_0} = M m^{l-1}$, and all other entries 0. In particular we see that $m \leq 1$, so $M \leq m^{-l+1} \leq m^{-n}$. \square

For each strong component \mathcal{K} of D , let $F_{\mathcal{K}}(z) = (F_c(z))_{c \in \mathcal{K}}$. Additionally for $\mathbf{y} = (y_c)_{c \in \mathcal{K}}$ let $P_{c', \mathcal{K}}(z, \mathbf{y})$ be obtained from $P_{c'}(z, (y_c)_{c \in \mathcal{C}})$ by substituting $y_c = F_c(z)$ for each $c \in \mathcal{C} \setminus \mathcal{K}$.

We are now ready to prove Lemma 5.8.

Proof of Lemma 5.8. Let us start by showing that

$$\lambda_{\mathcal{K}}(z) < 1 \text{ for every } 0 \leq z < r_{\mathcal{K}}. \quad (5)$$

We will argue as in [5, Proposition 4].

Assume for a contradiction that $\lambda_{\mathcal{K}}(z_0) \geq 1$ for some $0 \leq z_0 < r_{\mathcal{K}}$. Since $\lambda_{\mathcal{K}}(0) = 0$ and $\lambda_{\mathcal{K}}(z)$ is continuous, there is some $0 < s \leq z_0 < r_{\mathcal{K}}$ such that $\lambda_{\mathcal{K}}(s) = 1$. By the Perron-Frobenius Theorem for non-negative irreducible matrices, $\lambda_{\mathcal{K}}(s) = 1$ is a positive simple eigenvalue of $\mathfrak{J}_{\mathcal{K}}(s)$ for which we can find a left eigenvector $\mathbf{x} = (x_c)_{c \in \mathcal{K}}$ with only positive entries. Define the function

$$\mathcal{F}(z, \mathbf{y}) = \sum_{c \in \mathcal{K}} x_c (y_c - P_{c, \mathcal{K}}(z, \mathbf{y})),$$

where $\mathbf{y} = (y_c)_{c \in \mathcal{K}}$. Note that for $\zeta = (s, F_{\mathcal{K}}(s))$ there is some $\beta \geq 0$ such that

$$\mathcal{F}(\zeta) = 0, \quad \frac{\partial \mathcal{F}}{\partial z}(\zeta) = -\beta, \quad \frac{\partial \mathcal{F}}{\partial y_c}(\zeta) = 0.$$

We claim that β is strictly positive, that is, that there is some $c \in \mathcal{K}$ such that $P_{c, \mathcal{K}}(z, \mathbf{y})$ depends non-trivially on z . Recall that by Lemma 4.14 every c' -completion contains a non-virtual arc, and thus $F_{c'}(z)$ depends on z for every c' . If the strong component \mathcal{K} of D has an outgoing arc cc' , then $P_{c, \mathcal{K}}$ includes $F_{c'}(z)$, which depends on z . Otherwise for $c \in \mathcal{K}$, the function $P_{c, \mathcal{K}}(z, \mathbf{y}) = P_c(z, \mathbf{y})$ does not depend on any $y_{c'}$ for $c' \in \mathcal{C} \setminus \mathcal{K}$ and thus must depend on z , because $F_c(z) = P_c(z, \mathbf{F}(z))$ does.

Taylor expanding around ζ , replacing \mathbf{y} by $F_{\mathcal{K}}(z)$, and using that $F_c(z) = P_{c, \mathcal{K}}(z, F_{\mathcal{K}}(z))$ we obtain

$$\beta(s - z) = O((s - z)^2),$$

which is a contradiction. This proves (5).

We will now show that $\mathfrak{J}_{\mathcal{K}}(r_{\mathcal{K}})$ has finite entries. Let us write $m(z)$ for the smallest positive entry of $\mathfrak{J}_{\mathcal{K}}(z)$, and consider some $0 < z_1 < r_{\mathcal{K}}$. Then $\mathfrak{J}_{c, c'}(z) \leq m(z)^{-|\mathcal{K}|} \leq m(z_1)^{-|\mathcal{K}|}$ for every $c, c' \in \mathcal{K}$ and every $z_1 \leq z < r_{\mathcal{K}}$ by Lemma 5.9 and the monotonicity of $m(z)$. It follows from the monotone convergence theorem that $\mathfrak{J}_{c, c'}(r_{\mathcal{K}})$ is well-defined and $\mathfrak{J}_{c, c'}(r_{\mathcal{K}}) \leq m(z_1)^{-|\mathcal{K}|}$ for every $c, c' \in \mathcal{K}$, which proves the first assertion of the lemma.

For the second assertion, it remains to show that $\lambda_{\mathcal{K}}(r_{\mathcal{K}}) \leq 1$. This follows from (5) by taking the limit as z goes to $r_{\mathcal{K}}$. \square

6 Analyticity at the critical point

Keeping all definitions and assumptions from the previous section, the goal in this section is to prove the following analyticity results; recall that $R = \min_{c \in \mathcal{C}} r_c$ is the minimal radius of convergence of any $F_c(z)$.

Proposition 6.1. *For every $c \in \mathcal{U}$, $P_c(z, \mathbf{y})$ is analytic at $(z, \mathbf{y}) = (R, \mathbf{F}(R))$.*

Proposition 6.2. *We have that $\mathfrak{J}_{\mathcal{I}}(z, \mathbf{y})$ is analytic at $(z, \mathbf{y}) = (R, \mathbf{F}(R))$.*

We will prove Propositions 6.1 and 6.2 in a series of lemmas. We start by showing that $F_c(R)$ is finite for every $c \in \mathcal{U}$.

Lemma 6.3. *If T has more than two ends and Γ does not fix an end of T , then $F_c(R)$ is well-defined and finite for every $c \in \mathcal{U}$.*

Proof. First note that every $c \in \mathcal{U}$ is by definition non-boring and thus contains a virtual arc in $\mathcal{E}(\bar{x})$. If c contains $k \geq 2$ virtual arcs $e_1, \dots, e_k \in \mathcal{E}(\bar{x})$, then $F_c(z) \leq \prod_{i=1}^k F_{c_i}(z)$, where $c_i = (q_i, x, y)$ is the configuration such that q_i is the walk of length 1 with the single arc e_i . It thus suffices to prove the assertion for every $c = (q, x, y) \in \mathcal{U}$ that contains exactly two vertices and a virtual arc in $\mathcal{E}(\bar{x})$ connecting those vertices. Let c be such a configuration.

We claim that there are an open subtree S of T , an arrangement $A = (P, C)$ on S and three arcs $f, e_1, e_2 \in \partial E(S)$ satisfying the following conditions:

- (i) $C(f) = c$ and $X(f) = \bar{f}$.
- (ii) $C(e_1)$ and $C(e_2)$ are simple I-configurations such that $X(e_1) = e_1$ and $Y(e_2) = e_2$.
- (iii) $C(e')$ is boring if and only if $e' \in \partial E(S) \setminus \{f, e_1, e_2\}$.

If such an arrangement A exists, then A together with a c -completion A' forms a $C(e_1)$ – $C(e_2)$ -completion, whose length n is the distance of e_1 and e_2 . In particular

$$(\mathfrak{J}_p(R)^n)_{C(e_1), C(e_2)} \geq R^{\|A\|} F_c(R),$$

so Lemma 5.8 implies that $F_c(R)$ is finite.

For the proof of the claim we start by picking $f = \bar{x}$. Furthermore let e_1, e_2 be two different arcs having the same distance to f . We pick e_1 and e_2 such that the sets $\mathcal{V}(f)$, $\mathcal{V}(e_1)$ and $\mathcal{V}(e_2)$ are pairwise disjoint.

Let $q = (u, e, v) \in \mathcal{V}(f)$. As in the proof of Proposition 3.9 we can find two disjoint $\mathcal{V}(f)$ – $\mathcal{V}(e_2)$ -paths π_u and π_v in G , starting at u and v , respectively. Moreover, we find a $\mathcal{V}(f)$ – $\mathcal{V}(e_1)$ -path π in G starting at u . Let u' be the endpoint of π in $\mathcal{V}(e_1)$. We construct a walk w on $\mathcal{G}(K_f)$ as follows. Start at u' and follow π until its first intersection with one of the paths π_u and π_v . If we reached π_u , we follow it until u , add the arc e to reach v and follow π_v from v to $\mathcal{V}(e_2)$. Otherwise we follow π_v until v , add \bar{e} to reach u and follow π_u to reach $\mathcal{V}(e_2)$. In this case we reverse the obtained walk and exchange e_1 and e_2 . It is now easy to check that in both cases we end up with a SAW w on $\mathcal{G}(K_f)$ connecting $\mathcal{V}(e_1)$ and $\mathcal{V}(e_2)$ and containing the arc e . Let S be the support of w . By construction the arcs e_1, e_2 and f are in $\partial E(S)$. Theorem 4.13 provides us with a representation $A = (P, C)$ of w such that $X(e_1) = e_1$ and $Y(e_2) = e_2$. It is easy to check that A satisfies conditions (i)–(iii). \square

The following construction is essential in the upcoming proofs; it is sketched in Figure 3.




Figure 3: Reflection-extension with splitting point u . The last part of the original walk is drawn in gray, the modified walk is drawn in black. Note that the distance between u' and $g_e(u')$ is bounded by an absolute constant, so the increase in length will become negligible when the length of the original walk is large.

Construction 6.4. Let $a \in \{e_0, \bar{e}_0\}$ be a representative arc of the action of Γ on $E(T)$ and fix an odd integer N which is larger than or equal to the constant of Proposition 3.9. Then Proposition 3.8 provides us with an automorphism γ_a such that a and $\gamma_a(a)$ are linkable and have distance N . By Proposition 3.9 for any $v \in \mathcal{V}(a)$ there is a v - $\gamma_a v$ -path $\pi_{a,v}$ in G meeting $\mathcal{V}(a) \cup \mathcal{V}(\gamma_a a)$ only in its endpoints.

Moreover, for every $e \in E(T)$ fix an automorphism h_e mapping e to its representative $\rho(e) \in \{e_0, \bar{e}_0\}$.

Let $e \in E(T)$ and w be a self-avoiding walk on $\mathcal{G}(K_e)$ not intersecting $\mathcal{V}(e)$. A *reflection-extension* of w through e is constructed as follows. Let u be a vertex of w such that $d_G(u, \mathcal{V}(e)) = d_G(w, \mathcal{V}(e))$. We apply $g_e := h_e^{-1} \circ \gamma_{h_e(\bar{e})} \circ h_e$ to the sub-walk of w starting at u . Next we connect u to $\mathcal{V}(e)$ with a geodesic in G , which by definition intersects $\mathcal{V}(e)$ only in a single vertex u' . Furthermore we connect $g_e(u)$ to $\mathcal{V}(g_e(e))$ with the image of the latter geodesic under g_e . Finally, we connect u' to $g_e(u')$ with $h_e^{-1}(\pi_{h_e(\bar{e}), h_e(u')})$. Note that by our choice of the vertex u , the object we thus obtain is a SAW; we denote by $\text{Refl}_e(w)$ the set of all possible reflection-extensions of w through e . The vertex u is called *splitting point* of the reflection-extension.

Lemma 6.5. *Let $e \in E(T)$ and let w_1 and w_2 be different self-avoiding walks on $\mathcal{G}(K_e)$. Then $\text{Refl}_e(w_1) \cap \text{Refl}_e(w_2) = \emptyset$.*

Proof. We show that the walk w can be uniquely reconstructed from any $w' \in \text{Refl}_e(w)$ provided that e is known. Note that by construction w' consists of 5 parts: the sub-walk of w up to the vertex u , followed by a geodesic π_u from u to $u' \in \mathcal{V}(e)$, a walk from u' to $g_e(u')$, the image of π_u under g_e , and the image of the sub-walk of w starting at u under g_e . In particular, given w' it is easy to identify u as the first vertex of w' such that $g_e(u)$ is also contained in w' . Then w is obtained as the concatenation of the sub-walk of w' up to u and the image of the sub-walk of w' starting at $g_e(u)$ under the map g_e^{-1} . \square

Our next goal is to show that there are exponentially fewer SAWs that stay within a part than SAWs that use vertices in multiple parts. Consider some part $\mathcal{V}(t)$ and some vertex $o \in \mathcal{V}(t)$. We define $c_n(o, t)$ to be the number of SAWs w on G of length n starting at o visiting only vertices in $\mathcal{V}(t)$. Let $\mu_{o,t} = \limsup c_n(o, t)^{1/n}$.

Lemma 6.6. *If Γ does not fix an end of T , then $\mu_{o,t} < 1/R$ for every $t \in V(T)$ and $o \in \mathcal{V}(t)$.*

Proof. The statement is trivially satisfied if $\mathcal{V}(t)$ is finite, so we henceforth assume that $\mathcal{V}(t)$ is infinite. Consider a SAW w of length n starting at o and visiting only vertices in $\mathcal{V}(t)$. We will construct several SAWs starting at o that do not stay in $\mathcal{V}(t)$ by applying Construction 6.4 and then compare the corresponding generating functions.

We first need to define a suitable set of adhesion sets. Let $M > 0$ be the constant of Proposition 3.7. Let also $\mathcal{B}(w)$ be a set of arcs of $E(t)$ such that $1 \leq d_G(\mathcal{V}(e), w) \leq M$ for every $e \in \mathcal{B}(w)$, $d_G(\mathcal{V}(e), \mathcal{V}(f)) > 2M$ for every distinct $e, f \in \mathcal{B}(w)$, and $\mathcal{B}(w)$ is maximal with respect to the latter property.

We now associate to each $e \in \mathcal{B}(w)$ a vertex $u_e = u_e(w)$ of w such that $d_G(u_e, \mathcal{V}(e)) = d_G(w, \mathcal{V}(e)) \leq M$, and we denote the set thus obtained by U . Note that the vertices u_e are distinct because $d_G(\mathcal{V}(e), \mathcal{V}(f)) > 2M$ for every distinct $e, f \in \mathcal{B}(w)$. Using the ordering of the vertices u_e coming from the ordering of the vertices of w , we order the elements of $\mathcal{B}(w)$.

We claim that $|\mathcal{B}(w)| \geq m := n/\Delta^{4M+2D}$, where Δ is the maximal vertex-degree in G and D is maximal distance (in G) between two vertices in the same adhesion set. Indeed, let v be a vertex of w , and $f \in E(t)$ such that $d_G(v, \mathcal{V}(f)) \leq M$. Then

$$d_G \left(\mathcal{V}(f), \bigcup_{e \in \mathcal{B}(w)} \mathcal{V}(e) \right) \leq 2M$$

by the maximality of $\mathcal{A}(w)$. Hence $d_G(v, U) \leq 4M + 2D$. This implies that the balls of radius $4M + 2D$ around the vertices in U cover w . Since each ball of radius $4M + 2D$ has size at most Δ^{4M+2D} , the claim follows.

Let $\varepsilon \in (0, 1)$ be a small constant to be defined later, and consider a set $H \subseteq \mathcal{B}(w)$ of cardinality $k = \lfloor \varepsilon m \rfloor$. We write $e_i = e_i(H)$ for the i th element of H . We will successively apply Construction 6.4 to w . More precisely, let $w_1(H) \in \text{Refl}_{e_1}(w)$ be a reflection-extension of w through e_1 with splitting point u_{e_1} and let $g = g_{e_1}$ be the automorphism applied to the second part of w in the reflection-extension process.

Consider the second element $e_2 \in H$, and note that $g(u_{e_2})$ minimizes $d_G(v, \mathcal{V}(g(e_2)))$ among the vertices v of $w_1(H)$, since

$$d_G(\mathcal{V}(g(e_1)), \mathcal{V}(g(e_2))) > 2M.$$

Thus there is a reflection-extension $w_2(H) \in \text{Refl}_{g(e_2)}(w_1(H))$ of $w_1(H)$ through $g(e_2)$ with splitting point $g(u_{e_2})$. Continuing in this way, we obtain a sequence of SAWs $w_1(H), \dots, w_k(H)$. Note that $w_k(H)$ has length between n and $n + k\ell$ for some constant $\ell > 0$ independent of w .

Varying H over all possible subsets of $\mathcal{B}(w)$ of cardinality k , we obtain a map $(w, H) \mapsto w_k(H)$. We claim that this map is injective. Indeed, note first that the set H can be reconstructed uniquely from $w_k(H)$. The arc $e_1(H)$ is the unique arc in $E(t)$ such that the configuration defined by $w_k(H)$ on $e_1(H)$ is an I-configuration. This also defines uniquely the map $g = g_{e_1}$ used in the reflection-extension process. Similarly $g(e_2(H))$ is the unique arc in $E(g(t)) \setminus \{e_1(H)\}$ such that the configuration defined by $w_k(H)$ on $g(e_2(H))$ is an I-configuration. Proceeding in this way, we see that H is uniquely given by $w_k(H)$. But then the claim is a consequence of Lemma 6.5.

We have thus constructed at least

$$\binom{m}{k} c_n(o, t)$$

distinct SAWs that start at o and have length between n and $n + k\ell$. Since $R \leq 1$ by Remark 5.5, this implies that

$$R^{n+k\ell} \binom{m}{k} c_n(o, t) \leq \sum_{(w,H)} R^{\|w_k(H)\|},$$

where the sum runs over all SAWs w of length n starting at o and all subsets H of $\mathcal{B}(w)$ of cardinality k . We will now estimate this sum in terms of $\mathfrak{J}_p(R)$.

Each walk $w_k(H)$ is a self-avoiding walk starting at $\mathcal{V}(t)$, and ending at $\mathcal{V}(t')$, where $t' = (g_k \circ \dots \circ g_1)(t)$. Additionally, by construction $w_k(H)$ meets each of $e_1(H)$ and $e'_k(H) := (g_k \circ \dots \circ g_1)(e_k(H))$ only in a single vertex. Decomposing $w_k(H)$ at these two vertices, we see that it consists of a SAW $w_k^1(H)$ containing only vertices of $\mathcal{V}(t)$, a SAW $w_k^2(H)$ connecting $\mathcal{V}(e_1(H))$ and $\mathcal{V}(e'_k(H))$ and a SAW $w_k^3(H)$ containing only vertices of $\mathcal{V}(t')$. By Theorem 4.13 there is a representation $A_H = (P_H, C_H)$ of $w_k^2(H)$ which is a $C_H(e_1(H))$ - $C_H(e'_k(H))$ -completion for the simple configurations $C_H(e_1(H))$ and $C_H(e'_k(H))$. By Construction 6.4 the arcs $e_1(H)$ and $e'_k(H)$ have distance $kN + k - 1$. Thus the arrangement A_H is counted in some entry of $\mathfrak{J}_p(R)^{kN+k-1}$.

Letting $c_n^t = \max_{o \in \mathcal{V}(t)} c_n(o, t)$, we thus obtain

$$\sum_{(w,H)} R^{\|w_k(H)\|} \leq \|\mathfrak{J}_p(R)\|_1^{kN+k-1} \sum_{i=0}^n \sum_{j=0}^{n-i} c_i^t c_j^t R^{i+j}.$$

To estimate the latter sum, let $\mu_t = \max_{v \in \mathcal{V}(t)} \mu_{v,t}$. By taking cases according to whether $\mu_t R \leq 1$ or $\mu_t R > 1$ we obtain

$$\sum_{i=0}^n \sum_{j=0}^{n-i} c_i^t c_j^t R^{i+j} = \max\{\mu_t^n R^n, 1\} e^{o(n)}.$$

This implies that

$$c_n(o, t) \leq \frac{\max\{\mu_t^n, 1/R^n\}}{R^{k\ell} \binom{m}{k}} \|\mathfrak{J}_p(R)\|_1^{k(N+1)+o(n)}.$$

Recall that $k = \lfloor \varepsilon m \rfloor$. Hence by Stirling's approximation we obtain

$$\binom{m}{k} = \sqrt{\frac{m}{2\pi k(m-k)}} \frac{m^m}{k^k (m-k)^{m-k}} (1 + o(1)) = \left(\frac{1}{\varepsilon^\varepsilon (1-\varepsilon)^{1-\varepsilon}} \right)^{m+o(m)}. \quad (6)$$

For $\varepsilon = R^\ell \|\mathfrak{J}_p(R)\|_1^{-(N+1)}$ we have

$$c_n(o, t) \leq \max\{\mu_t^n, 1/R^n\} (1-\varepsilon)^{(1-\varepsilon)m+o(m)}.$$

Taking n -th roots and sending n to infinity we obtain that

$$\mu_{o,t} < \max\{\mu_t, 1/R\}.$$

Since this holds for all o , we obtain that $\mu_t < \max\{\mu_t, 1/R\}$, hence $\mu_t < 1/R$, as desired. \square

We are now ready to prove Proposition 6.1.

Proof of Proposition 6.1. By assumption all coefficients of $P_c(z, \mathbf{y})$ are non-negative. Thus it is enough to show that there is some constant $\delta > 0$ such that

$$F_{c,\delta}(R) := P_c((1+\delta)R, (1+\delta)\mathbf{F}(R)) < \infty.$$

As in the proof of Lemma 6.3, it suffices to consider configurations $c \in \mathcal{U}$ that contain only two vertices and a virtual arc connecting those vertices.

Let $c = (q, x, y)$ be such a configuration, and let $t = x^-$. Let us start by defining $\mathcal{A}(c, n)$ to be the set of arrangements $A = (P, C)$ on $S = \text{star}(t)$ such that $C(x) = c$ and $|P(t)| = n$, where here n counts the total number of arcs in $P(t)$ and not just the non-virtual ones. In particular every $A \in \mathcal{A}(c, n)$ satisfies

$$\|A\| + |\text{nb}(A, t) \setminus \{x\}| \leq n \leq \|A\| + K |\text{nb}(A, t) \setminus \{x\}|, \quad (7)$$

where K is the size of the adhesion sets of our tree decomposition. With this in mind, we can write

$$F_{c,\delta}(R) \leq \sum_{n=0}^{\infty} (1+\delta)^n \Phi_n(R),$$

where

$$\Phi_n(R) = \sum_{A \in \mathcal{A}(c, n)} R^{\|A\|} \prod_{e \in \text{nb}(A, t) \setminus \{x\}} F_{C(e)}(R).$$

It is enough find an upper bound for $\Phi_n(R)$ which decays exponentially in n because then for some $\delta > 0$ small enough, $F_{c,\delta}(R)$ is upper bounded by a geometric series. We split the inner sum into three parts and treat those parts individually. In order to split up the sum, we first need some definitions.

For $\varepsilon > 0$ we partition $\mathcal{A}(c, n)$ into two subsets according to whether they result in a 'small' or 'large' number of non-boring configurations on the edges of $E(S)$:

- (i) $\mathcal{S}(c, \varepsilon, n)$ is the set of $A \in \mathcal{A}(c, n)$ such that the number of arcs $e \in E(S) \setminus \{x\}$ for which $C(e)$ is non-boring, is smaller than εn ,
- (ii) $\mathcal{L}(c, \varepsilon, n) = \mathcal{A}(c, n) \setminus \mathcal{S}(c, \varepsilon, n)$.

Clearly,

$$\Phi_n(R) = \sum_{A \in \mathcal{S}(c, \varepsilon, n)} R^{\|A\|} \prod_{e \in \text{nb}(A, t) \setminus \{x\}} F_{C(e)}(R) + \sum_{A \in \mathcal{L}(c, \varepsilon, n)} R^{\|A\|} \prod_{e \in \text{nb}(A, t) \setminus \{x\}} F_{C(e)}(R). \quad (8)$$

We further split the second sum into a part summing over completions where ‘many’ of the $C(e)$ -completions are longer than some constant L , and a part summing over completions where ‘few’ of the $C(e)$ -completions are longer than L . More precisely, for $L > 0$ and $A \in \mathcal{L}(c, \varepsilon, n)$, let $\mathcal{M}(A, \varepsilon, n, L)$ be the set of sequences $(A_e)_{e \in E(S) \setminus \{x\}}$, where each $A_e = (P_e, C_e)$ is an arrangement on $K_{\bar{e}}$, such that $C_e(e) = C(e)$ for every $e \in E(S) \setminus \{x\}$, and for at least $\varepsilon n/2$ arcs $e \in E(S) \setminus \{x\}$, we have $\|A_e\| \geq L$. Let $\mathcal{F}(A, \varepsilon, n, L)$ be the set of sequences $(A_e)_{e \in E(S) \setminus \{x\}}$ not contained in $\mathcal{M}(A, \varepsilon, n, L)$. Write $\|(A_e)\| = \sum_{e \in E(S) \setminus \{x\}} \|A_e\|$.

Recall that

$$F_{C(e)}(R) = \sum_{A_e \text{ } C(e)\text{-completion}} R^{\|A_e\|}.$$

Using this identity on the factors of the product inside the second sum in the expression (8), we arrive at $\Phi_n(R) = \Sigma_n(R) + \Lambda_n^1(R) + \Lambda_n^2(R)$, where

$$\begin{aligned} \Sigma_n(R) &= \sum_{A \in \mathcal{S}(c, \varepsilon, n)} R^{\|A\|} \prod_{e \in \text{nb}(A, t) \setminus \{x\}} F_{C(e)}(R) \\ \Lambda_n^1(R) &= \sum_{A \in \mathcal{L}(c, \varepsilon, n)} R^{\|A\|} \sum_{(A_e) \in \mathcal{M}(A, \varepsilon, n, L)} R^{\|(A_e)\|} \\ \Lambda_n^2(R) &= \sum_{A \in \mathcal{L}(c, \varepsilon, n)} R^{\|A\|} \sum_{(A_e) \in \mathcal{F}(A, \varepsilon, n, L)} R^{\|(A_e)\|}. \end{aligned}$$

Let us first consider $\Sigma_n(R)$. Using the upper bound in (7) together with the fact that $R \leq 1$ by Remark 5.5, we immediately obtain

$$\Sigma_n(R) \leq t_1^{\varepsilon n} R^{(1-K\varepsilon)n} |\mathcal{S}(c, \varepsilon, n)|,$$

where $1 < t_1 := 1 + \max\{F_c(R) \mid c \in \mathcal{U}\}$ is a finite constant by Lemma 6.3. We will show that the right-hand side of the inequality decays exponentially, provided that ε is small enough.

Let us estimate the cardinality of $\mathcal{S}(c, \varepsilon, n)$. Each $A = (P, C) \in \mathcal{S}(c, \varepsilon, n)$ consists of a walk $P(t)$ and a collection of configurations $(C(e))_{e \in E(S)}$ with all but at most εn of them being boring. Note that $P(t)$ already uniquely determines A because c is a U-configuration. Indeed, the configurations $C(e) = (Q(e), X(e), Y(e))$ must be U-configurations and by Lemma 4.1 each $Q(e)$ is determined by $P(t)$. We derive an upper bound for the number of possible walks $P(t)$. We start by counting the possible walks

with exactly k virtual arcs, and whose walk components after removing the virtual arcs have lengths (n_0, \dots, n_k) . It is possible that some $n_i = 0$. Note that $\sum_{i=0}^k n_i = n - k$. Denote the set of such walks by $\mathcal{W}(n_0, \dots, n_k)$. We can construct such walks inductively. Start by choosing a walk W_0 of length n_0 starting at the initial vertex of the walk q of the configuration c . In the i -th step we attach a virtual arc e_i and a walk W_i of length n_i to the walk constructed so far; if no suitable virtual arcs exists, we do not proceed with the construction. There are at most $c_{n_i}(t)$ ways to choose the i -th walk, and there is a constant upper bound t_2 for the number of virtual arcs incident to a vertex. Thus

$$|\mathcal{W}(n_0, \dots, n_k)| \leq t_2^k \prod_{i=0}^k c_{n_i}(t),$$

where $c_{n_i}(t) = 1$ if $n_i = 0$. To bound this product, consider some $r \in (\mu_t, 1/R)$, where $\mu_t = \max_{o \in \mathcal{V}(t)} \mu_{o,t}$. Note that such an r exists by Lemma 6.6 and that there exists a constant $t_3 \geq 1$ such that $c_m(t) \leq t_3 r^m$ for every $m \geq 1$ and every $o \in \mathcal{V}(t)$. Moreover, every virtual arc lies in an adhesion set such that A is non-boring on the corresponding arc. This implies that $k \leq K\varepsilon n$. We conclude that

$$|\mathcal{W}(n_0, \dots, n_k)| \leq (t_2 t_3)^{K\varepsilon n} r^{n-k}.$$

Every $A \in \mathcal{S}(c, \varepsilon, n)$ is uniquely determined by some $W \in \mathcal{W}(n_0, \dots, n_k)$ for some $k \leq K\varepsilon n$ where n_0, \dots, n_k are given by the positions of the k virtual arcs in $P(t)$. We conclude that

$$|\mathcal{S}(c, \varepsilon, n)| \leq \sum_{k=0}^{\lfloor K\varepsilon n \rfloor} \binom{n}{k} (t_2 t_3)^{K\varepsilon n} r^{n-k}.$$

Using Stirling's approximation as in (6) and the fact that the binomial coefficients $\binom{n}{k}$ are monotonic in k , we see that

$$\binom{n}{k} \leq \left(\frac{1}{(K\varepsilon)^{K\varepsilon} (1 - K\varepsilon)^{1 - K\varepsilon}} \right)^{n+o(n)}$$

for every $0 \leq k \leq K\varepsilon n$ and every $0 < \varepsilon \leq \frac{1}{2K}$. Therefore, for some positive constant t_4 depending on t_1, t_2, t_3 and K ,

$$\Sigma_n(R) \leq \sum_{k=0}^{\lfloor K\varepsilon n \rfloor} R^{(1-K\varepsilon)n} t_4^{\varepsilon n} r^{n-k} \left(\frac{1}{(K\varepsilon)^{K\varepsilon} (1 - K\varepsilon)^{1 - K\varepsilon}} \right)^{n+o(n)}.$$

It is possible to choose $0 < \varepsilon \leq \frac{1}{2K}$ to be small enough so that

$$\frac{2^\varepsilon t_4^\varepsilon r^{1-k/n}}{R^{K\varepsilon} (K\varepsilon)^{K\varepsilon} (1 - K\varepsilon)^{1 - K\varepsilon}} \leq \frac{1}{R}$$

holds for n large enough because the left side goes to r as ε tends to 0. Hence, we obtain

$$\Sigma_n(R) \leq (K\varepsilon n + 1) R^{(1-K\varepsilon)n} 2^{-\varepsilon n} R^{K\varepsilon n} R^{-n+o(n)} \leq 2^{-\varepsilon n+o(n)},$$

as desired.

Let us next bound $\Lambda_n^1(R)$. Our aim is to estimate $\sum_{(A_e) \in \mathcal{M}(A, \varepsilon, n, L_0)} R^{\|(A_e)\|}$ for a certain length L_0 . Let $p = 2^{-4/\varepsilon}$ and note that there exist a constant $L_0 > 0$ such that

$$\sum_{\substack{A \text{ } c' \text{-completion} \\ \|A\| \geq L_0}} R^{\|A\|} \leq p F_{c'}(R).$$

for every $c' \in \mathcal{U}$.

For a fixed $H \subseteq \text{nb}(A, t) \setminus \{x\}$ with $|H| \geq \varepsilon n/2$, we have

$$\sum_{\substack{(A_e) \in \mathcal{M}(A, \varepsilon, n, L_0) \\ \|A_e\| \geq L_0 \text{ for } e \in H}} R^{\|(A_e)\|} \leq p^{\varepsilon n/2} \prod_{e \in \text{nb}(A, t) \setminus \{x\}} F_{C(e)}(R).$$

Each $(A_e) \in \mathcal{M}(A, \varepsilon, n, L_0)$ by definition satisfies $\|A_e\| \geq L_0$ for at least $\varepsilon n/2$ arcs $e \in E(S) \setminus \{x\}$. Summing over all possible H yields

$$\sum_{(A_e) \in \mathcal{M}(A, \varepsilon, n, L_0)} R^{\|(A_e)\|} \leq 2^n p^{\varepsilon n/2} \prod_{e \in \text{nb}(A, t) \setminus \{x\}} F_{C(e)}(R).$$

It follows that

$$\Lambda_n^1(R) \leq 2^n p^{\varepsilon n/2} \sum_{A \in \mathcal{L}(c, \varepsilon, n)} R^{\|A\|} \prod_{e \in \text{nb}(A, t) \setminus \{x\}} F_{C(e)}(R) \leq 2^n p^{\varepsilon n/2} F_c(R) \leq 2^{-n} F_c(R).$$

Finally we find an upper bound for $\Lambda_n^2(R)$. To this end, we employ a similar strategy as in Lemma 6.6.

For $e_0 \in E(T)$, Proposition 3.8 yields an arc $f_0 \in E(K_{e_0})$ in the orbit of \bar{e}_0 such that e_0 and f_0 are linkable. Let $M = d_G(\mathcal{V}(e_0), \mathcal{V}(f_0))$ and let $M' = d_T(e_0, f_0)$. We choose f_0 such that that $M > L_0$.

For a c -completion $A = (P, C)$, we denote by A_S and A_e the restrictions of A to S and $K_{\bar{e}}$, respectively. Assume that we are given a c -completion A such that $A_S \in \mathcal{L}(c, \varepsilon, n)$ and $(A_e)_{e \in E(S) \setminus \{x\}} \in \mathcal{F}(A, \varepsilon, n, L_0)$ and let $w = \varphi(A)$ be the walk represented by A . Consider the set of arcs $e \in \text{nb}(A, t) \setminus \{x\}$ such that $\|A_e\| < L_0$. As in the proof of Lemma 6.6 we can find a subset $\mathcal{B}(w)$ of these arcs of size at least $m := \lfloor \varepsilon' n \rfloor$ for some $0 < \varepsilon' < \varepsilon/2$ such that $d_G(\mathcal{V}(e), \mathcal{V}(f)) > 2M + 2D$ for distinct $e, f \in \mathcal{B}(w)$, where D is the diameter of the adhesion sets. Indeed, at least $\varepsilon n/2$ vertices of the walk w lie in adhesion sets $\mathcal{V}(e)$ for $e \in \text{nb}(A, t) \setminus \{x\}$ such that $\|A_e\| < L_0$, and the $(2M + 2D)$ -ball around each such vertex covers only a constant number of other vertices.

Let $\varepsilon'' \in (0, 1)$ be a constant to be defined. Consider a subset H of $\mathcal{B}(w)$ of cardinality $k := \lfloor \varepsilon'' m \rfloor$. Each $e \in H$ either lies in the orbit of e_0 or in the orbit of f_0 . If there is $\gamma \in \Gamma$ such that $\gamma e_0 = e$, then we choose $f(e, w) = \gamma f_0$. Otherwise there is $\gamma \in \Gamma$ such that $\gamma f_0 = e$ and we choose $f(e, w) = \gamma e_0$. Note that $M > L_0$ and the choice of e such that $\|A_e\| < L_0$ implies that $0 < d_G(w, \mathcal{V}(f(e, w))) < M + D$. Choose a vertex $u_e = u_e(w)$ of w such that $d_G(u_e, \mathcal{V}(f(e, w))) = d_G(w, \mathcal{V}(f(e, w)))$. The walk w induces an order on

the vertices $\{u_e \mid e \in H\}$, denote them by u_1, \dots, u_k . This order induces an order on H , denoted by e_1, \dots, e_k , and consequently also on the arcs $\{f(e, w) \mid e \in H\}$ which we denote by f_1, \dots, f_k . As in Lemma 6.6, we successively apply Construction 6.4 to w . Let $w_1(H) \in \text{Ref}_{f_1}(w)$ be a reflection-extension of w through f_1 with splitting point u_1 and let g_1 be the automorphism applied to the second part of w in the reflection-extension process.

Consider the second element $f_2 \in H$, and note that $g_1(u_2)$ minimizes $d_G(v, \mathcal{V}(g_1(f_2)))$ among the vertices v of $w_1(H)$, since

$$d_G(\mathcal{V}(g_1(f_1)), \mathcal{V}(g_1(f_2))) > 2M + 2D.$$

Thus there is a reflection-extension $w_2(H) \in \text{Ref}_{g_1(f_2)}(w_1(H))$ of $w_1(H)$ through $g_1(f_2)$ with splitting point $g_1(u_2)$. Continuing in this way, we obtain a sequence of SAWs $w_1(H), \dots, w_k(H)$.

Letting H vary over all possible subsets of $\mathcal{B}(w)$ of cardinality k , we obtain a map $(w, H) \mapsto w_k(H)$. We claim that this map is injective. Indeed, note first that the set H can be reconstructed uniquely from $w_k(H)$. The arc e_1 is the unique arc in $E(t)$ such that the configuration defined by $w_k(H)$ on e_1 is an I-configuration. Then, up to orientation, f_1 is the unique arc of $K_{\bar{e}_1}$ such that $d_G(\mathcal{V}(e_1), \mathcal{V}(f_1)) = M$ and $w_k(H)$ visits $\mathcal{V}(f_1)$. This also defines uniquely the map g_1 used in the reflection-extension process. Similarly $g_1(e_2)$ is the unique arc in $E(g_1(t)) \setminus \{e_1\}$ such that the configuration defined by $w_k(H)$ on $g_1(e_2)$ is an I-configuration. Proceeding in this way, we see that H is uniquely given by $w_k(H)$. But then the claim is a consequence of Lemma 6.5.

Note that each time we apply the local modification, the length of the SAW we obtain increases by at most ℓ , where $\ell > 0$ is a uniform constant, hence $|w_k(H)| \leq |w| + k\ell$. Therefore,

$$\Lambda_n^2(R) = \sum_{A \in \mathcal{L}(c, \varepsilon, n)} \sum_{(A_e)_{e \in \mathcal{F}(A, \varepsilon, n, L)}} R^{|w|} \leq \frac{1}{\binom{m}{k} R^{k\ell}} \sum_{(w, H)} R^{|w_k(H)|}.$$

where the sum index (w, H) ranges over the set of all walks w represented by a c -completion A such that $A_S \in \mathcal{L}(c, \varepsilon, n)$ and $(A_e)_{e \in E(S) \setminus \{x\}} \in \mathcal{F}(A, \varepsilon, n, L_0)$, and all possible choices of subsets H of $\mathcal{B}(w)$ of size k .

Each walk $w_k(H)$ is a self-avoiding walk starting at $\mathcal{V}(x)$, and ending at $\mathcal{V}(x')$, where $x' = (g_k \circ \dots \circ g_1)(x)$. Note that g_1, \dots, g_k , and thus also x' , depend on H . By Theorem 4.13 there is a representation $A_H = (P_H, C_H)$ of $w_k(H)$ such that $X_H(x) = x$ and $Y_H(x') = x'$. Note also that by construction $C_H(x)$ and $C_H(x')$ are simple configurations because w meets $\mathcal{V}(x)$ only in its endpoints. Observe that the distance of x and x' is $2kM' + kN' + k + 1$ because for every $i = 1, \dots, k$ the distance of e_i and f_i is M' and the distance of f_i and $g_i(f_i)$ is N' . Thus the arrangement A_H is counted in some entry of $\mathfrak{J}_p(R)^{k(2M'+N'+1)+1}$. We obtain

$$\sum_{(w, H)} R^{|w_k(H)|} \leq \left\| \mathfrak{J}_p(R)^{k(2M'+N'+1)+1} \right\|_1 \leq e^{t_5 k}$$

for some constant t_5 .

Using the bound (6) for the binomial coefficient, we see that we can choose ε'' small enough so that

$$\Lambda_n^2(R) \leq \frac{e^{t_5 k}}{\binom{m}{k} R^{k\ell}} \leq e^{-\varepsilon'' m + o(m)},$$

Since $m = \lfloor \varepsilon' n \rfloor$, the desired assertion follows readily. \square

Proposition 6.2 can be now proved in a similar way.

Proof of Proposition 6.2. For $c, c' \in \mathcal{I}_p$, the analyticity of $\mathfrak{J}_{c,c'}(z, \mathbf{y})$ can be proved by arguing as in Proposition 6.1. Indeed, we can split $\Phi_n(R)$ into three sums $\Sigma_n(R)$, $\Lambda_n^1(R)$ and $\Lambda_n^2(R)$, where each of them can be handled in the same way. We emphasise that this is indeed the case for $\Lambda_n^2(R)$ because of our assumption that $c, c' \in \mathcal{I}_p$, which allows us to obtain an arrangement that contributes to $\mathfrak{J}_p(R)^N$, for some $N > 0$, after applying a successive application of our construction.

In general, for $c = (q, x, y), c' = (q', x', y') \in \mathcal{I}$, we can see that c and c' contain simple configurations $c_1 = (q_1, x, y)$ and $c'_1 = (q'_1, x', y')$ and U-configurations $c_2 = (q_2, x, x), c'_2 = (q'_2, \bar{x}', \bar{x}')$. Then for $(z, \mathbf{y}) = ((1 + \delta)R, (1 + \delta)\mathbf{F}(R))$

$$\mathfrak{J}_{c,c'}(z, \mathbf{y}) \leq \mathfrak{J}_{c_1,c'_1}(z, \mathbf{y}) P_{c_2}(z, \mathbf{y}) P_{c'_2}(z, \mathbf{y}),$$

and each term in the product is finite for some $\delta > 0$ sufficiently small. \square

7 Spectral radius and radius of convergence

Our goal in this section is to show that $R < r_{\mathcal{K}}$ for every U-component \mathcal{K} . Our approach loosely follows Alm and Janson [1] of decomposing U-configurations into smaller U-configurations and pairs of I-configurations. In order to formalise this approach we first need a few definitions.

A *pair configuration* is a pair (c_1, c_2) of persistent I-configurations $c_1 = (q_1, x_1, y_1), c_2 = (q_2, x_2, y_2) \in \mathcal{I}_p$ such that $x_1 = x_2$; we write $\mathcal{I}_{\text{pair}} = \{(c_1, c_2) \in \mathcal{I}_p \times \mathcal{I}_p : x_1 = x_2\}$ for the set of all pair configurations.

Let (c_1, c_2) and (c'_1, c'_2) be pair configurations. A (c_1, c_2) – (c'_1, c'_2) –completion is a pair (A_1, A_2) consisting of a c_1 – c'_1 –completion A_1 and a c_2 – c'_2 –completion A_2 having the same target arc f . The length of (A_1, A_2) is the length of A_1 (which equals the length of A_2) and we call it *disjoint* if the walks p_1 and p_2 represented by A_1 and A_2 are disjoint. Denote by $\mathcal{A}(c_1, c_2, c'_1, c'_2, n)$ the set of all (c_1, c_2) – (c'_1, c'_2) –completions of length n and by $\mathcal{A}_{\text{disj}}(c_1, c_2, c'_1, c'_2, n)$ the subset of those which are disjoint. We define a matrix $\mathfrak{J}_{\mathcal{I}_{\text{pair}}}(z)$ with index set $\mathcal{I}_{\text{pair}} \times \mathcal{I}_{\text{pair}}$ entry-wise by

$$(\mathfrak{J}_{\mathcal{I}_{\text{pair}}}(z))_{(c_1, c_2), (c'_1, c'_2)} = \sum_{(A_1, A_2) \in \mathcal{A}(c_1, c_2, c'_1, c'_2, 1)} z^{\|A_1\| + \|A_2\|}.$$

Like in the case of single configurations, we define a dependency digraph D_{pair} for pair configurations. Its vertex set is $\mathcal{I}_{\text{pair}}$ and we have an arc from (c_1, c_2) to (c'_1, c'_2) if the entry at position $(c_1, c_2), (c'_1, c'_2)$ in the matrix $\mathfrak{J}_{\mathcal{I}_{\text{pair}}}(z)$ is non-zero.

A pair configuration (c_1, c_2) is called simple if both c_1 and c_2 are simple. It is called persistent, if it lies on a walk connecting two simple pair configurations in D_{pair} . The set of persistent pair configurations is denoted $\mathcal{I}_{\text{pair},p}$. Note that while every persistent pair configuration consists of two persistent I-configurations, the converse is not necessarily true. This is due to the fact that a c_1 - c'_1 -completion A_1 and a c_2 - c'_2 -completion A_2 might have the same source arc but different target arcs.

Finally we define the matrix $\mathfrak{J}_{\text{pair},\text{disj}}(z)$ with index set $\mathcal{I}_{\text{pair},p} \times \mathcal{I}_{\text{pair},p}$ entry-wise by

$$(\mathfrak{J}_{\text{pair},\text{disj}}(z))_{(c_1, c_2), (c'_1, c'_2)} = \sum_{(A_1, A_2) \in \mathcal{A}_{\text{disj}}(c_1, c_2, c'_1, c'_2, 1)} z^{\|A_1\| + \|A_2\|}.$$

Denote by $\lambda_{\text{pair},\text{disj}}(z)$ the spectral radius of $\mathfrak{J}_{\text{pair},\text{disj}}(z)$.

Lemma 7.1. *For every $z \in (0, R]$, $\mathfrak{J}_{\text{pair},\text{disj}}(z)$ has finite entries and furthermore, we have $\lambda_{\text{pair},\text{disj}}(z) < \lambda_{\mathcal{I}_p}(z)^2$.*

Proof. We start by noting that for $z \geq 0$ all matrices considered in this proof are non-negative, so whenever we restrict the index set or reduce some entries of a matrix, the spectral radius of the resulting matrix must be smaller than or equal to the spectral radius of the original matrix.

Consider the Kronecker product $\mathfrak{J}_{\mathcal{I}_p \times \mathcal{I}_p}(z) = \mathfrak{J}_{\mathcal{I}_p}(z) \otimes \mathfrak{J}_{\mathcal{I}_p}(z)$. The eigenvalues of the Kronecker product of two matrices M_1 and M_2 are precisely the products $\lambda_1 \lambda_2$ of eigenvalues λ_1 of M_1 and λ_2 of M_2 . Therefore the spectral radius of $\mathfrak{J}_{\mathcal{I}_p \times \mathcal{I}_p}(z)$ is $\lambda_{\mathcal{I}_p}(z)^2$. By Lemma 5.8, the matrix $\mathfrak{J}_{\mathcal{I}_p}(z)$ has finite entries, thus the same holds for $\mathfrak{J}_{\mathcal{I}_p \times \mathcal{I}_p}(z)$.

By definition every (c_1, c_2) - (c'_1, c'_2) -completion (A_1, A_2) consists of a c_1 - c'_1 -completion A_1 and a c_2 - c'_2 -completion A_2 , so for $z \in (0, R]$ all entries of the matrix $\mathfrak{J}_{\mathcal{I}_{\text{pair}}}(z)$ are smaller than or equal to the respective entries of $\mathfrak{J}_{\mathcal{I}_p \times \mathcal{I}_p}(z)$. We conclude that its spectral radius $\lambda_{\mathcal{I}_{\text{pair}}}(z)$ satisfies $\lambda_{\mathcal{I}_{\text{pair}}}(z) \leq \lambda_{\mathcal{I}_p}(z)^2$. By Lemma 5.4 all simple pair configurations are contained in the same strong component of D_{pair} , so the set $\mathcal{I}_{\text{pair},p}$ of persistent pair configurations is a strong component of D_{pair} . In particular for $z > 0$ the submatrix $\mathfrak{J}_{\mathcal{I}_{\text{pair},p}}(z)$ of $\mathfrak{J}_{\mathcal{I}_{\text{pair}}}(z)$ obtained by restricting the index set to $\mathcal{I}_{\text{pair},p}$ is irreducible. Finally, the entries of the matrix $\mathfrak{J}_{\text{pair},\text{disj}}(z)$ are smaller than or equal to the respective entries of $\mathfrak{J}_{\mathcal{I}_{\text{pair},p}}(z)$ with strict inequality for some entries, for instance the diagonal entries. Thus $\lambda_{\text{pair},\text{disj}}(z) < \lambda_{\mathcal{I}_{\text{pair}}}(z) \leq \lambda_{\mathcal{I}_p}(z)^2$ holds. \square

To simplify notation let $\lambda_{\mathcal{U}}(z) := \max\{\lambda_{\mathcal{K}}(z) \mid \mathcal{K} \subseteq \mathcal{U} \text{ strong component of } D\}$ and let $\lambda_{\mathcal{I}_t}(z) := \max\{\lambda_{\mathcal{K}}(z) \mid \mathcal{K} \subseteq \mathcal{I}_t \text{ strong component of } D\}$.

Lemma 7.2. *For every $z \in (0, R]$ we have $\lambda_{\mathcal{U}}(z) \leq \lambda_{\text{pair},\text{disj}}(z)$.*

Proof. Let $\mathcal{K} \subseteq \mathcal{U}$ be a strong component of D . Let c, c' be two configurations in \mathcal{K} and recall that $\mathcal{A}(c, c', n)$ denotes the set of c - c' -completions of length n . Then

$$(\mathfrak{J}_{\mathcal{K}}(z)^n)_{c, c'} = \sum_{A \in \mathcal{A}(c, c', n)} z^{\|A\|}. \quad (9)$$

By Corollary 3.4 there is a constant N' such that $\mathcal{V}(e) \cap \mathcal{V}(f) = \emptyset$ whenever $d_T(e, f) \geq N'$ holds for two edges $e, f \in E(T)$. Let $A \in \mathcal{A}(c, c', n)$ for some $n \geq N'$, let S be the support of A and let ι and τ be source and target arc of A , respectively. The arrangement A represents a unique SAW $\varphi(A)$ on $\mathcal{G}(S)$ starting and ending in $\mathcal{V}(\iota)$. Let v_0, v_1, \dots, v_k be the sequence of vertices of $\varphi(A)$ contained in at least one of $\mathcal{V}(\iota)$ and $\mathcal{V}(\tau)$, ordered by their appearance in $\varphi(A)$. We decompose $\varphi(A)$ as

$$\varphi(A) = v_0 w_1 v_1 w_2 v_3 \dots w_k v_k,$$

in other words, w_i is the sub-walk of $\varphi(A)$ starting at v_{i-1} and ending at v_i . Denote by K the size of an adhesion set of the tree decomposition. Then the number of sub-walks k is bounded from above by $2K$. We call a walk w_i virtual if it contains only a single arc and this arc is contained in $\mathcal{E}(\iota)$ or $\mathcal{E}(\tau)$, and non-virtual otherwise. Then each non-virtual w_i belongs to one of four possible classes.

- (i) w_i starts and ends in $\mathcal{V}(\iota)$
- (ii) w_i starts and ends in $\mathcal{V}(\tau)$
- (iii) w_i starts in $\mathcal{V}(\iota)$ and ends in $\mathcal{V}(\tau)$
- (iv) w_i starts in $\mathcal{V}(\tau)$ and ends in $\mathcal{V}(\iota)$

Observe that the number of walks of class (iii) coincides with the number of walks of class (iv) because $\varphi(A)$ starts and ends in $\mathcal{V}(\iota)$. Hence we can group them into pairs, each consisting of a walk of class (iii) and a walk of class (iv). Denote by \mathcal{W}_1 , \mathcal{W}_2 and \mathcal{W}_3 the sets of walks w_i of class (i), class (ii) and the set of pairs (w_i, w_j) of classes (iii) and (iv), respectively. For technical reasons, pairs (w_i, w_j) such that the final vertex of w_i coincides with the initial vertex of w_j (that is, $j = i + 1$) are not added to \mathcal{W}_3 ; instead their concatenation is included in \mathcal{W}_1 . Note that \mathcal{W}_3 contains at least one pair of walks, because c' is non-boring.

By Theorem 4.13 each $w \in \mathcal{W}_1$ is represented by a c_w -completion A_w , where c_w is a U-configuration on ι . Similarly, each $w \in \mathcal{W}_2$ is represented by a c_w -completion A_w , where c_w is a U-configuration on τ . Finally, each pair $(w, w') \in \mathcal{W}_3$ is represented by a disjoint $(c_w, c_{w'})$ - $(c'_w, c'_{w'})$ -completion $(A_w, A_{w'})$ of length n , where c_w and $c_{w'}$ are simple configurations on ι and c'_w and $c'_{w'}$ are simple configurations on τ .

Note that by construction

$$z^{\|A\|} = \prod_{w \in \mathcal{W}_1} z^{\|A_w\|} \prod_{w \in \mathcal{W}_2} z^{\|A_w\|} \prod_{(w, w') \in \mathcal{W}_3} z^{\|A_w\| + \|A_{w'}\|}.$$

Moreover, the arrangement A can be uniquely reconstructed from the pair c, c' , the c_w -completions A_w , and the $(c_w, c_{w'})$ - $(c'_w, c'_{w'})$ -completions $(A_w, A_{w'})$. Let L be the number of possible configurations on any fixed arc e . Then obviously for each $c \in \mathcal{C}$ there are at most L configurations on e which are Γ -equivalent to c , so at most L of the configurations c_w and c'_w are equivalent to c . We conclude that for fixed $k_1, k_2 \geq 0$ and $k_3 \geq 1$, the

subset of all $A \in \mathcal{A}(c, c', n)$ such that $|\mathcal{W}_i| = k_i$ for $i = 1, 2, 3$ gives a contribution to the sum (9) of at most

$$\left(L \sum_{c_1 \in \mathcal{U}} F_{c_1}(z) \right)^{k_1+k_2} \left(L^2 \sum_{(c_3, c_4), (c'_3, c'_4) \in \mathcal{I}_{\text{pair}, p}} (\mathfrak{J}_{\text{pair}, \text{disj}}(z)^n)_{(c_3, c_4), (c'_3, c'_4)} \right)^{k_3} \quad (10)$$

Lemma 6.3 tells us that $F_{c_1}(R)$ is finite for every $c_1 \in \mathcal{U}$. In particular, by monotonicity of the entries of $F_{c_1}(z)$, there is some constant L' which does not depend on n such that

$$L \sum_{c_1 \in \mathcal{U}} F_{c_1}(z) \leq L'$$

for all $z \in [0, R]$.

Recall that for any matrix norm $\|\cdot\|$ and any matrix A with spectral radius λ there are finite constants M, m such that $\|A^n\| \leq Mn^m\lambda^n$ holds for every $n \in \mathbb{N}$. An easy way to verify this is by computing powers of Jordan blocks in the Jordan normal form of A . Applying this to $\mathfrak{J}_{\text{pair}, \text{disj}}(z)$, we obtain

$$L^2 \sum_{(c_3, c_4), (c'_3, c'_4) \in \mathcal{I}_{\text{pair}, p}} (\mathfrak{J}_{\text{pair}, \text{disj}}(z)^n)_{(c_3, c_4), (c'_3, c'_4)} \leq Mn^m \lambda_{\text{pair}, \text{disj}}^n(z) \quad (11)$$

Utilizing these estimates and that $0 \leq k_1 + k_2 \leq 2K$ and $1 \leq k_3 \leq K$, from (9) we obtain

$$(\mathfrak{J}_{\mathcal{K}}(z)^n)_{c, c'} \leq \sum_{k_1=0}^{2K} \sum_{k_3=1}^K (L')^{k_1} (Mn^m \lambda_{\text{pair}, \text{disj}}^n(z))^{k_3}$$

Note that $\lambda_{\text{pair}, \text{disj}}(z) \leq 1$ by Lemma 7.1 combined with Lemma 5.8. Thus

$$(\mathfrak{J}_{\mathcal{K}}(z)^n)_{c, c'} \leq M' n^{m'} \lambda_{\text{pair}, \text{disj}}^n(z)$$

holds for some new constants M' and m' independent of n .

Taking n -th roots and sending n to infinity, an application of Gelfand's formula completes the proof:

$$\lambda_{\mathcal{K}}(z) = \lim_{n \rightarrow \infty} \| \mathfrak{J}_{\mathcal{K}}(z)^n \|_{\infty}^{1/n} \leq \lambda_{\text{pair}, \text{disj}}(z)$$

□

Lemma 7.3. *For every $z \in (0, R]$ we have $\lambda_{\mathcal{I}_t}(z) \leq \lambda_{\text{pair}, \text{disj}}(z)$.*

Proof. We follow a similar strategy as in the proof of Lemma 7.2. Let $\mathcal{K} \subseteq \mathcal{U}$ be a strong component of D and let c, c' be two configurations in \mathcal{K} . Then

$$(\mathfrak{J}_{\mathcal{K}}(z)^n)_{c, c'} = \sum_{A \in \mathcal{A}(c, c', n)} z^{\|A\|}. \quad (12)$$

As in the previous proof we translate every $A \in \mathcal{A}(c, c', n)$ to the SAW $\varphi(A)$ represented by A . Then we decompose $\varphi(A)$ at every vertex contained in the adhesion sets $\mathcal{V}(\iota) \cup \mathcal{V}(\tau)$,

where ι and τ are source and target arc of A . We end up with sub-walks w_1, \dots, w_k of $\varphi(A)$ and categorize them into classes (i) to (iv). However, because A is a c - c' -completion for two I-configurations c and c' , the walk $\varphi(A)$ starts at $\mathcal{V}(\iota)$ and ends at $\mathcal{V}(\tau)$. This means that class (iii) contains one more walk than class (iv). We denote by w_0 the walk w_i with maximal index i in class (iii). Excluding w_0 , as before we can build pairs, each consisting of a walk of class (iii) and a walk of class (iv). Denote by \mathcal{W}_1 , \mathcal{W}_2 and \mathcal{W}_3 the sets of walks w_i of class (i), class (ii) and the set of pairs (w_i, w_j) of classes (iii) and (iv), respectively.

As before, for every $w \in \mathcal{W}_1 \cup \mathcal{W}_2$ there is a c_w -completion A_w representing w provided by Theorem 4.13 such that c_w is a U-configuration. Furthermore, for every $(w, w') \in \mathcal{W}_3$ there is a $(c_w, c_{w'})$ - $(c'_w, c'_{w'})$ -completion $(A_w, A_{w'})$ representing (w, w') such that $c_w, c_{w'}, c'_w$ and $c'_{w'}$ are simple configurations. Finally, there is a c_0 - c'_0 -completion A_{w_0} representing w_0 such that c_0 and c'_0 are simple configurations.

The main difference to the previous proof is that the size k_3 of the set \mathcal{W}_3 need not necessarily be larger than 0. We denote by $\mathcal{A}_{=0}(c, c', n)$ the subset of all $A \in \mathcal{A}(c, c', n)$ such that $k_3 = 0$ and by $\mathcal{A}_{>0}(c, c', n)$ its complement $\mathcal{A}(c, c', n) \setminus \mathcal{A}_{=0}(c, c', n)$. Let us first deal with arrangements in $\mathcal{A}_{>0}(c, c', n)$, which works similarly to the previous proof. The only difference is an additional factor

$$L \sum_{c_0, c'_0 \in \mathcal{I}_p} (\mathfrak{J}_{\mathcal{I}_p}(z)^n)_{c_0, c'_0}$$

in the analogue of equation (10) coming from the walk w_0 . As in (11), we can bound this factor from above by $M_0 n^{m_0} \lambda_{\mathcal{I}_p}^n(z)$, for some constants M_0 and m_0 independent from n . By Lemma 5.8 we have that $\lambda_{\mathcal{I}_p}(z) \leq 1$, thus we again end up with

$$\sum_{A \in \mathcal{A}_{>0}(c, c', n)} z^{|A|} \leq M' n^{m'} \lambda_{\text{pair, disj}}^n(z) \quad (13)$$

for some constants M', m' independent of n .

We are left to deal with arrangements in $\mathcal{A}_{=0}(c, c', n)$. Here we need to treat the c_w -completions A_w for $w \in \mathcal{W}_1 \cup \mathcal{W}_2$ more carefully. Let S_w be the support of A_w and denote by W the open path in T connecting ι and τ . Then $W \cap S_w$ is an open path of W containing ι if $w \in \mathcal{W}_1$ and τ if $w \in \mathcal{W}_2$. Denote by $j_w \in [n]$ the length of this open path. Then A_w consists of a c_w - c'_w -completion of length $j_w - 1$ for some non-boring U-configuration c'_w and a c'_w -completion. We claim that there is a constant N independent of n such that

$$\sum_{w \in \mathcal{W}_1 \cup \mathcal{W}_2} j_w \geq n - N \quad (14)$$

holds for every $A \in \mathcal{A}_{=0}(c, c', n)$. Indeed, by the choice of A for every interior arc e of W the configuration $\rho(C(e))$ lies in the component \mathcal{K} and thus must be transient by Lemma 5.6. In particular $\mathcal{V}(e)$ must be visited by some $w \in \mathcal{W}_1 \cap \mathcal{W}_2$, otherwise $C(e) = C_{w_0}(e)$ holds and thus $\rho(C(e))$ lies in the same component as the simple configurations c_0 and c'_0 and thus is persistent. By Corollary 3.4 there is a constant N' such that

$\mathcal{V}(e) \cap \mathcal{V}(f) = \emptyset$ whenever $d_T(e, f) \geq N'$. In particular, $w \in \mathcal{W}_1$ cannot visit $\mathcal{V}(e)$ for any arc e such that $d_T(t, e) > j_w + N'$. This implies in particular that (14) holds for $N = 2N'$.

For fixed $k \geq 0$ and integers $j_1, \dots, j_k \in [n]$, the subset of all $A \in \mathcal{A}_{=0}(c, c', n)$ such that \mathcal{W}_i contains k elements w_1, \dots, w_k with $j_{w_i} = j_i$ gives a contribution to the sum (12) of at most

$$\left(L \sum_{c_0, c'_0 \in \mathcal{I}_p} (\mathfrak{J}_{\mathcal{I}_p}(z)^n)_{c_0, c'_0} \right) \prod_{i=1}^k \left(L \sum_{c_1, c'_1 \in \mathcal{U}} (\mathfrak{J}_{\mathcal{U}}(z)^{j_i-1})_{c_1, c'_1} F_{c'_1}(z) \right)$$

The first factor is bounded by $M_0 n^{m_0}$ as in the previous case. For the other factors, we use again that by Lemma 6.3 there is a constant L' such that $F_c(z) \leq L'$ holds for every $c \in \mathcal{U}$ and every $z \in [0, R]$. Therefore, we bound each factor as in (11) to obtain the upper bound

$$\prod_{i=1}^k \left(L \sum_{c_1, c'_1 \in \mathcal{U}} (\mathfrak{J}_{\mathcal{U}}(z)^{j_i-1})_{c_1, c'_1} F_{c'_1}(z) \right) \leq (M_1)^k \prod_{i=1}^k (j_i - 1)^{m_1} \lambda_{\mathcal{U}}(z)^{j_i-1} \leq M_2^k n^{m_2} \lambda_{\mathcal{U}}(z)^n,$$

where M_1, M_2, m_1 and m_2 are constants independent of n and for the last inequality we use (14) and that $\lambda_{\mathcal{U}}(z) \leq 1$ by Lemma 7.2.

With these estimates we obtain

$$\sum_{A \in \mathcal{A}_{=0}(c, c', n)} z^{\|A\|} \leq M_0 n^{m_0} \sum_{k=0}^{2K} \sum_{j_1, \dots, j_k=1}^n M_2^k n^{m_2} \lambda_{\mathcal{U}}^n(z) \leq M_3 n^{m_3} \lambda_{\mathcal{U}}^n(z) \quad (15)$$

for some new constants M_3 and m_3 .

Combining equations (13) and (15) and using Lemma 7.2, we end up with

$$(\mathfrak{J}_{\mathcal{K}}(z)^n)_{c, c'} \leq M'' n^{m''} \lambda_{\text{pair, disj}}^n(z)$$

for some new constants M'' and m'' . Taking the n -th root and sending n to infinity completes the proof. \square

Corollary 7.4. *If Γ does not fix an end of T , then $R_{\mathcal{U}} > R$.*

Proof. Since $P_c(z, \mathbf{y})$ is analytic at $(z, \mathbf{y}) = (R, (F_{c'}(R))_{c' \in \mathcal{C}})$ for all $c \in \mathcal{U}$ by Proposition 6.1, the matrix $\mathfrak{J}_{\mathcal{U}}(R)$ is well-defined and has finite entries (for the strong components of \mathcal{U} this is already implied by Lemma 5.8). It follows from Lemmas 5.8, 7.1, and 7.2 that $\lambda_{\mathcal{U}}(R) < 1$, hence $I - \mathfrak{J}_{\mathcal{U}}(R)$ is invertible. Since $P_c(z, \mathbf{y})$ is analytic at $(z, \mathbf{y}) = (R, (F_{c'}(R))_{c' \in \mathcal{C}})$ for all $c \in \mathcal{U}$, it follows from the analytic Implicit Function Theorem that the functions $F_c(z)$ for $c \in \mathcal{U}$ are analytic in a neighbourhood of R . \square

8 Proof of main results

In this section we will prove the main results of this paper building on the results developed in the previous sections.

Our overall strategy is to exploit the close connection between c -completions and SAWs to obtain an expression for the SAW-generating function $F_{\text{SAW}}(z)$ in terms of the functions $F_c(z)$. To do this, recall that Theorem 4.15 allows us to bijectively translate every SAW w on G into a complete arrangement $A = (P, C)$ on the support of w representing w such that $\|A\|$ coincides with the length of w . By Lemma 4.17 the shape $P(s_0)$ of the source s_0 of A starts with the first arc of p . In particular, even when fixing the starting vertex o of our SAWs, s_0 may still vary depending on the specific choice of w , the only restriction is that $\mathcal{V}(s_0)$ has to contain o . To get rid of this small inconvenience, let us first manipulate our tree decomposition.

Fix $o \in V(G)$ and let S be the subtree of T induced by all vertices $t \in V(T)$ for which $\mathcal{V}(t)$ contains a vertex at distance at most 2 from o . Note that S is finite by Proposition 3.3 and that $\mathcal{G}(S)$ contains all edges of G incident to o or one of its neighbours. Let $\mathcal{T}' = (T', \mathcal{V}') = \mathcal{T}/S$ be the tree decomposition obtained from \mathcal{T} by contracting all interior edges of S and let s_0 be the vertex of T' representing $V(S)$. Clearly the tree decomposition \mathcal{T}' is not Γ -invariant, as the new part $\mathcal{V}'(s_0)$ plays a special role and cannot be mapped to any other parts. However, every arc $e \in E(T')$ pointing towards s_0 was also present in T and the open cone K_e in T' coincides with the respective cone in T , so c -completions and their generating functions remain the same. By Theorem 4.15 any SAW w starting at o is represented by a unique complete arrangement A on the support of w in T' . Clearly $P(s_0)$ starts at o . Moreover, all edges incident to o are contained in the part $\mathcal{V}'(s_0)$ by construction of \mathcal{T}' . Thus by Lemma 4.17 the source of A has to be s_0 . We obtain

$$F_{\text{SAW}}(z) = \sum_{w \text{ SAW starting at } o} z^{|w|} = \sum_{\substack{A \text{ compl. arr.} \\ \text{with source } s_0 \text{ s.t.} \\ P(s_0) \text{ starts at } o}} z^{\|A\|}.$$

By decomposing the tree T' into an open star centered at s_0 and possibly infinitely many cones $K_{\bar{e}}$ for $e \in E(s_0)$ and the respective arrangement on the open subtree S into an arrangement $A = (P, C)$ on $\text{star}(s_0)$ and $C(e)$ -completions for all non-boring configurations on arcs e of $\text{star}(s_0)$, we conclude

$$F_{\text{SAW}}(z) = \sum_{\substack{A \text{ arr. on } \text{star}(s_0): \\ P(s_0) \text{ starts at } o}} z^{\|A\|} \prod_{e \in \text{nb}(A, s_0)} F_{C(e)}(z).$$

We point out that all arrangements A included in the sum satisfy $X(e) = \bar{e}$ for every $e \in E(s_0)$ because by construction o is not included in any adhesion set.

To analyse $F_{\text{SAW}}(z)$ we need to first isolate the terms that determine its radius of convergence. To this end we write

$$F_{\text{SAW}}(z) = F_{\text{SAW}}^1(z) + F_{\text{SAW}}^2(z), \tag{16}$$

where

$$F_{\text{SAW}}^1(z) = \sum_{\substack{A \text{ arr. on } \text{star}(s_0): \\ P(s_0) \text{ starts at } o \text{ and} \\ \forall e \in E(s_0): Y(e) = \bar{e}}} z^{\|A\|} \prod_{e \in \text{nb}(A, s_0)} F_{C(e)}(z)$$

and

$$F_{\text{SAW}}^2(z) = \sum_{\substack{A \text{ arr. on } \text{star}(s_0): \\ P(s_0) \text{ starts at } o \\ \exists e_0 \in E(s_0): Y(e_0) = e_0}} z^{\|A\|} \prod_{e \in \text{nb}(A, s_0)} F_{C(e)}(z).$$

Proposition 8.1 below implies that $F_{\text{SAW}}^1(z)$ is analytic at $z = R$. Hence only $F_{\text{SAW}}^2(z)$ is relevant. In order to study the latter we further write

$$F_{\text{SAW}}^2(z) = \mathbf{F}^-(z) \sum_{k=0}^{\infty} \mathfrak{J}_{\mathcal{I}}(z)^k \mathbf{F}^+(z) \quad (17)$$

where $\mathbf{F}^-(z) = (F_c^-(z))_{c \in \mathcal{I}}$ is the vector with entries

$$F_c^-(z) = \sum_{\substack{A \text{ arr. on } \text{star}(s_0): \\ P(s_0) \text{ starts at } o \\ \exists e_0 \in E(s_0): C(e_0) = c}} z^{\|A\|} \prod_{e \in \text{nb}(A, s_0) \setminus \{e_0\}} F_{C(e)}(z)$$

and $\mathbf{F}^+(z) = (F_c^+(z))_{c \in \mathcal{I}}$ is the vector with entries

$$F_c^+(z) = \sum_{\substack{A \text{ arr. on } \text{star}(x^-): \\ C(x) = c \text{ and} \\ \forall e \in E(x^-): Y(e) = \bar{e}}} z^{\|A\|} \prod_{e \in \text{nb}(A, x^-) \setminus \{x\}} F_{C(e)}(z).$$

By Lemma 5.8, we know that

$$F_{\text{SAW}}^2(z) = \mathbf{F}^-(z) (I - \mathfrak{J}_{\mathcal{I}}(z))^{-1} \mathbf{F}^+(z) \quad (18)$$

for $z < R$.

In order to prove Theorem 1.1, by [41, Theorem 4] it suffices to show that $\mu_p < \mu_w$, where μ_p and μ_w are the connective constants for SAPs and SAWs, respectively. To stay within the framework of our definitions, it will be convenient to work with self-avoiding returns (SARs) instead of SAPs. By construction of \mathcal{T}' all edges incident to neighbours of o are contained in the part $\mathcal{V}'(s_0)$. Thus by Lemma 4.17 every SAR w starting at o is represented by a unique complete arrangement A on the support of w in T' with source and target s_0 such that $P(s_0)$ starts at o and ends at a neighbour of o . Let

$$F_{\text{SAR}}(z) = \sum_{w \text{ SAR starting at } o} z^{|w|}.$$

As above, we obtain

$$F_{\text{SAR}}(z) = \sum_{\substack{A \text{ arr. on } \text{star}(s_0): \\ P(s_0) \text{ starts at } o \text{ and ends in } N(o)}} z^{\|A\|} \prod_{e \in \text{nb}(A, s_0)} F_{C(e)}(z). \quad (19)$$

Observe that the condition $X(e) = Y(e) = \bar{e}$ for every $e \in E(s_0)$ implies that each configuration $C(e)$ is a U-configuration.

We will now prove the following result.

Proposition 8.1. *Each of the functions $F_{\text{SAW}}^1(z)$, $F_{\text{SAR}}(z)$, $F^-(z)$ and $F^+(z)$ is analytic at $z = R$. Moreover, the function $F_{\text{SAW}}(z)$ is analytic in the interval $(0, R)$.*

Proof. Let $F(z)$ be any of the functions $F_{\text{SAW}}^1(z)$, $F_{\text{SAR}}(z)$, $F_c^-(z)$ and $F_c^+(z)$. Then there is some $s \in V(T)$ such that each of the walks counted by $F(z)$ starts at some given vertex v_0 in $V(s)$ and ends in $V(s)$. Let w be one of these walks. We decompose w as follows.

Let e_1 and e_2 be the first and last arc of w not in $\mathcal{E}(s)$, respectively, and let $v_1 = e_1^-$ and $v_2 = e_2^+$. If no such arcs exist, we set $v_1 = v_2 = w^+$. Then wv_1 and v_2w are SAWs on $\mathcal{G}(s)$. If $v_1 \neq v_2$, let $f_1, f_2 \in E(s)$ be such that e_i is an edge of $\mathcal{G}(K_{\bar{f}_i})$. Then v_1wv_2 is a SAW on G starting in $\mathcal{V}(f_1)$ and ending in $\mathcal{V}(f_2)$. If $f_1 = f_2$, then v_1wv_2 can be represented by a pair consisting of a c -completion on the cone K_{f_1} and a c' -completion on the cone $K_{\bar{f}_1}$. Note that c and c' are U-configurations which only differ in their entry and exit direction. Assume now that $f_1 \neq f_2$. Then v_1wv_2 can be represented by a triple consisting of a c_1 -completion on the cone $K_{\bar{f}_1}$, a c_2 -completion on $K_{\bar{f}_2}$, and a c'_1 - c'_2 -completion of length 1 on $K_{f_1} \cap K_{f_2}$ which contributes to $\mathfrak{J}_{\mathcal{I}}$. Here c_1 and c_2 are U-configurations, c'_1 and c'_2 are I-configurations, c_1 and c'_1 only differ in their entry direction while c_2 and c'_2 only differ in their exit direction.

With this decomposition in hand, the analyticity of $F(z)$ follows from Proposition 6.1, Proposition 6.2 and Lemma 6.6.

To see that $F_{\text{SAW}}(z)$ is analytic in the interval $(0, R)$, recall that by (16) and (18) we have that

$$F_{\text{SAW}}(z) = F_{\text{SAW}}^1(z) + \mathbf{F}^-(z)(I - \mathfrak{J}_{\mathcal{I}}(z))^{-1}\mathbf{F}^+(z).$$

is finite for $z \in (0, R)$. □

With the above result in hands we can show that R is determined by \mathcal{I}_p .

Proposition 8.2. *If Γ does not fix an end of T , then we have $R_{\mathcal{I}_p} = R$.*

Proof. First, we claim that $\lambda_{\mathcal{I}_p}(R) = 1$. Assume for a contradiction that $\lambda_{\mathcal{I}_p}(R) < 1$. Then $\lambda_{\mathcal{I}}(R) < 1$ by Lemmas 5.8, 7.1, and 7.3. Recall that $\mathfrak{J}_{\mathcal{I}}(z)$ is analytic at $z = R$ by Proposition 6.2. Then $\lambda_{\mathcal{I}}(z) < 1$ for every z in a neighbourhood of R . Now, observe that

$$(F_c(z))_{c \in \mathcal{I}} = \sum_{n=0}^{\infty} \mathfrak{J}_{\mathcal{I}}(z)^n \mathbf{F}^+(z) \tag{20}$$

provided that $\mathfrak{J}_{\mathcal{I}}(z)$ and $\mathbf{F}^+(z)$ are analytic. If $\lambda_{\mathcal{I}}(z) < 1$, this implies that $(F_c(z))_{c \in \mathcal{I}} = (I - \mathfrak{J}_{\mathcal{I}}(z))^{-1}\mathbf{F}^+(z)$.

Since this holds at $z = R$, it follows that $(F_c(z))_{c \in \mathcal{I}}$ is analytic at $z = R$, hence $R_{\mathcal{I}} > R$. The latter together with Corollary 7.4 implies that $R > R$, which is absurd. This proves the claim.

Since $\lambda_{\mathcal{I}_p}(R) = 1$, equation (20) together with the fact that \mathcal{I}_p is a strong component implies that $F_c(R)$ is infinite for every $c \in \mathcal{I}_p$. The desired assertion follows. □

We are now ready to prove the main results of this paper. We first show that SAW is ballistic. In order to prove Theorem 1.2, we will show the strict inequality $\mu_p < \mu_w$. The fact that this implies the ballisticity of SAW was proved in [41, Theorem 4]. This result was proved for transitive graphs but as explained in [41, Remark 4.1], it can be extended to graphs which are not necessarily transitive but satisfy

$$\limsup_{n \rightarrow \infty} \left(\sup_{x \in V} p_n(x) \right)^{1/n} < \liminf_{n \rightarrow \infty} \left(\inf_{x \in V} c_n(x) \right)^{1/n}.$$

In particular, the result of [41, Theorem 4] holds for quasi-transitive graphs.

Proof of Theorem 1.2. By Proposition 8.1, the generating function $F_{\text{SAR}}(z)$ is analytic at $z = R$. Since there are fewer self-avoiding polygons of length n containing o than self-avoiding returns starting at o , it follows that $\mu_p < 1/R$.

Let p be a shortest walk from o to some adhesion set $\mathcal{V}(e)$, where $e^+ = s_0$. Let c be the simple configuration (q, x, y) where q is the walk consisting only of the terminal vertex of p , $x = e$, and $y = \bar{e}$. Then $z^{|p|} F_c(z) \leq F_{\text{SAW}}(z)$ because p together with the walk corresponding to any c -completion is a self-avoiding walk. This together with Proposition 8.2 implies that $\mu_w \geq 1/R_{\mathcal{I}_p} = 1/R$. The desired result follows. \square

Remark 8.3. As mentioned in the proof of Theorem 1.2, we have $\mu_w \geq 1/R$. The reverse inequality $\mu_w \leq 1/R$ follows from the fact that $F_{\text{SAW}}(z)$ is finite for every $z < R$ by Proposition 8.1. This proves that $\mu_w = 1/R$.

We will now prove that c_n grows asymptotically like μ_w^n , that is, the subexponential factor is $O(1)$.

Proof of Theorem 1.1. We will show that all singularities of $F_{\text{SAW}}(z)$ on the circle $|z| = R$ are simple poles and they are located at the complex k -th roots of R^k for some k , from which the desired result will follow. To this end, let $r > R$ be such that each of $F_{\text{SAW}}^1(z), F^-(z), F^+(z)$ and $\mathfrak{J}_{\mathcal{I}}(z)$ is analytic for every $z \in \mathbb{C}$ such that $|z| < r$. Such an r exists by Propositions 6.2 and 8.1. Then the points of singularity of $F_{\text{SAW}}(z)$ for $|z| < r$ are the points at which $\det(I - \mathfrak{J}_{\mathcal{I}}(z)) = 0$ by (16) and (18).

Since $\det(I - \mathfrak{J}_{\mathcal{I}}(z))$ is a non-constant analytic function, its zeros are isolated, hence we can assume without loss of generality that $\det(I - \mathfrak{J}_{\mathcal{I}}(z)) \neq 0$ for every $R < |z| < r$. For every $|z| < R$ we have $\lambda_{\mathcal{I}}(z) < 1$ by Lemma 5.8, thus the only singularities of $F_{\text{SAW}}(z)$ for $|z| < r$ are at the points of the circle $|z| = R$ at which 1 is an eigenvalue of $\mathfrak{J}_{\mathcal{I}}(z)$. Let us denote this set by S and consider some $w \in S$.

We claim that 1 is a simple eigenvalue of $\mathfrak{J}_{\mathcal{I}}(w)$. Indeed, note that 1 is a simple eigenvalue of the irreducible matrix $\mathfrak{J}_{\mathcal{I}_p}(R)$ by the Perron-Frobenius theorem, and that $\lambda_{\mathcal{I}_t}(R) < 1$ by Lemmas 7.3 and 7.1. Since $|\mathfrak{J}_{c,c'}(w)| \leq \mathfrak{J}_{c,c'}(R)$ for every $c, c' \in \mathcal{I}$ by the triangle inequality we have $\lambda_{\mathcal{I}_p}(w) = 1$ and $\lambda_{\mathcal{I}_t}(w) < 1$. It thus suffices to show that 1 is a simple eigenvalue of $\mathfrak{J}_{\mathcal{I}_p}(w)$.

Note that $\mathfrak{J}_{\mathcal{I}_p}(w)$ is not a real matrix in general, and we cannot apply Perron-Frobenius. Instead, we argue as follows. The matrix $\mathfrak{J}_{\mathcal{I}_p}(R)$ is irreducible, $|\mathfrak{J}_{c,c'}(w)| \leq \mathfrak{J}_{c,c'}(R)$ for $c, c' \in \mathcal{I}_p$, and $\mathfrak{J}_{\mathcal{I}_p}(w)$ and $\mathfrak{J}_{\mathcal{I}_p}(R)$ have the same spectral radius. Because 1

is the eigenvalue of $\mathfrak{J}_{\mathcal{I}_p}(w)$ of maximal modulus, it follows from [43, Chapter I, Proposition 6.4] that $\mathfrak{J}_{\mathcal{I}_p}(w)$ and $\mathfrak{J}_{\mathcal{I}_p}(R)$ are similar matrices and thus have the same spectrum. Hence 1 is a simple eigenvalue of $\mathfrak{J}_{\mathcal{I}_p}(w)$. Since $\lambda_{\mathcal{I}_t}(w) < 1$ we deduce that 1 is a simple eigenvalue of $\mathfrak{J}_{\mathcal{I}}(w)$ thus proving our claim.

Our aim is to deduce that $(I - \mathfrak{J}_{\mathcal{I}}(z))^{-1}$ has a simple pole at every point in S , which implies that the same holds for $F_{\text{SAW}}(z)$.

Applying [1, Lemma 9] we obtain for any $w \in S$ that

$$\lim_{z \rightarrow w} (z - w)(I - \mathfrak{J}_{\mathcal{I}}(z))^{-1} = (\eta^T(w)\mathfrak{J}'_{\mathcal{I}}(w)\xi(w))^{-1} \xi(w)\eta^T(w), \quad (21)$$

provided that $\eta^T(w)\mathfrak{J}'_{\mathcal{I}}(w)\xi(w) \neq 0$, where $\eta(w)$ and $\xi(w)$ are left and right eigenvectors of $\mathfrak{J}(w)$ corresponding to 1 normalised so that $\eta^T(w)\xi(w) = 1$.

To verify the condition note that since $\lambda_{\mathcal{I}_t}(w) < 1$, the only non-zero entries of $\eta(w)$ and $\xi(w)$ correspond to elements of \mathcal{I}_p , hence $\eta^T(w)\mathfrak{J}'_{\mathcal{I}}(w)\xi(w) = \eta_p^T(w)\mathfrak{J}'_{\mathcal{I}_p}(w)\xi_p(w)$, where $\eta_p(w)$ and $\xi_p(w)$ are the restrictions to \mathcal{I}_p . In the particular case $w = R$, we have $\eta_p^T(R)\mathfrak{J}'_{\mathcal{I}_p}(R)\xi_p(R) \neq 0$ because by the Perron-Frobenius theorem the entries of $\eta_p(R)$ and $\xi_p(R)$ are strictly positive and at least one entry of $\mathfrak{J}_{\mathcal{I}_p}(z)$ is strictly increasing in z . Thus it suffices to show that $\eta_p^T(w)\mathfrak{J}'_{\mathcal{I}_p}(w)\xi_p(w) = \eta_p^T(R)\mathfrak{J}'_{\mathcal{I}_p}(R)\xi_p(R)$.

To this end, let $D = D(w)$ be an invertible matrix such that $D^{-1}\mathfrak{J}_{\mathcal{I}_p}(w)D = \mathfrak{J}_{\mathcal{I}_p}(R)$. Then $D^{-1}\xi_p(w) = \alpha\xi_p(R)$ and $\eta_p^T(w)D = \alpha^{-1}\eta_p^T(R)$ for some $\alpha \neq 0$ due to the fact that each of $D^{-1}\xi_p(w)$, $\xi_p(R)$, $\eta_p^T(w)D$, $\eta_p^T(R)$ is an eigenvector of an one-dimensional eigenspace of $\mathfrak{J}_{\mathcal{I}_p}(R)$. The scalar factor α is due to the normalisation. Applying again [43, Chapter I, Proposition 6.4] we obtain that $|\mathfrak{J}_{c,c'}(w)| = \mathfrak{J}_{c,c'}(R)$ for all $c, c' \in \mathcal{I}_p$, in other words, we have equality in the triangle inequality. This implies that there exist integers k, ℓ such that t_n is a non-zero Taylor coefficient of $\mathfrak{J}_{c,c'}(z)$ only if $n \in k\mathbb{Z} + \ell$ and w/R is a k th root of unity.

Thus $D^{-1}\mathfrak{J}_{\mathcal{I}_p}(\beta w)D = \mathfrak{J}_{\mathcal{I}_p}(\beta R)$ for every $\beta \in [0, 1]$, hence $D^{-1}\mathfrak{J}'_{\mathcal{I}_p}(w)D = \mathfrak{J}'_{\mathcal{I}_p}(R)$. This in turn implies that $\eta_p^T(w)\mathfrak{J}'_{\mathcal{I}_p}(w)\xi_p(w) = \eta_p^T(R)\mathfrak{J}'_{\mathcal{I}_p}(R)\xi_p(R)$, and we can conclude that (21) holds.

Therefore there exist $b_1, b_2, \dots, b_k \in \mathbb{C}$ such that the function

$$H(z) = F_{\text{SAW}}(z) - \sum_{i=1}^k \frac{b_i}{w_i - z}$$

is analytic in the open disk $|z| < r$, where w_1, w_2, \dots, w_k are the elements of S . This implies that the n th Taylor coefficient h_n of $H(z)$ around 0 satisfies $h_n = O((r - \varepsilon)^{-n})$ for $\varepsilon = (r - R)/2 > 0$. Taylor expanding we see that

$$c_n - \sum_{i=1}^k \frac{b_i}{w_i^{n+1}} = h_n = O((r - \varepsilon)^{-n}).$$

Since each w_i/R is a k th root of unity, $\mu_w = 1/R$ by Remark 8.3, and $c_n \geq 0$, it follows that (1) holds for some $a_1, a_2, \dots, a_k \geq 0$. Since on a quasi-transitive graph $c_n^{1/n}$ converges to μ_w [28], it follows that each a_i is strictly positive.

The second part of the theorem follows now immediately, recalling that on a transitive graph we have $c_n \geq \mu_w^n$ for every $n \geq 1$. \square

References

- [1] S. E. Alm and S. Janson. Random self-avoiding walks on one-dimensional lattices. *Comm. Statist. Stochastic Models*, 6(2):169–212, 1990.
- [2] R. Bauerschmidt, H. Duminil-Copin, J. Goodman, and G. Slade. Lectures on self-avoiding walks. In *Probability and statistical physics in two and more dimensions*, volume 15 of *Clay Math. Proc.*, pages 395–467. Amer. Math. Soc., Providence, RI, 2012.
- [3] I. Benjamini. Self avoiding walk on the seven regular triangulation. *arXiv preprint arXiv:1612.04169*, 2016.
- [4] I. Benjamini and C. Panagiotis. Hyperbolic self avoiding walk. *Electron. Commun. Probab.*, 26:1–5, 2021.
- [5] T. Ceccherini-Silberstein and W. Woess. Growth and ergodicity of context-free languages. *Trans. Am. Math. Soc.*, 354(11):4597–4625, 2002.
- [6] H. Duminil-Copin, S. Ganguly, A. Hammond, and I. Manolescu. Bounding the number of self-avoiding walks: Hammersley-welsh with polygon insertion. *Ann. Probab.*, 48(4), 2020.
- [7] H. Duminil-Copin, A. Glazman, A. Hammond, and I. Manolescu. On the probability that self-avoiding walk ends at a given point. *Ann. Probab.*, 44(2):955–983, 2016.
- [8] H. Duminil-Copin and A. Hammond. Self-avoiding walk is sub-ballistic. *Comm. Math. Phys.*, 324(2):401–423, 2013.
- [9] H. Duminil-Copin and S. Smirnov. The connective constant of the honeycomb lattice equals $\sqrt{2 + \sqrt{2}}$. *Ann. of Math. (2)*, 175(3):1653–1665, 2012.
- [10] M. J. Dunwoody. Cutting up graphs. *Combinatorica*, 2(1):15–23, 1982.
- [11] M. J. Dunwoody and B. Krön. Vertex cuts. *J. Graph Theory*, 80(2):136–171, 2015.
- [12] P. J. Flory. *Principles of Polymer Chemistry*. Cornell University Press, 1953.
- [13] A. Georgakopoulos and A. Wendland. A study of 2-ended graphs via harmonic functions. *arXiv preprint arXiv:2304.13317*, 2023.
- [14] L. A. Gilch and S. Müller. Counting self-avoiding walks on free products of graphs. *Discrete Math.*, 340(3):325–332, 2017.
- [15] G. Grimmett and Z. Li. Self-Avoiding Walks and the Fisher Transformation. *Electron. J. Combin.*, 20(P47), 2013.

- [16] G. Grimmett and Z. Li. Strict inequalities for connective constants of transitive graphs. *SIAM J. Discrete Math.*, 28(3):1306–1333, 2014.
- [17] G. Grimmett and Z. Li. Bounds on connective constants of regular graphs. *Combinatorica*, 35(3):279–294, 2015.
- [18] G. Grimmett and Z. Li. Connective constants and height functions for Cayley graphs. *Trans. Amer. Math. Soc.*, 369(8):5961–5980, 2017.
- [19] G. Grimmett and Z. Li. Self-Avoiding Walks and Amenability. *Electron. J. Combin.*, 24(4):4–38, 2017.
- [20] G. Grimmett and Z. Li. Locality of connective constants. *Discrete Math.*, 341(12):3483–3497, 2018.
- [21] G. Grimmett and Z. Li. Cubic graphs and the golden mean. *Discrete Math.*, 343(111638), 2020.
- [22] G. Grimmett, Y. Peres, and A. Holroyd. Extendable self-avoiding walks. *Ann. Inst. Henri Poincaré D*, 1:61–75, 2014.
- [23] G. R. Grimmett and Z. Li. Self-avoiding walks and connective constants. In *Sojourns in probability theory and statistical physics. III. Interacting particle systems and random walks, a festschrift for Charles M. Newman*, pages 215–241. Singapore: Springer; Shanghai: NYU Shanghai, 2019.
- [24] R. Halin. Über die Maximalzahl fremder unendlicher Wege in Graphen. *Math. Nachr.*, 30:63–85, 1965.
- [25] R. Halin. Automorphisms and endomorphisms of infinite locally finite graphs. *Abh. Math. Sem. Univ. Hamburg*, 39:251–283, 1973.
- [26] R. Halin. S -functions for graphs. *J. Geom.*, 8(1-2):171–186, 1976.
- [27] M. Hamann, F. Lehner, B. Miraftab, and T. Rühmann. A Stallings’ type theorem for quasi-transitive graphs, 2018.
- [28] J. M. Hammersley. Percolation processes. II. The connective constant. *Proc. Cambridge Philos. Soc.*, 53:642–645, 1957.
- [29] J. M. Hammersley and D. J. A. Welsh. Further results on the rate of convergence to the connective constant of the hypercubical lattice. *Q. J. Math.*, 13(1):108–110, 1962.
- [30] T. Hara and G. Slade. The lace expansion for self-avoiding walk in five or more dimensions. *Reviews in Mathematical Physics*, 4(02):235–327, 1992.
- [31] T. Hara and G. Slade. Self-avoiding walk in five or more dimensions I. The critical behaviour. *Communications in Mathematical Physics*, 147(1):101–136, 1992.

- [32] T. Hutchcroft. Self-avoiding walk on nonunimodular transitive graphs. *Ann. Probab.*, 47(5):2801–2829, 2019.
- [33] H. Kesten. On the Number of Self-Avoiding Walks. *Journal of Mathematical Physics*, 4(7):960–969, 1963.
- [34] H. Kesten. On the Number of Self-Avoiding Walks. II. *Journal of Mathematical Physics*, 5(8):1128–1137, 1964.
- [35] D. Krachun and C. Panagiotis. Quantitative sub-ballisticity of self-avoiding walk on the hexagonal lattice. *arXiv preprint arXiv:2310.17299*, 2023.
- [36] F. Lehner and C. Lindorfer. Self-avoiding walks and multiple context-free languages. *Comb. Theory*, 3(1):50, 2023. Id/No 18.
- [37] Z. Li. Positive speed self-avoiding walks on graphs with more than one end. *Journal of Combinatorial Theory, Series A*, 175:105257, 2020.
- [38] N. Madras and G. Slade. *The self-avoiding walk*. Modern Birkhäuser Classics. Birkhäuser/Springer, New York, 2013. Reprint of the 1993 original.
- [39] N. Madras and C. Wu. Self-avoiding walks on hyperbolic graphs. *Combin. Probab. Comput.*, pages 523–548, 2005.
- [40] A. Nachmias and Y. Peres. Non-amenable Cayley graphs of high girth have $p_c < p_u$ and mean-field exponents. *Electron. Commun. Probab.*, 17, 2012.
- [41] C. Panagiotis. Self-avoiding walks and polygons on hyperbolic graphs. *Journal of Graph Theory (to appear)*, 2019.
- [42] N. Robertson and P. D. Seymour. Graph minors. III. Planar tree-width. *J. Combin. Theory Ser. B*, 36(1):49–64, 1984.
- [43] H. H. Schaefer. *Banach Lattices and Positive Operators*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1974.