Navigating Directed Cayley Graphs of Small Diameter: A Potent Solovay-Kitaev Procedure

Henry Bradford

May 21, 2019

Abstract

Let Γ be a group and $(\Gamma_n)_{n=1}^{\infty}$ be a descending sequence of finite-index normal subgroups. We establish explicit upper bounds on the diameters of the directed Cayley graphs of the Γ/Γ_n, under some natural hypotheses on the behaviour of power and commutator words in Γ. The bounds we obtain do not depend on a choice of generating set. Moreover under reasonable conditions our method provides a fast algorithm for navigating directed Cayley graphs. The proof is closely analogous to the the *Solovay-Kitaev procedure*, which only uses commutator words, but also only constructs small-diameter *undirected* Cayley graphs. We apply our procedure to give new directed diameter bounds on finite quotients of a large class of regular branch groups, and of $\mathrm{SL}_2(\mathbb{F}_q[[t]])$ (for q even).

Mathematics Subject Classification (2010). 05C12, 05C20, 20E18.

Keywords. Profinite Groups, Diameters, Spectral Gap

1 Introduction

Let G be a finite group, and $S \subseteq G$ be a generating set. We denote by $B_S^+(n)$ the set of elements of G expressible as positive words of length at most n in S. The *directed diameter of G with respect to S* is defined to be:

$$\mathrm{diam}^+(G,S) = \min\{n \in \mathbb{N} : B_S^+(n) = G\}.$$

The *directed diameter of G*, denoted $\mathrm{diam}^+(G)$, is now defined to be the maximal value of $\mathrm{diam}^+(G,S)$ as S ranges over all generating subsets of G. By contrast, the *(undirected) diameter of G with respect to S* is $\mathrm{diam}(G,S) = \mathrm{diam}^+(G, S \cup S^{-1})$, and the *diameter of G* is the maximal value of $\mathrm{diam}(G,S)$ over S. Clearly $\mathrm{diam}(G) \leq \mathrm{diam}^+(G)$ for any G. The purpose of this paper is to give new upper bounds on $\mathrm{diam}^+(G)$ for certain families of familiar finite groups, to provide fast algorithms for writing elements as positive words of length satisfying this bound, and to outline a procedure for proving results of this type in a more general setting.

1.1 Statement of Results

For the sake of concision in describing the algorithmic aspects of our work, we introduce the following terminology.

Definition 1.1. *Let $(G_n)_n$ be a sequence of finite groups. Let $d, l_n, t_{n,d} \in \mathbb{N}$ with:*

$$\operatorname{diam}^+(G_n) \leq l_n \tag{1}$$

for all n. We say that the directed navigation problem for G_n is solvable for the bound (1) in time $t_{n,d}$ *if there is a deterministic algorithm which, given an index n, a generating set $S_n \subseteq G_n$ satisfying $|S_n| \leq d$ and an element $g \in G_n$, outputs in time at most $t_{n,d}$ a positive word w in S_n of length at most l_n which is equal to g in G_n.*

The technical core of our work is Theorem 2.3 below, which gives bounds of the form (1), and a corresponding fast solution to the directed navigation problem, under natural hypotheses on the groups G_n. We then apply Theorem 2.3 to various concrete sequences $(G_n)_n$ which satisfy these hypotheses. Our first application concerns congruence quotients of the $\mathbb{F}_q[[t]]$-analytic group $\mathrm{SL}_2(\mathbb{F}_q[[t]])$ (q even). In [10] upper bounds on the (undirected) diameter were given for congruence quotients of many analytic (virtually) pro-p groups, including $\mathrm{SL}_d(\mathbb{F}_q[[t]])$ for q odd or $d \geq 3$. For technical reasons related to the structure of the associated Lie algebras, the case $d = 2$, q even fell beyond the scope of the methods of [10]. Therefore our result here is new even for undirected diameters.

Theorem 1.2. *Let $\mathbb{F}_q$ be the finite field of even order q. Let $G(n,q) = \mathrm{SL}_2(\mathbb{F}_q[t]/(t^n))$. Let $\epsilon > 0$. There exist an absolute constant $C > 0$ such that for all $n \in \mathbb{N}$,*

$$\operatorname{diam}^+\big(G(n,q)\big) = O_{q,\epsilon}\big(\log^{C+\epsilon}|G(n,q)|\big). \tag{2}$$

Moreover there exists an absolute constant $C' > 0$ such that the directed navigation problem for $G(n,q)$ is solvable for the bound (2) in time $O_{q,\epsilon}\big(d^{O_q(1)}\log^{C'+\epsilon}|G(n,q)|\big)$.

The proof presented here yields the explicit constants $C = \log(7)/\log(4/3) \approx 6.764$; $C' = 2 + \log(4)/\log(4/3) \approx 6.819$.

We turn next to groups of automorphisms of regular rooted trees. In [11] the (undirected) diameters of congruence quotients of *branch* groups acting on rooted trees were studied. Polylogarithmic upper bounds were obtained in two cases: Grigorchuk's first group and the Gupta-Sidki p-groups. Our work here covers a broad class of branch groups (see Theorem 4.8 below for a full statement). One consequence of our investigations is the following.

Theorem 1.3. *Let p be a prime; $\mathcal{T}$ be the p-ary rooted tree, and Γ be just-infinite regular branch over $K \lhd \Gamma$. Suppose Γ has the congruence subgroup property, and that K/K^m is a p-group. Then there exists $C > 0$ such that for all $n \in \mathbb{N}$,*

$$\operatorname{diam}^+\big(\Gamma/\operatorname{Stab}_\Gamma(n)\big) = O_\Gamma\big(\log^C|\Gamma : \operatorname{Stab}_\Gamma(n)|\big). \tag{3}$$

Moreover there exists $C' > 0$ such that the directed navigation problem for $\Gamma/\operatorname{Stab}_\Gamma(n)$ is solvable for the bound (3) in time $O_\Gamma\big(d^{O_\Gamma(1)}\log^{C'}|\Gamma : \operatorname{Stab}_\Gamma(n)|\big)$.

Regular branch groups, the subgroups $\mathrm{Stab}_\Gamma(n)$ and the embedding $K^m \hookrightarrow K$ will be defined in Section 4. For now let us simply note that $\Gamma/\mathrm{Stab}_\Gamma(n)$ is a transitive imprimitive permutation group of degree p^n.

The hypotheses of Theorem 1.3 are known to hold for many branch groups (see for instance Example 4.11 below). The constants C and C' are given explicitly in terms of certain structural data associated to the group Γ, which in particular cases may be computed. To illustrate this we work through a specific example: the group of automorphisms of the ternary rooted tree introduced by Fabrykowski and Gupta [21].

Theorem 1.4. *Let $\Gamma \leq \mathrm{Aut}(\mathcal{T}_3)$ be the Fabrykowski-Gupta group. Then for all $n \in \mathbb{N}$,*

$$\mathrm{diam}^+\big(\Gamma/\mathrm{Stab}_\Gamma(n)\big) = O\big(\log^C|\Gamma/\mathrm{Stab}_\Gamma(n)|\big) \tag{4}$$

where $C = 6\log(19)/\log(3) \approx 16.081$. Moreover the directed navigation problem for $\Gamma/\mathrm{Stab}_\Gamma(n)$ is solvable for the bound (4) in time $O\big(d+\log^{C'}|\Gamma/\mathrm{Stab}_\Gamma(n)|\big)$, where $C' = 1 + 6\log(7)/\log(3) \approx 11.627$.

Γ will be defined in Section 5. Once again a polylogarithmic bound for this group is new even for undirected diameters.

It is very likely that Theorem 2.3, or variants thereof, will also be applicable to many other groups than those covered by Theorems 1.2, 1.3 and 1.4.

As noted above, for any finite group G we have $\mathrm{diam}(G) \leq \mathrm{diam}^+(G)$. Somewhat surprisingly, there is also a converse inequality due to Babai.

Theorem 1.5 ([3] Corollary 2.3)**.** *Let G be a finite group. Then:*

$$\mathrm{diam}^+(G) = O\big(\mathrm{diam}(G)\log|G|^2\big).$$

As a result, all groups with polylogarithmic diameter also have polylogarithmic directed diameter, and where the degree of the polylogarithm in the former is explicitly known, so is that in the latter. In spite of this, there are advantages to deriving directed diameter bounds without the use of Theorem 1.5, even when good (undirected) diameter bounds are known. In particular, the proof of Theorem 1.5 is non-constructive, so does not yield any non-trivial solution to the directed navigation problem (see [3] Section 5 for a discussion of this and related problems). One can be very confident that the potent SKP (either in the form of Theorem 2.3 or with modifications) will provide solutions to the directed navigation problem for many of the other $\mathbb{F}_q[[t]]$-analytic groups and Nottingham groups of finite fields studied in [10]. These solutions will moreover witness directed diameter bounds qualitatively similar, if quantitatively weaker, than those obtained by combining Theorem 1.5 with the results of [10, 11]. Nevertheless, owing to the availability of Theorem 1.5, we have opted predominantly to illustrate the implementation of the potent SKP with examples for which polylogarithmic *undirected* diameter bounds were not previously known.

1.2 Background and Outline of the Proof

Estimating the diameters of finite Cayley graphs has been a subject of widespread interest for many years. Motivation comes from the problem of constructing efficient communication networks [24]; analysis of algorithms in computational group theory [2], and various combinatorial puzzles (card-shuffling; generalizations of the Rubik's cube; the towers of Hanoi; pebble motions on graphs and so on) [14, 27]. Owing to the concrete nature of these applications, one often seeks not only good diameter bounds, but also fast algorithms that express a group element as a word in a generating set, the length of which satisfies the bound. This is the navigation problem. Fortunately many of the results on diameter in permutation groups have essentially algorithmic proofs [4, 5]. Meanwhile the navigation problem in $\mathrm{SL}_d(\mathbb{F}_p)$ (and more generally Chevalley groups over $\mathbb{F}_p$ and other finite rings) was studied in [28, 32, 26], where fast (sometimes probabilistic) algorithms were described and analyzed for particular generating sets (though a good solution to the navigation problem for groups of Lie type realizing the best known diameter bounds for arbitrary or generic generators remains elusive). The navigation problem is also of relevance in cryptography, in that efficient solutions are an obstruction to the construction of secure Cayley hash functions (see [12, 31] for a discussion).

In spite of this impressive progress, much less is known about the directed navigation problem, as was noted in [3]. This is an unsatisfactory state of affairs, as solutions to many combinatorial puzzles are better modeled by directed as opposed to undirected navigation (consider for instance the practical difficulty of inverting a large-order riffle shuffle of a deck of cards). Further, directed navigation is more relevant to the cryptanalysis of Cayley hash functions, in which a bit-stream is encoded as a *positive* word in generators. Of the few results available, one of the most impressive is [33], which addresses the directed navigation problem for the symmetric group with respect to random pairs of generators. In this paper we introduce a set of tools that allow one to attack the directed navigation problem under certain group-theoretic conditions.

The inspiration for our results comes from the *Solovay-Kitaev procedure*. Given a compact metric group Γ and a subset S generating a dense subgroup, the SKP provides a framework for constructing a word w in S which approximates a given element $g \in \Gamma$ to a prescribed level of accuracy. Moreover, the length of w in the word metric defined by S is bounded in terms of the distance in Γ between g and w. The first examples to which the SKP was applied were the groups $\mathrm{SU}(k)$, where the problem of approximating arbitrary elements by words in a generating set was motivated by considerations coming from quantum computation [13]. The SKP has since been applied to other Lie groups (for instance by Dolgopyat [20], who independently discovered a version of the SKP and employed it to elucidate spectral properties of semisimple Lie groups). It was however also soon noticed that similar techniques were relevant to finitely generated (abstract or profinite) groups Γ equipped with a profinite metric, and that in this setting approximating elements by short words is equivalent to proving good diameter bounds for finite quotients of Γ. This idea has been

exploited in several papers [23, 16, 18, 10, 11]. Moreover the SKP gives a fast solution to the navigation problem: this is described explicitly in [13, 23, 18], and can easily be derived from the proofs of the results in [16, 10, 11].

How does the SKP work? We assume that there is a neighbourhood U of the identity in Γ satisfying two hypotheses. The first hypothesis is that every element z of U lying sufficiently close to the identity is approximable by a product of (a bounded number of) commutators $[x_i, y_i]$, where $x_i, y_i \in U$ are significantly further from the identity than z is. The second, complementary, hypothesis is that for $x, y \in U$, the commutator $[x, y]$ is significantly closer to the identity than x and y. It follows from the latter that if the pairs (x, y) and $(\tilde{x}, \tilde{y})$ are close, then $[x, y]$ and $[\tilde{x}, \tilde{y}]$ are even closer. If $z \in U$ is the error in our existing verbal approximation $\tilde{g}$ to $g \in \Gamma$; $[x_1, y_1] \cdots [x_A, y_A]$ is an approximation to z (which exists by the first hypothesis) and $\tilde{x}_i, \tilde{y}_i$ are verbal approximations to x_i, y_i (which we may assume exist by induction), then $\tilde{g}[\tilde{x}_1, \tilde{y}_1] \cdots [\tilde{x}_A, \tilde{y}_A]$ is a better verbal approximation to g.

In the present paper we modify this strategy, in that we replace the first hypothesis by the requirement that z is approximable by a product of kth powers y_i^k, for $y_i \in U$ and $k \geq 2$ fixed. To implement the induction step, we must then also strengthen the second hypothesis, by requiring that taking kth powers moves elements of U closer to the identity, as well as commutators. As we shall see below (Remark 2.2), a very natural setting in which the second hypothesis holds is when $k = p$ is a prime and Γ is a residually p-finite group, equipped with the profinite metric defined by the mod-p dimension series. Because it relies heavily on properties of proper powers, it seems appropriate to term the new method a *potent* Solovay-Kitaev procedure. The fact that it yields a *directed* diameter bound follows from the fact that the proper powers used to express elements close to the identity are *positive words*.

A version of the SKP was also used by Bourgain and Gamburd [8, 9] (in conjunction with other tools) to produce new examples of *expander Cayley graphs*. Expanders are sparse finite regular graphs with very strong connectivity and mixing properties. For instance they have logarithmic (undirected) diameter and, which is more, the endpoints of paths of logarithmic length are *equidistributed* over the graph. Expanders have remarkable and diverse applications across pure mathematics, communication theory and theoretical computer science; we refer the reader to the excellent survey articles [25, 30] for an overview of these. It would be very interesting to investigate the possibility of adapting the *potent* SKP to construct new examples of expanders.

In spite of the obvious analogies between the original SKP and our new potent variant, and the relevance of the former to approximation problems in real and complex Lie groups, the potent SKP appears to be predominantly a "non-analytic" phenomenon: raising elements of a real or complex Lie group to a proper power does not generically move them closer to the identity. Indeed the problem of approximating an arbitrary element in a Lie group by a short positive word in an arbitrary generating set appears to be open. As noted in [13], a solution to this problem for $\mathrm{SU}(d)$ would be of interest in the context of quantum computation: the hypothesis of a symmetric generating set, although

group-theoretically natural, has no clear justification when the set of generating matrices is interpreted as the instruction set of a quantum computer. The potent SKP *does* yield directed diameter bounds for quotients of p-adic analytic groups, by exploiting their connection with *powerful pro-p groups*, but the diameter bounds are rather weak: for instance for the groups $G(d,p,n) = \mathrm{SL}_d(\mathbb{Z}/p^n\mathbb{Z})$ we would obtain $\mathrm{diam}^+\big(G(d,p,n)\big) = O_{p,d}\big(|G(d,p,n)|^{1/(d^2-1)}\big)$, which compares poorly with the polylogarithmic undirected diameter bounds for these groups in [10]. We discuss the relevance of the potent SKP to p-adic analytic groups futher in Section 6.

The paper is structured as follows. In Section 2 we develop the potent Solovay Kitaev procedure in an abstract setting, giving sufficient conditions on the behaviour of power and commutator words in the sequence $(\Gamma_i)_i$ for a good upper bound on the $\mathrm{diam}^+(\Gamma/\Gamma_i)$ to hold. In Section 3 we prove Theorem 1.2. In Section 4 we investigate the potent SKP in relation to branch groups, and prove Theorem 1.3. In Section 5 we deduce 1.4 from the resuts of Section 4. In Section 6 we derive from the potent SKP a weak upper bound on directed diameters in quotients of p-adic analytic groups. In Section 7 we discuss some implications of our results for spectral gaps and mixing times of random walks. Finally in Section 8 we discuss some complementary lower bounds on diameters of finite groups.

2 The Procedure

In this Section we describe the potent Solovay-Kitaev Procedure in an abstract group-theoretic context. The Procedure is expressed in Theorem 2.3. Sections 3 and 5 will then be devoted to proving that the hypotheses of Theorem 2.3 hold in the relevant settings such that Theorems 1.2 and 1.4 follow immediately.

We start with an observation to the effect that, given an approximation to a group element, the kth power of the element is well-approximated by the kth power of the approximation. For $N \leq G$ denote by $\mathcal{U}_k(N)$ the subgroup of G generated by all kth powers of elements of N. Note that $\mathcal{U}_k(N)$ is normal in G whenever N is.

Lemma 2.1. *Let Γ be a group; let $M, N \lhd \Gamma$ and let $k \in \mathbb{N}_{\geq 2}$. Then for all $g \in M, h \in N$,*

$$(gh)^k g^{-k} \in [M,N]\mathcal{U}_k(N).$$

Proof. Let $\langle [g,h] \rangle^\Gamma$ be the normal closure of $[g,h]$ in Γ. Then $\langle [g,h] \rangle^\Gamma \leq [M,N]$ (since $M, N \lhd \Gamma$), and $(gh)^k g^{-k} h^{-k}$ is clearly trivial in $\Gamma/\langle [g,h] \rangle^\Gamma$ (since the images of g and h in the latter quotient commute). Thus $(gh)^k g^{-k} h^{-k} \in [M,N]$ and the result follows. □

The conclusion of Lemma 2.1 will be useful in situations where $[M,N]\mathcal{U}_k(N)$ is much smaller than N.

Example 2.2. *Let $(\Gamma_n)_{n=1}^\infty$ be a descending sequence of finite-index normal subgroups of Γ. Suppose that for all $m, n \in \mathbb{N}$:*

(i) $[\Gamma_n, \Gamma_m] \subseteq \Gamma_{n+m}$;

(ii) $\mho_k(\Gamma_n) \subseteq \Gamma_{kn}$.

Let $n \leq m$ and let $g \in \Gamma_n$, $h \in \Gamma_m$. Then by Lemma 2.1:

$$(gh)^k \equiv g^k \mod \Gamma_{n+m}.$$

It is a classical fact that $(\Gamma_n)_{n=1}^{\infty}$ satisfies conditions (i) and (ii) above with $k = p$ a prime when Γ_n is the mod-p dimension series *of Γ. Recall that the latter is the sequence $(D_n(\Gamma))_{n=1}^{\infty}$ of normal subgroups of Γ given by:*

$$D_n(\Gamma) = \{g \in \Gamma : g - e \in I^n\}$$

where I is the augmentation ideal *of the group algebra $\mathbb{F}_p\Gamma$, defined to be the kernel of the* augmentation mapping $\phi : \mathbb{F}_p\Gamma \to \mathbb{F}_p$, *which is given by:*

$$\phi(\sum{}' \lambda_g \cdot g) = \sum{}' \lambda_g.$$

Alternatively, $D_n(\Gamma)$ may be defined recursively by $D_1(\Gamma) = \Gamma$, $D_{n+1}(\Gamma) = [\Gamma, D_n(\Gamma)]\mho_p(D_{\lceil (n+1)/p \rceil}(\Gamma))$.

Another example of a sequence $(\Gamma_n)_{n=1}^{\infty}$ in which conditions (i) and (ii) above hold, and which will be relevant to Theorem 1.2, is given below (see Lemma 3.2).

Theorem 2.3. *Let $(M_n)_{n=1}^{\infty}$, $(N_n)_{n=1}^{\infty}$ be sequences of finite-index normal subgroups in Γ. Let $(A_n)_{n=1}^{\infty}, (k_n)_{n=1}^{\infty}$ be a sequence of positive integers. Suppose that for all $n \in \mathbb{N}$:*

(i) $N_n \leq M_n$;

(ii) $[M_n, N_n] \leq N_{n+1}$;

(iii) $\mho_{k_n}(N_n) \leq N_{n+1}$;

(iv) For all $z \in N_n$, there exist $y_1, \ldots, y_{A_n} \in M_n$ such that:

$$y_1^{k_n} \cdots y_{A_n}^{k_n} z^{-1} \in N_{n+1}. \tag{5}$$

Then for all $n \in \mathbb{N}$:

$$\operatorname{diam}^+(\Gamma/N_n) \leq l_n = |\Gamma : N_1| \prod_{i=1}^{n-1} (1 + A_i k_i). \tag{6}$$

Further suppose that for all $m \in \mathbb{N}$, the times needed to compute:

(a) The product gh of given input elements $g, h \in \Gamma/N_m$;

(b) The inverse g^{-1} of a given input element $g \in \Gamma/N_m$;

(c) $N_m y_1, \ldots, N_m y_{A_n}$, given input $1 \leq n \leq m$ and $N_m z$, where $y_i \in M_n$ and $z \in N_n$ are as in (5)

are at most $f(m)$. Then the directed navigation problem for Γ/N_n is solvable for the bound (6) in time:

$$f(n)\Big(C|S|^{|\Gamma:N_1|+1}\prod_{i=1}^{n-1}(A_i+1)+\sum_{i=1}^{n-1}(A_ik_i+3)\prod_{j=i}^{n-2}(A_j+1)\Big) \tag{7}$$

for $C>0$ an absolute constant.

Proof. First let us establish the diameter bound. For $n=1$ the conclusion is trivial. Suppose by induction that $\mathrm{diam}^+(\Gamma/N_n)\leq l_n$. Let $S_{n+1}\subseteq\Gamma/N_{n+1}$ be a generating set and let S_n be the image of S_{n+1} in Γ/N_n. Then S_n generates Γ/N_n. Let $g\in\Gamma/N_{n+1}$. By inductive hypothesis there exists $w\in B^+_{S_{n+1}}(l_n)$ such that $z=w^{-1}g\in N_n$. By hypothesis (iv) there exist $y_1,\ldots,y_{A_n}\in M_n$ such that $y_1^{k_n}\cdots y_{A_n}^{k_n}z^{-1}\in N_{n+1}$.

By inductive hypothesis there exist, for $1\leq i\leq A_n$, $\tilde{y}_i\in B^+_{S_{n+1}}(l_n)$ such that $y_i\tilde{y}_i^{-1}\in N_n$. Combining hypotheses (ii) and (iii) with Lemma 2.1, we have $y_i^{k_n}(\tilde{y}_i)^{-k_n}\in N_{n+1}$. Then:

$$g=wz\equiv w(\tilde{y}_1)^{k_n}\cdots(\tilde{y}_{A_n})^{k_n} \mod N_{n+1}$$

and $w(\tilde{y}_1)^{k_1}\cdots(\tilde{y}_{A_n})^{k_n}\in B^+_{S_{n+1}}(l_n(1+A_nk_n))$. The diameter bound follows by induction.

We now describe and analyze an algorithm $\texttt{APPROX}(n,i,g,S)$, which takes as input $n,i\in\mathbb{N}$ with $i\leq n$, $g\in\Gamma/N_n$ and $S\subseteq\Gamma/N_n$, and outputs both a positive word $\tilde{w}\in F(S)$ of length at most l_i and the evaluation w of $\tilde{w}$ in Γ/N_n, with the property that $g\equiv w \mod N_i$. The algorithm required by the statement of the Theorem will be $\texttt{APPROX}(n,n,g,S)$.

First note that $\texttt{APPROX}(n,1,g,S)$ runs in time $O(|S|^{|\Gamma:N_1|+1}f(n))$: we may simply compute all products of elements in S of length at most $|\Gamma:N_1|$; one of these will agree with g modulo N_1.

Now we employ recursion. Given $1\leq i\leq n-1$, let $(\tilde{w}_i,w_i)$ be the output of $\texttt{APPROX}(n,i,g,S)$. Then $z=w_i^{-1}g\in N_i$. Compute $y_1,\ldots,y_{A_i}\in M_i$ as in (5); as hypothesized in (c) above, this requires time at most $f(n)$.

Let $(\tilde{v}_{i,j},v_{i,j})$ be the output of $\texttt{APPROX}(n,i,y_j,S)$. The output of $\texttt{APPROX}(n,i+1,g,S)$ is $(\tilde{w}_{i+1},w_{i+1})$, where $\tilde{w}_{i+1}=\tilde{w}_i\tilde{v}_{i,1}^{k_i}\cdots\tilde{v}_{i,A_i}^{k_i}$ and $w_{i+1}=w_iv_{i,1}^{k_i}\cdots v_{i,A_i}^{k_i}$. Our proof of the diameter bound above witnesses that $\tilde{w}_{i+1},w_{i+1}$ have the required properties.

Finally take $t_{n,i}\in\mathbb{N}$ such that $\texttt{APPROX}(n,i,g,S)$ runs in time at most $t_{n,i}$ for all g,S. As noted above, we may take:

$$t_{n,1}=C|S|^{|\Gamma:N_1|+1}f(n)$$

For $1\leq i\leq n-1$ note that to implement $\texttt{APPROX}(n,i+1,g,S)$ we must call $\texttt{APPROX}(n,i,h,S)$ for A_i+1 elements h, and carry out A_ik_i+1 computations of type (a) and one each of type (b) and (c). We may therefore take:

$$t_{n,i}=(A_i+1)t_{n,i}+(A_ik_i+3)$$

and the conclusion (7) follows. □

Remark 2.4. (i) The statement of Theorem 2.3 is more general than we shall need in the setting of Theorems 1.2 and 1.4, where (k_n) will be a constant sequence, and (A_n) will be periodic. We state Theorem 2.3 in this general form to emphasize the adaptability of the potent SKP, and its potential applicability to problems much more diverse than the applications we give here.

(ii) Equally, additional refinements to Theorem 2.3 are possible, which slightly improve the diameter bounds and the runtime of our algorithm. For instance, suppose there exists a constant $n_0 \in \mathbb{N}$ such that for all n, $M_{n+n_0} \leq N_n$. Then for any generating set $S \subseteq \Gamma/N_n$ and any $1 \leq i \leq n-1$, we have $N_i/N_{i+1} \subseteq B_S^+(L_i)N_{i+1}/N_{i+1}$, where $L_0 = |\Gamma : N_1|$, and $L_i = A_i k_i(L_{i-n_0} + \cdots + L_{i-1})$ for $i \geq 1$ (with $L_i = 0$ for negative indices). Thus:

$$\operatorname{diam}^+(\Gamma/N_n) \leq L_0 + \cdots + L_{n-1}. \tag{8}$$

To see that this is a stronger upper bound, note that the bound (6) may be expressed as $l_n = L'_0 + \cdots L'_{n-1}$, where $L'_0 = |\Gamma : N_1|$ and $L'_i = A_i k_i(L'_0 + \cdots + L'_{i-1})$ for $i \geq 1$.

(iii) The initial step of our induction, which yields the trivial bounds $\operatorname{diam}^+(\Gamma/N_1) \leq |\Gamma : N_1|$ and a solution to the directed navigation problem for Γ/N_1 in time $O(|S|^{|\Gamma:N_1|+1})$, is far from optimal in many cases. For instance $\operatorname{diam}^+(\mathrm{SL}_2(q)) = O(\log(q)^c)$ for an absolute constant c [17], which enables improvements to the constants appearing in our Theorem 1.2.

(iv) In principle there is nothing to prevent one from using more general word maps than powers in Theorem 2.3. Indeed for F a free group and (w_n) a sequence of non-trivial reduced words of lengths (k_n) in F, suppose we replace hypothesis (iii) of Theorem 2.3 with the claim that all evaluations of w_n in N_n lie in N_{n+1}, and hypothesis (iv) with the claim that every element z of N_n may be written, modulo N_{n+1}, as the product of at most A_n elements $w_n(\mathbf{y}_i)$, where $\mathbf{y}_i$ is a tuple of elements in M_n. Then arguing, mutatis mutandis, as in the proof of Theorem 2.3, we have $\operatorname{diam}(\Gamma/N_n) \leq l_n$ as in (6). If we further assume the w_n to be positive words, then we have the same bound for $\operatorname{diam}^+(\Gamma/N_n)$ (with corresponding statements for navigation under amended hypotheses). It would be interesting to investigate the applications of a procedure using more general words in contexts where using only power or commutator words proves ineffective.

3 Proofs for $\mathrm{SL}_2(\mathbb{F}_q[[t]])$

Let $\mathbb{F}_q$ be a finite field of even order q, let $\mathbb{F}_q[[t]]$ be the power series ring of $\mathbb{F}_q$ and let $\Gamma = \mathrm{SL}_2(\mathbb{F}_q[[t]])$. For $n \in \mathbb{N}$, let:

$$K_n = \Gamma \cap (I_2 + t^n \mathbb{M}_2(\mathbb{F}_q[[t]])) = \ker(\pi_n),$$

where $\mathbb{M}_2$ denotes the algebra of 2-by-2 matrices over a given ring and $\pi_n : \Gamma \twoheadrightarrow \mathrm{SL}_2(\mathbb{F}_q[t]/(t^n))$ is the congruence map. Hence $(K_n)_n$ is a descending chain of finite-index normal subgroups of Γ.

Lemma 3.1. *Let $n, m \in \mathbb{N}$. Then:*

(i) $[K_n, K_m] \subseteq K_{n+m}$;

(ii) $\mathfrak{V}_2(K_n) \subseteq K_{2n}$.

Proof. Let $X, \tilde{X}, Y, \tilde{Y} \in \mathbb{M}_2(\mathbb{F}_q[[t]])$ be such that $g = I_2 + t^n X$, $g^{-1} = I_2 + t^n \tilde{X}$, $h = I_2 + t^m Y$, $h^{-1} = I_2 + t^m \tilde{Y}$.

(i) Since $g^{-1} \cdot g = h^{-1} \cdot h = I_2$,

$$X + \tilde{X} + t^n \tilde{X} X = Y + \tilde{Y} + t^m \tilde{Y} Y = 0. \tag{9}$$

Thus:

$$\begin{aligned} [g, h] &= (I_2 + t^n \tilde{X})(I_2 + t^m \tilde{Y})(I_2 + t^n X)(I_2 + t^m Y) \\ &\equiv I_2 + t^n (X + \tilde{X}) + t^m (Y + \tilde{Y}) + t^{2n} \tilde{X} X + t^{2m} \tilde{Y} Y \\ &\equiv I_2 \mod t^{n+m} \text{ (by (9))} \end{aligned}$$

so $[g, h] \in K_{n+m}$.

(ii) $g^2 = (I_2 + t^n X)^2 = I_2 + t^{2n} X^2 \in K_{2n}$ (since $\mathrm{char}(\mathbb{F}_q) = 2$).

□

Lemma 3.2. *Let $z \in K_{3n}$. Then there exist $y_1, y_2, y_3 \in K_n$ such that:*

$$y_1^2 \cdot y_2^2 \cdot y_3^2 \cdot z^{-1} \in K_{4n}.$$

Proof. There exist $a, b, c, d \in \mathbb{F}_q[[t]]$ such that:

$$z = I_2 + t^{3n} \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

Then $1 = \det(z) = 1 + t^{3n}(a + d) + t^{6n}(ad - bc)$, so $a \equiv d \mod t^{3n}$. Set:

$$\begin{aligned} y_1 &= \begin{pmatrix} 1 + \overline{a} t^{3n} & t^n \\ \overline{a} t^{2n} & 1 \end{pmatrix}, \\ y_2 &= \begin{pmatrix} 1 + t^n & \overline{b} t^n \\ 0 & (1 + t^n)^{-1} \end{pmatrix}, \\ y_3 &= \begin{pmatrix} (1 + t^n)^{-1} & 0 \\ \overline{c} t^n & 1 + t^n \end{pmatrix} \end{aligned}$$

(for any $\overline{a} \equiv a, \overline{b} \equiv b, \overline{c} \equiv c \mod t^n$), so that $y_1, y_2, y_3 \in K_n$. We compute:

$$y_1^2 \cdot y_2^2 \cdot y_3^2 \equiv \begin{pmatrix} 1 + t^{3n} a & t^{3n} b \\ t^{3n} c & 1 + t^{3n} a \end{pmatrix} \equiv z \mod t^{4n}$$

as required. □

Remark 3.3. *It is clear from the proof of Lemma 3.2 that there is an algorithm which, given $z \in K_{3n}$, computes the y_1, y_2, y_3 in time $O(n)$ (by reading the coefficients a, b, c modulo t^n in our expression for z and substituting into our expressions for y_1, y_2, y_3).*

Proof of Theorem 1.2. Let $(\alpha_n)_n$, $(\beta_n)_n$ be ascending sequences of integers such that (a) $\alpha_n + \beta_n \geq \beta_{n+1}$ and (b) $\beta_n \geq 3\alpha_n$. Note that (a) and (b) together imply (c) $4\beta_n/3 \geq \beta_{n+1}$.

We define $M_n = K_{\alpha_n}, N_n = K_{\beta_n} \leq \Gamma$ and set $A_n = 3$, $k_n = 2$. We check that these sequences satisfy the hypotheses of Theorem 2.3. Hypothesis (i) is clear; hypotheses (ii) and (iii) follow from Lemma 3.1 and the above conditions, and hypothesis (iv) follows from Lemma 3.2 and condition (c) above.

We therefore have:

$$\begin{aligned}\mathrm{diam}^+(\Gamma/N_n) &\leq 7^{n-1}|\Gamma : N_1| \\ &= O_{\beta_1,q}\big(\log|\Gamma : N_n|^{n\log(7)/\log(\beta_n)}\big)\end{aligned}$$

(since $|SL_2(\mathbb{F}_q[t]/(t^m))| = (q^2-1)q^{3m-2}$). The bound (2) for this subsequence of $G(n,q) = \Gamma/K_n$ follows from the easy observation that for all $\epsilon > 0$ we may take $\alpha_n, \beta_n = \Omega_\epsilon((\frac{4}{3} - \epsilon)^n)$.

For the directed navigation problem, we observe that multiplying two elements of $\mathrm{SL}_2(\mathbb{F}_q[t]/(t^n))$ involves $O(n^2)$ multiplications and additions of pairs of elements of $\mathbb{F}_q$, so may be achieved in time $O_q(n^2)$. Inversion involves only the rearrangement of co-ordinates so may be accomplished in linear time, as may computing the approximations y_i to a given z (by Remark 3.3). We therefore satisfy conditions (a), (b) and (c) of Theorem 2.3 with $f(n) = \beta_n^2$, so from (7), we have a solution in time:

$$O\big(\beta_n^2 4^n |S|^{|G:K_{\beta_1}|+1}\big) = O\big(|S|^{O_{q,\epsilon}(1)}\log|\Gamma : N_n|^{2+\frac{\log(4)}{\log(4/3-\epsilon)}}\big).$$

The conclusions of Theorem 1.2 for general $G(n,q) = \Gamma/K_n$ follow from the above bounds for Γ/N_m by taking m such that $N_m \leq K_n \leq N_{m-1}$ and comparing the indices of N_m and K_n in Γ. □

Remark 3.4. Set $\Gamma = \mathrm{SL}_d(\mathbb{F}_q[[t]])$ for $d \geq 3$, q even, and again take:

$$K_n = \Gamma \cap (I_2 + t^n \mathbb{M}_2(\mathbb{F}_q[[t]])).$$

Slightly modifying the above construction for SL_2, it is easy to show that every element of K_{3n} may be written modulo K_{4n} as the product of four squares of elements in K_n (provided q is sufficiently large, depending on d). It follows that $\mathrm{diam}^+(\Gamma/K_n) = O_{d,q,\epsilon}(\log|\Gamma : K_n|^{C+\epsilon})$ for all $\epsilon > 0$, where $C = \log(9)/\log(4/3) \approx 7.638$. For comparison, the results of [10] yield $\mathrm{diam}(\Gamma/K_n) = O_{d,q}(\log|\Gamma : K_n|^{C'})$, where $C' = \log(44)/\log(2) \approx 5.459$. Thus the bound for diam^+ obtained by combining the latter bound for diam with Theorem 1.5 is asymptotically very slightly better than that obtained by applying the potent Solovay-Kitaev procedure directly, but does not provide a solution to the directed navigation problem, which the potent SKP does.

4 Branch Groups

For $m \geq 2$ define the *m-ary rooted tree* to be the graph $\mathcal{T}_\mathcal{A}$ with vertex set $\mathcal{A}^*$ the set of formal positive words on alphabet $\mathcal{A}$, a set of cardinality m, and edges (w, wa) for $w \in \mathcal{A}^*$ and $a \in \mathcal{A}$. The set $V_n = \mathcal{A}^n$ of words of length n in $\mathcal{A}$ (that is, the set of vertices of $\mathcal{T}_\mathcal{A}$ at distance n from the *root vertex*, represented by the empty word) is known as the *nth level set* of $\mathcal{T}_\mathcal{A}$. In particular $V_n V_m = V_{n+m}$ and $V_1 = \mathcal{A}$, and we will use these notations interchangeably.

The group $\mathrm{Aut}(\mathcal{T}_\mathcal{A})$ of graph automorphisms of $\mathcal{T}_\mathcal{A}$ is precisely the set of permutations of $\mathcal{A}^*$ which respect prefixes, and in particular fixes the root vertex. The kernel of the action of $\mathrm{Aut}(\mathcal{T}_\mathcal{A})$ on the nth level set V_n will be called the *nth level stabiliser* and denoted $\mathrm{Stab}(n)$; it is naturally isomorphic to the direct product $\mathrm{Aut}(\mathcal{T}_\mathcal{A})^{V_n}$. If $\Gamma \leq \mathrm{Aut}(\mathcal{T}_\mathcal{A})$ we write $\mathrm{Stab}_\Gamma(n)$ for $\Gamma \cap \mathrm{Stab}(n)$.

Definition 4.1. *A subgroup Γ of $\mathrm{Aut}(\mathcal{T}_\mathcal{A})$ is said to possess* the congruence subgroup property *if, for every $H \leq \Gamma$ of finite index, there exists $n \in \mathbb{N}$ such that $\mathrm{Stab}_\Gamma(n) \leq H$.*

For any $\phi \in \mathrm{Aut}(\mathcal{T}_\mathcal{A})$, there exists a unique $\sigma_\phi \in \mathrm{Sym}(\mathcal{A})$ such that for any $x \in \mathcal{A}$, there exists $\phi_x \in \mathrm{Aut}(\mathcal{T}_\mathcal{A})$ such that:

$$\phi(xw) = \sigma_\phi(x)\phi_x(w), \text{ for all } w \in \mathcal{A}^*. \tag{10}$$

Moreover the ϕ_x are uniquely determined by (10). The induced map $\psi : \phi \mapsto (\phi_x)_{x\in\mathcal{A}} \cdot \sigma_\phi$ gives an isomorphism $\mathrm{Aut}(\mathcal{T}_\mathcal{A}) \to \mathrm{Aut}(\mathcal{T}_\mathcal{A}) \wr \mathrm{Sym}(\mathcal{A})$. Note that the level stabilisers may be described recursively by $\mathrm{Stab}(0) = \mathrm{Aut}(\mathcal{T}_\mathcal{A})$ and $\mathrm{Stab}(n+1) = \psi^{-1}(\mathrm{Stab}(n)^\mathcal{A})$.

Of particular interest among the subgroups of $\mathrm{Aut}(\mathcal{T}_\mathcal{A})$ are those whose action on $\mathcal{T}_\mathcal{A}$ is *branch*. Here we focus specifically on *regular* branch groups.

Definition 4.2. *Let $\Gamma \leq \mathrm{Aut}(\mathcal{T}_\mathcal{A})$. Γ is* regular branch *if:*

(i) The action of Γ on $\mathcal{A}$ is transitive;

(ii) For all $x \in \mathcal{A}$, $\{\phi_x : \phi \in \mathrm{Stab}_\Gamma(1)\} = \Gamma$;

(iii) Γ has a finite-index subgroup K such that $K^\mathcal{A} \leq \psi(K)$.

We will simply say that a group Γ branches over K *when the alphabet $\mathcal{A}$ and the action of Γ on $\mathcal{A}^*$ is clear.*

It follows from (i) and (ii) that Γ is transitive on every V_n. Henceforth we usually suppress the map ψ from expressions and identify subgroups of Γ with their image under ψ, so we may for instance speak of $K^\mathcal{A}$ as a subgroup of K; $\mathrm{Stab}_\Gamma(n)$ as a subgroup of Γ^{V_n} and so on.

Lemma 4.3. *Let $H \lhd \Gamma$ with $H \leq K$. Then $H^{V_n} \lhd \Gamma$ for all $n \geq 1$.*

Proof. By induction (and replacing H by $H^{V_{n-1}}$) we may assume $n = 1$. Let $\mathbf{h} = (h_x)_{x\in\mathcal{A}} \in H^\mathcal{A}$ and let $g \in \Gamma$. Then there exists $\sigma \in \mathrm{Sym}(\mathcal{A})$ and $\mathbf{g} = (g_x)_{x\in\mathcal{A}} \in \Gamma^\mathcal{A}$ such that $g = \mathbf{g}\sigma$. Then for all $x \in \mathcal{A}$, $(\mathbf{h}^g)_x = h_{\sigma(x)}^{g_{\sigma(x)}} \in H$. □

Corollary 4.4. *If Γ is regular branch, then Γ branches over a finite-index* normal *subgroup.*

Proof. Suppose Γ branches over the finite-index subgroup L. Let $K \leq L$ be the normal core of L in Γ. Then $K^{\mathcal{A}} \leq L^{\mathcal{A}} \leq L$. By Lemma 4.3, $K^{\mathcal{A}}$ is normal in Γ, so by definition of the normal core, $K^{\mathcal{A}} \leq K$. □

$\mathrm{Aut}(\mathcal{T}_{\mathcal{A}})$ is naturally a profinite group, with $\mathrm{Stab}(n)$ a neighbourhood basis at the identity. For $\Gamma \leq \mathrm{Aut}(\mathcal{T}_{\mathcal{A}})$, the closure $\overline{\Gamma}$ of Γ in this topology is isomorphic to $\operatorname{proj\,lim} \Gamma/\operatorname{Stab}_\Gamma(n)$. It is clear that if Γ branches over K, then $\overline{\Gamma}$ branches over $\overline{K}$. From this description, we have the next fact, which we will need to guarantee that our power-word approximations exist.

Proposition 4.5. *Let Γ be regular branch. Then Γ does not have finite exponent.*

Proof. Let $m \in \mathbb{N}$ and suppose (for a contradiction) that the exponent $\exp(\Gamma)$ divides m. Then so too does every $\exp(\Gamma/\operatorname{Stab}_\Gamma(n))$, and hence also $\exp(\overline{\Gamma})$. But by [1][Corollary 1.4] $\overline{\Gamma}$ contains a free subgroup. □

Lemma 4.6. *If Γ is regular branch, then there exist $C_1, C_2 > 0$ such that for all $m \in \mathbb{N}$,*

$$|\Gamma : \operatorname{Stab}_\Gamma(m)| \geq \exp\big(C_1 \exp(C_2 m)\big).$$

Proof. Let $g \in K$ be non-trivial. There exists $m_0 \in \mathbb{N}$ such that $g \notin \operatorname{Stab}_\Gamma(m_0)$. Then as $\mathbf{e}$ ranges over $\{0,1\}^{V_{m-m_0}}$, the elements $(g^{e_v})_{v \in V_{m-m_0}}$ are distinct modulo $\operatorname{Stab}_\Gamma(m)$. Thus:

$$|\Gamma : \operatorname{Stab}_\Gamma(m)| \geq 2^{|V_{m-m_0}|} = \exp\big(\log(2)|V_{m_0}|^{-1} \exp(|\mathcal{A}|m)\big).$$

□

The next Lemma will allow us to use the branch structure of Γ to help construct our approximations by power words.

Lemma 4.7. *Let $N_1, N_2, M \lhd \Gamma$ with $N_2 \leq N_1 \leq M \leq K$. Let $A, k \in \mathbb{N}$. Suppose that for all $z \in N_1$, there exist $y_1, \ldots, y_A \in M$ such that:*

$$y_1^k \cdots y_A^k z^{-1} \in N_2. \tag{11}$$

Then for all $m \in \mathbb{N}$ and all $\mathbf{z} \in N_1^{V_m}$, there exist $\mathbf{y}_1, \ldots, \mathbf{y}_A \in M^{V_m}$ such that:

$$\mathbf{y}_1^k \cdots \mathbf{y}_A^k \mathbf{z}^{-1} \in N_2^{V_m}. \tag{12}$$

Proof. Write $\mathbf{z} = (z_v)_{v \in V_m}$, for $z_v \in N_1$. For each z_v, choose corresponding $y_{v,1}, \ldots, y_{v,A} \in M$ as in (11). Then $\mathbf{y}_i = (y_{v,i})_{v \in V_m} \in M^{V_m}$ satisfy (12). □

Theorem 4.8. *Let p be a prime number. Let $\Gamma \leq \mathrm{Aut}(\mathcal{T}_{\mathcal{A}})$ be regular branch over $K \lhd \Gamma$. Suppose that:*

(i) There exists $a \in \mathbb{N}$ such that $|K : K^{\mathcal{A}}| = p^a$;

(ii) There exists $n_0 \in \mathbb{N}$ such that $K^{V_{n_0}} \leq \mathfrak{V}_p(K)$.

Then for all $n \in \mathbb{N}$,

$$\operatorname{diam}^+\big(\Gamma/K^{V_n}\big) = O_\Gamma\big(\exp(Cn)\big) \tag{13}$$

where $C = an_0 \log(an_0 p + 1)$.

Proof. Define normal subgroups H_i of K by $H_1 = K^{V_{n_0}}$; $H_{i+1} = [K, H_i]\mathfrak{V}_p(H_i)K^{V_{n_0+1}}$. Note that for each i, H_i is contained in the ith term of the lower central p-series of $K/K^{V_{n_0+1}}$. Since $K/K^{V_{n_0+1}}$ is a finite p-group, there exists i such that $H_i = K^{V_{n_0+1}}$. Let i_0 be the minimal index i with this last property. Note that if i is such that $H_{i+1} = H_i$, then $H_j = H_i$ for all $j \geq i$. Thus for $i < i_0$, p divides $|H_i : H_{i+1}|$. Now $|K^{V_{n_0}}/K^{V_{n_0+1}}| = p^{an_0}$, so $H_{an_0+1} = K^{V_{n_0+1}}$.

By construction, for each $1 \leq i \leq an_0$ H_i/H_{i+1} is an elementary abelian p-group of rank at most an_0. By hypothesis (ii), H_i is generated by pth powers of elements of K, so for each i there exist $y_1, \ldots, y_{an_0} \in K$ such that $y_1^p, \ldots, y_{an_0}^p \in H_i$ span H_i/H_{i+1}. Thus for all $z \in H_i$, there exist $\alpha_i \in \{0, 1, \ldots, p-1\}$ such that:

$$(y_1^{\alpha_1})^p, \ldots, (y_{an_0}^{\alpha_{an_0}})^p z^{-1} \in H_{i+1}. \tag{14}$$

We will apply Theorem 2.3. For $n \geq 1$, let $q, r \in \mathbb{N}$ be such that $n = an_0 q + r$, with $1 \leq r \leq an_0$. Set $M_n = K^{V_q}$ and $N_n = H_r^{V_q}$, so that for $m \geq n_0$, $K^{V_m} = N_{an_0(m-n_0)+1}$. Hypotheses (i), (ii) and (iii) of Theorem 2.3 are now immediate, and hypothesis (iv) follows by applying Lemma 4.7 to (14), with $A_n = an_0$ and $k_n = p$ for all n. The bound (13) then follows from (6). □

Remark 4.9. In the above proof, an approximation to an element of N_n, modulo N_{n+1}, as a product of at most an_0 pth powers of elements of M_n, may be computed in time $O_\Gamma\big(|\mathcal{A}|^{n/an_0}\big)$. For this, we encode elements of Γ/K^{V_n} via a *branch portrait*. That is, let $1 \in T$ be a transversal to $K^{\mathcal{A}}$ in Γ and let $1 \in U$ be a transversal to $K^{\mathcal{A}}$ in K. Given T and U, then for all $g \in \Gamma/K^{V_n}$ there exist unique data $\big(t, (u_v)_{1 \leq m \leq n-1; v \in V_m}\big)$ with $t \in T$, $u_v \in U$ such that:

$$g = t \cdot (u_{v_1})_{v_1 \in V_1} \cdots (u_{v_{n-1}})_{v_{n-1} \in V_{n-1}}.$$

Write $n = an_0 q + r$, with $1 \leq r \leq an_0$, so that $M_n = K^{V_q}$, $N_n = H_r^{V_q}$ and $N_{n+1} = H_{r+1}^{V_q}$. $K^{V_{n_0+1}} \leq H_r \leq K^{V_{n_0}}$, so for $g \in \Gamma$, $g \in N_n$ iff $u_v = 1$ for $v \in V_m$ and $1 \leq m \leq q + n_0 - 1$ and for all $v \in V_q$, $(u_{vw})_{w \in V_{n_0}} \in H_r$. This can be verified in time linear in $|V_q|$. For fixed $v \in V_q$, writing $z_v = (u_{vw})_{w \in V_{n_0}} \in H_r$ as a product of an_0 pth powers in K modulo H_{r+1} as in (14) is accomplished in bounded time, so as in Lemma 4.7, writing $(z_v)_{v \in V_q}$ as a product of an_0 pth powers in M_n modulo N_{n+1} takes times at most linear in $|V_q|$.

Hence if $\tilde{f}$ is a function such that elements of Γ/K^{V_m}, described by their branch portraits, may be multiplied and inverted in time at most $\tilde{f}(m)$, then the directed navigation problem for Γ/K^{V_m} is solvable for the bound (13) in time $O_\Gamma\big((\tilde{f}(m) + |\mathcal{A}|^m) \cdot |S|^{|\Gamma : K^{V_{n_0}}|+1} \cdot (an_0 + 1)^{an_0 m}\big)$.

Proof of Theorem 1.3. We have $\mathcal{T} = \mathcal{T}_\mathcal{A}$, with $|\mathcal{A}| = p$. We will apply Theorem 4.8. Since $K/K^\mathcal{A}$ is a finite p-group, hypothesis (i) holds.

For hypothesis (ii), note that $\mho_p(K)$ is a normal subgroup of Γ. By Proposition 4.5, Γ does not have exponent dividing $p|\Gamma : K|$. Thus $\mho_p(K)$ is non-trivial, hence of finite-index (since Γ is just-infinite). By the congruence subgroup property, there exists n_0 such that $\mathrm{Stab}_\Gamma(n_0) \leq \mho_p(K)$. Thus (13) holds.

Since $K^{V_n} \leq \mathrm{Stab}_\Gamma(n)$ we have $\mathrm{diam}^+\big(\Gamma/\mathrm{Stab}_\Gamma(n)\big) = O_\Gamma\big(\exp(Cn)\big)$. Finally, by Lemma 4.6, $\exp(Cn) = O\big(\log^{O(1)}|\Gamma : \mathrm{Stab}_\Gamma(n)|\big)$. □

Remark 4.10. *The constants appearing in Theorem 1.3 depend on p; on a and n_0 from Theorem 4.8, and on the constants appearing in Lemma 4.6.*

Example 4.11. Let p be an odd prime and set $\mathcal{A} = \{0, 1, \ldots, p-1\}$. Let $\Gamma \leq \mathrm{Aut}(\mathcal{T}_\mathcal{A})$ be the *GGS group* with defining vector $\mathbf{e} \in \mathbb{F}_p^{p-1}$. It is shown in [22] (Theorem 2.7) that if $\mathbf{e}$ is non-constant, then Γ is just-infinite and has the congruence subgroup property. It further follows from Propositions 2.2 and 2.3 of [22] that Γ branches over $K = \gamma_3(\Gamma) \leq \mathrm{Stab}_\Gamma(1)$, of index p^3 in Γ. Thus $K/K^\mathcal{A}$, which is naturally embeddable into $(\Gamma/K)^\mathcal{A}$, is a p-group, and Theorem 1.3 applies. The Gupta-Sidki p-groups studied in [11] correspond to the case $\mathbf{e} = (1, -1, 0 \ldots, 0)$. For these groups, a polylogarithmic bound for $\mathrm{diam}^+\big(\Gamma/\mathrm{Stab}_\Gamma(n)\big)$ follows from the results of [11] and Theorem 1.5 (albeit without any non-trivial solution to the directed navigation problem). For all other GGS groups considered the diameter bound coming from Theorem 1.3 is new, even for undirected diameters. The Fabrykowski-Gupta group of Theorem 1.4 is also a GGS group (with $p = 3$, $\mathbf{e} = (1, 0)$). Nevertheless we see in the next Section that quantitatively stronger bounds follow by applying Theorem 4.8 rather than going via Theorem 1.3.

5 The Fabrykowski-Gupta Group

In this Section we deduce Theorem 1.4 from Theorem 4.8. First let us define the Fabrykowski-Gupta group Γ. Let $\mathcal{A} = \{0, 1, 2\}$ and write $\mathcal{T}_\mathcal{A} = \mathcal{T}_3$. The *Fabrykowski-Gupta* group is the subgroup Γ of $\mathrm{Aut}(\mathcal{T}_3)$ which is generated by the two automorphisms a, b defined by:

$$\begin{array}{lll} a(0w) = 1w, & a(1w) = 2w, & a(2w) = 0w, \\ b(0w) = 0(aw), & b(1w) = 1w, & b(2w) = 2(bw). \end{array} \tag{15}$$

That is, a cyclically permutes the subtrees rooted at 0, 1 and 2, while $b \in \mathrm{Stab}_\Gamma(1)$ is defined recursively by $b = (a, 1, b)$. It is easily seen that a and b have order 3. Using (15) we compute the following, which will be useful subsequently.

Lemma 5.1. *Let $i, j, k \in \{0, 1, 2\}$. Then for $w \in \{0, 1, 2\}^*$,*

(i) $a^i b^j (b^a)^k 000w = ijkw$;

(ii) $(b^{ab})^i 000w = 000(a^i w)$;

(iii) $(b^{aba^{-1}ba})^i 000w = 000(b^i w)$.

Let $K = [\Gamma, \Gamma]$ be the derived subgroup of Γ. We have $K \leq \mathrm{Stab}_\Gamma(1)$, since $\Gamma / \mathrm{Stab}_\Gamma(1) \cong C_3$ is abelian. Consider the following elements of K:

$$x_1 = [a, b] = (b^{-1}a, a^{-1}, b)$$
$$x_2 = [a, x_1] = (ba, a^{-1}ba^{-1}, ab).$$

Proposition 5.2. *Let $\Gamma, K, a, b, x_1, x_2$ be as above.*

(i) $K = \langle x_1 \rangle^\Gamma$;

(ii) Γ *branches over* K;

(iii) $\Gamma/K \cong C_3 \times C_3$, *with basis* Ka, Kb;

(iv) $K/K^{\mathcal{A}} \cong C_3 \times C_3$, *with basis* $K^{\mathcal{A}}x_1, K^{\mathcal{A}}x_2$.

Proof. (i) is clear since Γ is generated by a and b. (ii) is proved as Proposition 6.2 in [6]. (iii) and (iv) also follow easily from the results of [6] Section 6, however for the sake of completeness we give a self-contained proof.

For (iii), Γ/K is certainly a quotient of $C_3 \times C_3$, since Γ is generated by two elements of order 3. On the other hand, there is a natural homomorphism $\Gamma \to C_3 \wr C_3$ (with kernel $\mathrm{Stab}_\Gamma(2)$). Inspection of the action of a and b on $\mathcal{T}_3$ confirms that this homomorphism is surjective. But $(C_3 \wr C_3)^{\mathrm{Ab}} \cong C_3 \times C_3$.

For (iv), note that by embedding $K \leq \mathrm{Stab}_\Gamma(1) \hookrightarrow \Gamma^{\mathcal{A}}$, $K/K^{\mathcal{A}}$ is naturally a subgroup of $(\Gamma/K)^{\mathcal{A}}$, so by (ii) is an elementary abelian 3-group. Moreover by (i), K is generated by the conjugates of x_1. Consider the action of Γ on $K/K^{\mathcal{A}}$ by conjugation. Since b acts trivially, $K/K^{\mathcal{A}}$ is generated by the images of x_1, x_1^a and $x_1^{a^2}$.

Now x_1, x_1^a are non-zero and independent modulo $K^{\mathcal{A}}$ (x_1^a has non-zero a-component in the 3rd co-ordinate, which x_1 does not, for instance). However $x_1^{a^2} \equiv (x_1^a x_1)^{-1} \mod K^{\mathcal{A}}$. Hence $K/K^{\mathcal{A}} \cong C_3 \times C_3$ is spanned by x_1 and x_1^a, or equivalently by x_1 and $(x_1^a)^{-1}x_1 = x_2$. □

Thus we have a descending sequence of finite-index normal subgroups:

$$\Gamma \geq K \geq K^{\mathcal{A}} \geq K^{V_2} \geq \ldots \geq K^{V_m} \geq \ldots$$

with $|\Gamma : K^{V_m}| = 3^{3^m+1}$. Moreover $K^{V_m} \leq \mathrm{Stab}_\Gamma(m+1)$ for all $m \in \mathbb{N}$. Since, by [6] Proposition 6.5:

$$|\Gamma : \mathrm{Stab}_\Gamma(m+1)| = 3^{3^m+1} \tag{16}$$

we conclude the following.

Corollary 5.3. $K^{V_m} = \mathrm{Stab}_\Gamma(m+1)$ *for all* $m \geq 1$.

The key to the proof of Theorem 1.4 is the next result, which provides an effective version of hypothesis (ii) of Theorem 4.8.

Proposition 5.4. $K^{V_3} \leq \mho_3(K)$.

Proof of Theorem 1.4. By Corollary 5.3, $\Gamma/\operatorname{Stab}_\Gamma(m) = \Gamma/K^{V_{m-1}}$. We apply Theorem 4.8 with $p = 3$. By Proposition 5.2 (iv) and Proposition 5.4 we may take $a = 2$ and $n_0 = 3$, respectively. Therefore, for C as in the statement of Theorem 1.4:

$$\operatorname{diam}^+\big(\Gamma/\operatorname{Stab}_\Gamma(m)\big) = O(19^{6m}) = O(3^{Cm}) = O\big(\log^C|\Gamma/\operatorname{Stab}_\Gamma(m)|\big)$$

where the last equality holds by (16).

For the directed navigation problem we observe that we may take $\tilde{f}(m) = O(3^m)$ in Remark 4.9. By Corollary 5.3, the description of elements of $\Gamma/\operatorname{Stab}_\Gamma(m) = \Gamma/K^{V_{m-1}}$ via a branch portrait is clearly equivalent to a description as permutations of V_m, and passing between these descriptions takes time at most linear in $|V_m|$. Composition and inversion of elements of $\operatorname{Sym}(V_m)$ may be accomplished in time linear in $|V_m| = 3^m$. Thus the runtime is:

$$O\big(|S|^{O(1)}3^m(7^6)^m\big) = O\big(|S|^{O(1)}\log^{C'}|\Gamma/\operatorname{Stab}_\Gamma(m)|\big).$$

Finally, we reduce the dependence of the runtime of our solution to the directed navigation problem on the order d of a generating set S. $G_m = \Gamma/\operatorname{Stab}_\Gamma(m)$ is a finite 3-group, and is 2-generated (since Γ is). Thus S contains a 2-element subset T which still generates G_m; indeed any subset which generates G_m modulo $\Phi(G_m) = [G_m, G_m]\mho_p(G_m) = K/\operatorname{Stab}_\Gamma(m)$ has this property. Computing the image of each $s \in S$ in $G_m/\Phi(G_m) \cong C_p^2$ in turn, we find a generating pair T in time linear in d, and having reduced from S to T, the runtime is: $O\big(\log^{C'}|\Gamma/\operatorname{Stab}_\Gamma(n)|\big)$, as desired. □

It therefore suffices to prove Proposition 5.4. At this point let us introduce some further notation. For $v \in V_n$ and $g \in \operatorname{Aut}(\mathcal{T}_\mathcal{A})$, let $p_v(g) \in \operatorname{Stab}(n)$ be the unique tree automorphism satisfying $p_v(g)_v = g$ and $p_v(g)_u = \mathrm{id}$ for all $u \in V_n \setminus \{v\}$. Note that for all $k \in K$ and $v \in \mathcal{A}^*$, $p_v(k) \in K$.

Lemma 5.5. $p_{000}(x_1) \in \mho_3(K)$.

Proof. We claim that:

$$p_{000}(x_1) = \big(x_1^{ba}p_0(x_1)^{b^a}\big)^3\big(x_1^{ba}\big)^{-3}\big(p_0(x_1)^{b^a}\big)^{-3}$$

(a product of three cubes in K). Recall that $b = (a, 1, b)$ and $x_1 = (b^{-1}a, a^{-1}, b)$, so $b^a = (b, a, 1)$ and:

$$x_1^{ba} = (a^{-1}b^{-1}a^{-1}, a^{-1}, b)^a = (b, a^{-1}b^{-1}a^{-1}, a^{-1})$$

$$p_0(x_1)^{b^a} = p_0(x_1^b)$$

$$x_1^{ba}p_0(x_1)^{b^a} = (x_1b, a^{-1}b^{-1}a^{-1}, a^{-1})$$

Thus:

$$(x_1^{ba})^{-3} = (1, (aba)^3, 1)$$

$$(x_1^{ba}p_0(x_1)^{b^a})^3 = ((x_1b)^3, (aba)^{-3}, 1)$$

so:

$$\big(x_1^{ba}p_0(x_1)^{b^a}\big)^3\big(x_1^{ba}\big)^{-3}\big(p_0(x_1)^{b^a}\big)^{-3} = p_0\big((x_1b)^3(x_1^b)^{-3}\big) \tag{17}$$

and:

$$x_1b = (b^{-1}a^{-1}, a^{-1}, b^{-1})$$
$$x_1^b = (a^{-1}b^{-1}a^{-1}, a^{-1}, b)$$

hence:

$$(x_1b)^3(x_1^b)^{-3} = p_0\big((b^{-1}a^{-1})^3(aba)^3\big) \tag{18}$$

while:

$$(b^{-1}a^{-1})^3(aba)^3 = b^{-1}(b^{-1})^ab(b^a) = p_0(x_1). \tag{19}$$

Combining (17), (18) and (19), we have the required conclusion. □

Lemma 5.6. *For all $y \in K$, $p_{000}(y) \in \mho_3(K)$.*

Proof. Let $L = \{y \in K : p_{000}(y) \in \mho_3(K)\}$. It is clear that L is closed under composition and inversion. By Lemma 5.5, $x_1 \in L$. By Lemma 5.1 (ii) and (iii), L is invariant under conjugation by a and b. Hence $L \supseteq \langle x_1\rangle^\Gamma = K$ (by Proposition 5.2 (i)). □

Proof of Proposition 5.4. Let $\mathbf{y} \in K^{V_3}$. Then:

$$\mathbf{y} = \prod_{v\in V_3} p_v(y_v)$$

It therefore suffices to check that for $y \in K$ and $v \in V_3$, $p_v(y) \in \mho_3(K)$. Let $i, j, k \in \{0, 1, 2\}$ be such that $v = ijk$. Then by Lemma 5.1 (i), $p_v(y) = p_{000}(y)^g$, where $g = (b^a)^k b^j a^i$. We are done by Lemma 5.6. □

6 p-adic Analytic Groups

In this Section we prove a directed diameter bound for a sequence of quotients of an arbitrary compact p-adic group, and observe that our bound is an instance of the potent SKP. We assume that this bound is well-known, but we are not aware of an existing reference. Before stating the result we require some background on p-adic analytic groups. Our exposition here is based on [19].

Definition 6.1. *Let Γ be a finitely generated pro-p group.* Γ *is* powerful *if $\Gamma/\overline{\mho_{p^e}(\Gamma)}$ is abelian, where $e = 2$ when $p = 2$ and $e = 1$ when p is odd.* Γ *is* uniform *if it is powerful and torsion-free. The* rank *of a uniform group is the minimal size of a topological generating set.*

There are many characterizations of p-adic analytic groups. For compact groups, perhaps the easiest to visualize is this: a compact topological group Γ is p-adic analytic iff it is isomorphic to a closed subgroup of some $\mathrm{SL}_n(\hat{\mathbb{Z}}_p)$. Equivalently, Γ is p-adic analytic iff it has the structure of a p-adic analytic manifold, such that the group operations are analytic functions. The *dimension* of Γ in this case is its dimension as a p-adic analytic manifold.

Theorem 6.2. *Let Γ be a compact p-adic analytic group. Then Γ has an open characteristic uniform pro-p subgroup K.*

Proof. By [19] Corollary 8.33, Γ has an open subgroup H which is pro-p of *finite rank* (see [19] Definition 3.12). The characteristic core J or H in Γ is also open pro-p of finite rank. By [19] Corollary 4.3, J has an open characteristic uniform pro-p subgroup K. In particular K is characteristic in Γ. □

Proposition 6.3. *Let Γ be a pro-p group. Let $(\Gamma_i)_i$ be the lower central p-series of Γ. Then for all i, j:*

(i) $[\Gamma_i, \Gamma_j] \leq \Gamma_{i+j}$;

(ii) $\mho_p(\Gamma_i) \leq \Gamma_{i+1}$.

Proof. (i) is proved as [19] Proposition 1.16 (ii). (ii) is immediate from the definition of the lower central p-series. □

Theorem 6.4 ([19] Theorem 3.6 (ii) and (iv))**.** *Let H be a finitely generated powerful pro-p group. Let $(H_i)_i$ be the lower central p-series of H.*

(i) For all i, j, $(H_{i+1})_{j+1} = H_{i+j+1}$;

(ii) For all i, j, the map $x \mapsto x^{p^j}$ induces an epimorphism $H_i/H_{i+1} \to H_{i+j}/H_{i+j+1}$.

Theorem 6.5 ([19] Theorem 8.36)**.** *Let Γ be a compact p-adic analytic group of dimension d. Let K be an open uniform subgroup of Γ. Then K has rank d.*

Lemma 6.6. *Let K be a uniform pro-p group of rank d. Let $(K_i)_i$ be the lower central p-series of K. Then $K_i/K_{i+1} \cong C_p^d$.*

Proof. By Theorem 6.4 (i), $K_{i+1} = \Phi(K_i)$, so $K_i/K_{i+1} \cong C_p^{d(K_i)}$. By [19] Proposition 4.4, $d(K_i) = d$. □

We are now ready to state and prove our diameter bound.

Theorem 6.7. *Let Γ be a compact p-adic analytic group of dimension d. Let K be an open characteristic uniform pro-p subgroup. Let $(K_i)_i$ be the lower central p-series of K. Then for all n,*

$$\operatorname{diam}^+(\Gamma/K_n) \leq |\Gamma : K_2|(p^{n-1} - 1)/(p - 1) = O_\Gamma\big(|\Gamma : K_n|^{1/d}\big).$$

Proof. Let $S \subseteq \Gamma/K_n$ be a generating set. Then $K_1/K_2 \subseteq B_S^+(|\Gamma : K_2|)K_2/K_2$, so by Theorem 6.4 (iii), $K_i/K_{i+1} \subseteq B_S^+(p^{i-1}|\Gamma : K_2|)$ for all $i \leq n - 1$,

$$\operatorname{diam}^+(\Gamma/K_n, S) \leq |\Gamma : K_2|(1 + p + \cdots + p^{n-2}).$$

We may also interpret this bound as an instance of the of the potent SKP. We apply Theorem 2.3 with $M_i = K_i$, $N_i = K_{i+1}$, $A_i = 1$ and $k_i = 1$. Hypothesis (i) of Theorem 2.3 is clear; hypotheses (ii) and (iii) follow from Theorem 6.4 (i), and hypothesis (iv) follows from Theorem 6.4 (iii). Moreover $N_i = M_{i+1}$ so the

improvement described in Remark 2.4 (ii) is available to us, with $n_0 = 1$, and the required bound follows from (8).

For the second equality, it suffices to note that $|K_i : K_{i+1}| \geq p^d$ for all i, which is immediate from Lemma 6.6. □

Remark 6.8. The conclusion of Theorem 6.7 is best possible in general: this is witnessed by the example $\Gamma = F \times \hat{\mathbb{Z}}_p^d$, where F is a finite group (which may be chosen to be of arbitrarily large diameter). Under the assumption that Γ is *Fab* (that is: every open subgroup has finite abelianisation) much stronger, indeed polylogarithmic, diameter bounds for Γ/K_n are provided by [10]. These may then be extended to the directed diameter by Theorem 1.5. Nevertheless, the degree of the polylogarithmic upper bound for $\mathrm{diam}(\Gamma/K_n)$ from [10] in general grows like $\log(d)$ in the dimension d of Γ, so the conclusion of Theorem 6.7 does improve upon the results of [10] for certain groups Γ/K_n when d is large compared with n and p (say $\log(d) \gg \log(p)n/\log(n)$).

7 Spectral Gap and Mixing Time

Let G be a finite group and $S \subseteq G$. Let A_S be the (symmetric, normalized) adjacency operator on the Cayley graph $\mathrm{Cay}(G, S)$. A_S is a self-adjoint operator of norm one; let its spectrum be:

$$1 = \lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_{|G|} \geq -1.$$

The eigenvalue λ_1 corresponds to the constant functions on G; it is a simple eigenvalue iff S generates G. In this case, the quantity $1 - \lambda_2$ is the *spectral gap* of the pair (G, S).

In many applications it is desirable for a Cayley graph to have large spectral gap. In particular, a family of bounded-valence Cayley graphs whose spectral gaps are uniformly bounded away from zero form an *expander family*. There is also a close relationship between spectral gap and diameter.

Proposition 7.1 ([15] Corollary 3.1)**.** *The spectral gap of (G, S) is at least $(2|S|\,\mathrm{diam}(G, S)^2)^{-1}$.*

From this inequality and our diameter bounds, we obtain substantial lower bounds on spectral gap for Cayley graphs of our groups (albeit weaker bounds than would be needed to verify that our Cayley graphs are expanders).

A second invariant of great interest in both practical and theoretical contexts is the *mixing time* of the pair (G, S), which measures the time taken for a (symmetric) lazy random walk on $\mathrm{Cay}(G, S)$ to closely approximate the uniform distribution (with respect to some metric). Here we follow the following convention: let δ_e be the Dirac mass at the identity of G, and let $T_S = (A_S + I)/2$, where I is the identity operator on G.

Definition 7.2. *The* ℓ^∞-mixing time *of the pair (G, S) is the smallest $l \in \mathbb{N}$ such that:*

$$\left\|T_S^l \delta_e - \tfrac{1}{|G|}\chi_G\right\|_\infty \leq \tfrac{1}{2|G|}.$$

It is clear that the ℓ^∞-mixing time of (G,S) is an upper bound for the diameter. Via the spectral gap, we also have a converse inequality.

Proposition 7.3 ([29] Theorem 5.1)**.** *Suppose the pair (G,S) has spectral gap $\epsilon > 0$. Then there exists an absolute constant $C > 0$ such that the ℓ^∞-mixing time of (G,S) is at most $(C/\epsilon)\log|G|$.*

Using our diameter bounds, we therefore also obtain new upper bounds on ℓ^∞-mixing time.

Corollary 7.4. *Let q be a power of 2 and let $S_n \subseteq G_n = \mathrm{SL}_2(\mathbb{F}_q[t]/(t^n))$ be a generating set. Then for all $\epsilon > 0$ the spectral gap of (G_n, S_n) is $\Omega_{q,\epsilon}\big(|S_n|^{-1}\log^{-C-\epsilon}|G_n|\big)$ and the ℓ^∞-mixing time of (G_n,S_n) is $O_{q,\epsilon}\big(|S_n|\log^{1+C+\epsilon}|G_n|\big)$, where $C = 2\log(7)/\log(4/3) \approx 13.528$.*

Corollary 7.5. *Let Γ be the Fabrykowski-Gupta group and let $S_n \subseteq G_n = \Gamma/\operatorname{Stab}_\Gamma(n)$ be a generating set. Then the spectral gap of (G_n,S_n) is $\Omega\big(|S_n|^{-1}\log^{-C}|G_n|\big)$ and the ℓ^∞-mixing time of (G_n,S_n) is $O\big(|S_n|\log^{1+C}|G_n|\big)$, where $C = 12\log(19)/\log(3) \approx 32.162$.*

Corollary 7.6. *Let Γ be a compact p-adic analytic group of dimension d; let K be an open characteristic powerful pro-p subgroup; let $(K_i)_i$ be the lower central p-series of K, and let $S_n \subseteq G_n = \Gamma/K_n$ be a generating set. Then the spectral gap of (G_n,S_n) is $\Omega_\Gamma\big(|S_n|^{-1}|G_n|^{-2/d}\big)$ and the ℓ^∞-mixing time of (G_n,S_n) is $O_\Gamma\big(|S_n||G_n|^{2/d}\log|G_n|\big)$.*

8 Lower Bounds

Proving non-trivial lower bounds on diameters of Cayley graphs is in general a difficult problem. An elementary counting argument shows that, if a finite group G is generated by a subset S with $|S| \geq 2$, then $\mathrm{diam}^+(G,S) \geq \log|G|/\log|S| - 1$. In particular, if (G_n) is a sequence of finite groups which admit generating sets of bounded size, then:

$$\log|G_n| = O\big(\mathrm{diam}(G_n)\big). \tag{20}$$

A stronger lower bound is available for finite groups arising as quotients of an infinite finitely generated group of *subexponential growth*. Recall that for Γ a group with a finite generating set S, the *growth function* of Γ is $\gamma_\Gamma^S(n) = |B_S(n)|$. If $\pi : \Gamma \to G$ is a finite quotient of Γ, then $\pi(B_S(n)) = B_{\pi(S)}(n)$, from which we have the next bound.

Proposition 8.1. *Let Γ be an infinite group with a finite generating set S, let (G_n) be a sequence of finite quotients of Γ and let $f : \mathbb{N} \to \mathbb{N}$ be a monotone increasing function with $\gamma_\Gamma^S \leq f$. Then for all n, $\mathrm{diam}(G_n) \geq f^{-1}\big(|G_n|\big)$.*

In [11] we used Proposition 8.1 to relate the growth of the first Grigorchuk group to the diameters of its finite quotients. In the case of the Fabrykowski-Gupta group, we can improve slightly on (20), using the work of Bartholdi and Pochon.

Theorem 8.2 ([7] Theorem 1)**.** *Let Γ be as in Section 5. Then for any finite generating set S of Γ, $\gamma_\Gamma^S(n) = \exp\big(O(n(\log\log n)^2/\log n)\big)$.*

Corollary 8.3. *There exists an absolute constant $C > 0$ such that for any finite quotient G of Γ, $\mathrm{diam}(G) \geq C \log|G| \log\log|G|/(\log\log\log|G|)^2$.*

Acknowledgements

This work was supported by ERC grant no. 648329 "GRANT". I am grateful to Alejandra Garrido for helpful discussions. I would also like to thank two anonymous referees, whose insightful commentary on a previous version of this paper helped me both to strengthen the results contained herein and to improve their presentation.

References

[1] M. Abért. *Group laws and free subgroups in topological groups.* Bull. London Math. Soc. 37 (2005), Issue 4, 525–534.

[2] L. Babai. *Local expansion of vertex-transitive graphs and random generation in finite groups.* In Proceedings of the twenty-third annual ACM symposium on theory of computing, 164–174, ACM, New York, 1991.

[3] L. Babai. *On the diameter of Eulerian orientations of graphs.* In Proceedings of the seventeenth annual ACM-SIAM symposium on discrete algorithms, 822–831, ACM, New York, 2006.

[4] L. Babai, R. Beals, Á. Seress. *On the diameter of the symmetric group: polynomial bounds.* In Proceedings of the fifteenth annual ACM-SIAM symposium on discrete algorithms, 1108–1112, SIAM, Philadelphia, 2004.

[5] L. Babai, T.P. Hayes. *Near-independence of permutations and an almost sure polynomial bound on the diameter of the symmetric group.* Proceedings of the sixteenth annual ACM-SIAM symposium on discrete algorithms, 1057–1066, SIAM, Philadelphia, 2005.

[6] L. Bartholdi, R. Grigorchuk. *On parabolic subgroups and Hecke algebras of some fractal groups.* Serdica Math. J. 28, No. 1 (2002) 47–90

[7] L. Bartholdi, F. Pochon. *On growth and torsion of groups.* Groups Geom. Dyn. 3 (2009), 525–539

[8] J. Bourgain, A. Gamburd. *Expansion and random walks in $SL_d(\mathbb{Z}/p^n\mathbb{Z})$: I.* J. Eur. Math. Soc. 10, Issue 4 (2008) 987–1011.

[9] J. Bourgain, A. Gamburd. *Expansion and random walks in $SL_d(\mathbb{Z}/p^n\mathbb{Z})$: II.* J. Eur. Math. Soc. 11, Issue 5 (2009) 1057–1103.

[10] H. Bradford. *New Uniform Diameter Bounds in Pro-p Groups.* Groups Geom. Dyn. 12 (2018), Issue 3, 803–836

[11] H. Bradford. *Uniform Diameter Bounds in Branch Groups.* arXiv:1703.05852 [math.GR]

[12] L. Bromberg, V. Shpilrain, A. Vdovina. *Navigating in the Cayley graph of $\mathrm{SL}_2(\mathbb{F}_p)$ and applications to hashing.* Semigroup Forum 94, Issue 2, (2017) 314–324.

[13] C.M. Dawson, M.A. Nielsen. *The Solovay-Kitaev algorithm.* Quantum Information & Computation 6, Issue 1 (2006) 81–95.

[14] P. Diaconis. *Mathematical Developments from the Analysis of Riffle Shuffling.* In Groups, Combinatorics And Geometry: Durham 2001 (2003) 73–97.

[15] P. Diaconis, L. Saloff-Coste. *Comparison techniques for random walk on finite groups.* Ann. Probab. 21, Issue 4 (1993) 2131–2156.

[16] O. Dinai. *Uniform poly-log diameter bounds for some families of finite groups.* Proc. Amer. Math. Soc. 134, Issue 11 (2006) 3137–3142.

[17] O. Dinai. *Growth in SL_2 over finite fields.* Journal of Group Theory 14, Issue 2 (2011) 273–297.

[18] O. Dinai. *Diameters of Chevalley groups over local rings.* Archiv der Mathematik 99, Issue 5 (2012) 417–424.

[19] J.D. Dixon, M.P.F. Du Sautoy, A. Mann, D. Segal. *Analytic pro-p groups (2nd Edition).* Cambridge Studies in Advanced Mathematics. Cambridge University Press (1999).

[20] D. Dolgopyat. *On mixing properties of compact group extensions of hyperbolic systems.* Israel J. Math. 130, Issue 1 (2000) 157–205.

[21] J. Fabrykowski, N.D. Gupta. *On groups with sub-exponential growth functions. II.* J. Indian Math. Soc. (N.S.) 56 (1991), 217–228.

[22] G.A. Fernández-Alcober, A. Garrido, J. Uria-Albizuri. *On the congruence subgroup property for GGS-groups.* Proc. Amer. Math. Soc. 145 (2017) 3311–3322.

[23] A. Gamburd, M. Shahshahani. *Uniform diameter bounds for some families of Cayley graphs.* Int. Math. Res. Notices 71 (2004), 3813–3824.

[24] M.C. Heydemann. *Cayley graphs and interconnection networks.* In Graph Symmetry, 167–224. Springer Netherlands, 1997.

[25] S. Hoory, N. Linial, A. Wigderson. *Expander Graphs and their Applications.* Bull. Amer. Math. Soc. 43, Issue 4 (2006) 439–561.

[26] M. Kassabov, T.R. Riley. *Diameters of Cayley graphs of Chevalley groups.* European Journal of Combinatorics 28, Issue 3 (2007), 791–800.

[27] D. Kornhauser, G. Miller, P. Spirakis. *Coordinating pebble motion on graphs, the diameter of permutation groups, and applications.* MS thesis, M.I.T., Dept. of Electrical Engineering and Computer Science, 1984.

[28] M. Larsen. *Navigating the Cayley graph of* $\mathrm{SL}_2(\mathbb{F}_p)$. Int. Math. Res. Not., Volume 2003, Issue 27 (2003), 1465–1471.

[29] L. Lovász. *Random walks on graphs: a survey.* Combinatorics, Paul Erdös is Eighty 2 (1993), 146.

[30] A. Lubotzky. *Expander Graphs in Pure and Applied Mathematics.* Bull. Amer. Math. Soc. 49, Issue 1 (2012) 113–162.

[31] C. Petit, J.J. Quisquater. *Rubik's for Cryptographers.* IACR Cryptology ePrint Archive 638 (2011).

[32] T.R. Riley. *Navigating in the Cayley Graphs of* $\mathrm{SL}_N(\mathbb{Z})$ *and* $\mathrm{SL}_N(\mathbb{F}_p)$. Geom. Dedicata 113, Issue 1, (2005) 215–229.

[33] J.C. Schalge-Puchta. *Applications of character estimates to statistical problems for the symmetric group.* Combinatorica 32, Issue 3 (2012) 309–323.

H. Bradford, Georg-August-Universität Göttingen, Germany
E-mail address: henry.bradford@mathematik.uni-goettingen.de