PROD. TYPE: COM
PP: 1-16 (col.fig.: Nil)
JPAA3227 DTD VER: 5.0.1
ED: Sabitha
PAGN: Murthy.N – SCAN: Hema

ARTICLE IN PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

ELSEVIER

Journal of Pure and Applied Algebra ▮▮▮ (▮▮▮▮) ▮▮▮–▮▮▮

JOURNAL OF PURE AND APPLIED ALGEBRA

www.elsevier.com/locate/jpaa

Bass–Tits minimization of automata, quotients of trees and diameters

Lisa Carbone[*,1], Dennis Clark[2]

Department of Mathematics, Rutgers, The State University, 110 Freling Huysen Road, Piscataway, NJ 08854-8019, USA

Received 21 July 2004; received in revised form 5 January 2005

Communicated by C.A. Weibel

Abstract

Let X be a tree and let $G=Aut(X)$, Bass and Tits have given an algorithm to construct the 'ultimate quotient' of X by G starting with any quotient of X, an 'edge-indexed' graph. Using a sequence of integers that we compute at consecutive steps of the Bass–Tits (BT) algorithm, we give a lower bound on the diameter of the ultimate quotient of a tree by its automorphism group. For a tree X with finite quotient, this gives a lower bound on the minimum number of generators of a uniform X-lattice whose quotient graph coincides with $G\backslash X$. This also gives a criterion to determine if the ultimate quotient of a tree is infinite. We construct an edge-indexed graph (A, i) for a deterministic finite state automaton and show that the BT algorithm for computing the ultimate quotient of (A, i) coincides with state minimizing algorithm for finite state automata. We obtain a lower bound on the minimum number of states of the minimized automaton. This gives a new proof that language for the word problem in a finitely generated group is regular if and only if the group is finite, and a new proof that the language of the membership problem for a subgroup is regular if and only if the subgroup has finite index.

MSC: primary 20F32; secondary 22F50

* Corresponding author.
E-mail address: carbonel@math.rutgers.edu (L. Carbone).

[1] The first author was supported in part by NSF Grant #DMS-0401107.

[2] The second author was supported in part by a research award for undergraduates at Harvard, and an NSF Graduate Research Fellowship at the University of Michigan.

0022-4049/$ - see front matter © 2005 Published by Elsevier B.V.
doi:10.1016/j.jpaa.2005.04.013

1. Introduction

In [3], Bass and Tits gave an algorithm for computing the *ultimate quotient* of a tree X by its full automorphism group starting with any quotient of the tree, an 'edge-indexed' graph. In this work we give a number of applications of the BT algorithm. We show that the BT algorithm can be used to obtain a lower bound on the diameter of the ultimate quotient of a tree (Section 5). Given an edge-indexed graph we obtain a sequence $(d_s)_{s\geqslant 0}$ of positive integers, where the term d_s of the sequence is determined at step s of the Bass–Tits algorithm. At each step s, the term d_s is a lower bound for the diameter of the ultimate quotient. The sequence $(d_s)_{s\geqslant 0}$ converges to the diameter if the ultimate quotient is finite. If the sequence diverges our results imply that the ultimate quotient is infinite. In many cases this can be determined in a finite number of steps even though the BT algorithm may not terminate.

For finite edge-indexed graphs equal to their ultimate quotients, we can compute their diameters. Our diameter computation is a slight generalization of a standard method which computes the largest height of a shortest paths tree. Here we replace the distance between two vertices in a graph by the distance between a vertex and an equivalence class arising from the BT algorithm.

As another application we show that the BT algorithm can be viewed as a generalization of the state minimizing algorithm for a finite state automaton.

More precisely, to each deterministic automaton we associate a directed edge-indexed graph and we show that the edge-indexed graph of the minimized automaton coincides (up to graph isomorphism preserving edge-indices) with the BT ultimate quotient of the edge-indexed graph of the automaton (Section 8). The correspondence between these algorithms is natural given that they can both be viewed as using refinements of an equivalence relation on edge-indexed graphs.

The sequence $(d_s)_{s\geqslant 0}$ when computed for the edge-indexed graph of a finite state automaton gives a lower bound on the minimum number of states of the automaton (Section 9).

When applied to the automaton of the Cayley graph of a group G with fixed finite generating set S, our results give a new proof of Anisimov and Seifert's theorem [1] that states that the word problem language for G is regular if and only if the edge-indexed graph of the automaton, and hence the group G, is finite. We also give a new proof that the language of the membership problem for subgroups is not regular (Theorem 7).

In Section 2 we outline the notions of graphs and edge-indexed graphs. In Section 3, we describe the BT algorithm. In Section 4 we examine the structure of the ultimate quotient by determining the images of paths in an edge-indexed graph under the quotient morphism. We also give a description of the fiber over a vertex in the ultimate quotient. In Section 5 we describe the sequence $(d_s)_{s\geqslant 0}$ of an edge-indexed graph which determines the diameter of the ultimate quotient.

The authors would like to thank Paul Schupp for pointing out that the BT algorithm should coincide with the state minimizing algorithm for a finite state automaton, and to Ilya Kapovich for helpful discussions.

2. Graphs and edge-indexed graphs

In this section and throughout the paper, by a *graph* we mean a *quasi-graph* in the sense [3]. That is, we allow for the presence of *self-inverse loops* in the sense of [3]. Some arguments presented here may require slight modification for the explicit presence of self-inverse loops. Such details are routine and are left to the reader.

Let A denote a graph, with vertices VA, oriented edges EA. We assume that all graphs are connected. A path in A is called *reduced* if it contains no backtracking. A *morphism* $\phi : A \longrightarrow B$ of graphs takes vertices to vertices, edges to edges and satisfies:

$$\overline{\phi(e)} = \phi(\bar{e}),$$

$$\partial_0\phi(e) = \phi(\partial_0(e)),$$

$$\partial_1\phi(e) = \phi(\partial_1(e)),$$

where $\partial_0(e)$ and $\partial_1(e)$ denote the initial and terminal vertices of an edge, respectively. An *isomorphism* of graphs is a morphism which is bijective on both vertices and edges.

Let A be a graph. The *diameter* of A is

$$diam(A) := \max_{a,b\in VA} d(a,b),$$

where $d(a,b)$ is the length of the shortest reduced path between a and b. Let $v \in VA$. The *star* of v, denoted $E_0(v)$ is

$$E_0(v) := \{e \in EA | \partial_0(e) = v\}.$$

Let $C \subseteq VA$ be a subset of vertices. Let $v \in VA$ be a vertex. A *geodesic* from v to C, denoted $[v, C]$, is a shortest reduced path from v to a vertex in C, and the *distance* from v to C is defined as

$$d(v,C) := \min_{x\in C} d(v,x) = |[v,C]|.$$

An *edge-indexed graph* (A, i) consists of an underlying graph A, and an assignment of a positive integer $i(e) > 0$ to each oriented edge $e \in EA$. An edge-indexed graph (A, i) is *finite* if it has finitely many vertices and finitely many edges but we allow for the possibility that for some $e \in EA$, $i(e) = \infty$. For $v \in VA$, the *degree* of v in (A, i) is defined as

$$deg_{(A,i)}(v) := \sum_{e\in E_0(v)} i(e).$$

An edge-indexed graph (A, i) determines its universal covering tree $X = \widetilde{(A, i)}$ up to isomorphism ([2], Chapter 2). Every edge-indexed graph arises as a quotient of its universal covering tree $X = \widetilde{(A, i)}$ by a subgroup of $G = Aut(X)$.

The diameter of an edge-indexed graph (A, i) is the diameter of its underlying graph A.

An *isomorphism* $\rho : (A, i) \longrightarrow (B, j)$ of edge-indexed graphs is a morphism $\rho : A \longrightarrow B$ such that for each $v \in VA$

$$deg_{(A,i)}(v) = deg_{(B,j)}\rho(v).$$

If $\phi : A \longrightarrow B$ is a graph isomorphism, then ϕ satisfies the following *continuity rule*: if $v \in VA$, then ϕ maps the neighbours (the vertices at distance 1 from v) of v bijectively to the neighbours of $\phi(v)$.

3. The BT degree refinement algorithm

Following [3], let V be a set, and let $Eq(V)$ denote the set of all equivalence relations on V. For $R \in Eq(V)$ and $x \in V$, let x_R denote the R-class of x.

BT introduced the *degree refinement operator* on (A, i)

$$\rho : Eq(VA) \longrightarrow Eq(VA)$$

defined on $R \in Eq(VA)$ as follows:

$$(a, b) \in \rho R \iff (a, b) \in R \quad \text{and} \quad i(a, c_R) = i(b, c_R) \quad \text{for } c \in VA,$$

where for $C, D \subseteq VA$ we set

$$i(C, D) := \sum_{e \in E(C,D)} i(e),$$

with

$$E(C, D) := \{e \in EA | \partial_0 e \in C,\ \partial_1 e \in D\}.$$

When $C = \{a\} \in VA$ we write

$$i(a, D) = i(\{a\}, D).$$

Next we define $R_N = R_N(A, i)$ inductively as follows:

$$R_0 = VA \times VA,$$

$$R_{N+1} = \rho R_N \subset R_N$$

for $N \geqslant 0$, and we put

$$R_* = R_*(A, i) = \cap_{N \geqslant 0} R_N.$$

Thus

$$(a, b) \in R_0 \quad \text{for all } a, b \in VA$$

and

$$(a, b) \in R_1 \iff deg_{(A,i)}(a) = deg_{(A,i)}(b).$$

We will refer to $N \in \mathbb{N} \cup \{*\}$ as *step N* of the BT degree refinement algorithm, and to the elements of R_N as *classes at step N*, or as *R_N-classes*.

Proposition 1 (*Bass and Tits [3], (6.6)*). *We have*

(a) $\rho R_* = R_*$.
(b) *if* $R \in Eq(VA)$ *and* $\rho R = R$ *then* $R \subset R_*$.
(c) $R_* = R_N$ *if* $N \geqslant |VA|$.

Let

$$(A_*, i_*) := (A, i)/R_*.$$

We call (A_*, i_*) the *ultimate quotient* of (A, i), or of X modulo G. The following theorem justifies the use of this terminology.

Theorem 1 (*Bass and Tits [3], (6.6)*). *Let* (A, i) *be an edge-indexed graph, let* $X = \widetilde{(A, i)}$, *and let* $G = Aut(X)$. *Then*

$$(A_*, i_*) = I(G\backslash\backslash X),$$

where $G\backslash\backslash X$ *denotes the quotient graph of groups for X modulo G, and* $I(G\backslash\backslash X)$ *denotes its edge-indexed quotient graph.*

The algorithm terminates when the equivalence classes stabilize in a finite number of steps. In this case

$$\rho(R_n) = \rho(R_{n+1}) = R_*$$

for some $n < \infty$. If (A, i) is finite, this occurs in a finite number of steps.

4. Diameter of the ultimate quotient

In this section we start with any edge-indexed graph, denoted (A, i), and using the BT algorithm we produce an invariant which is a sequence of positive integers, denoted $(d_s)_{s \geqslant 0}$, where each element of the sequence is determined at a step s of the algorithm. This gives a simple and effective way to determine if the ultimate quotient of (A, i) is infinite, often requiring only steps of the BT algorithm, even though the BT algorithm may not terminate. The authors would like to thank the referee whose comments clarified and simplified our ideas in this section.

We begin with the following observation. Let $H \leqslant G = Aut(X)$ be a subgroup of G and let $Y \subseteq VX$ be an H-invariant subset of vertices of X. Let $x \in VX$. Then

$$d(x, Y) = d(hx, Y)$$

for every $h \in H$. In particular this is true if Y is an equivalence class of vertices at some step of the BT algorithm. It follows easily that if $H \leqslant G = Aut(X)$ and $A = H\backslash X$ is the quotient graph, then a shortest path γ from a vertex $x \in VX$ to an H-invariant subset $Y \subseteq VX$ maps injectively to a shortest path γ_0 in the ultimate quotient (A_0, i_0).

Definition.

Let (A, i) be an edge-indexed graph. Let $s \in \mathbb{Z}_{>0}$ be a step of the BT algorithm applied to (A, i). Let $v \in VA$ and let C be an equivalence class of vertices that occurs at step s. Define

$$d_s := \max_{v \in VA, C} d(v, C),$$

where

$$d(v, C) = \min_{x \in C} d(x, v)$$

and $d(x, v)$ denotes the length of the shortest reduced path in A between x and v in VA. Let $\mathscr{D}$ denote the sequence $(d_s)_{s \in \mathbb{Z}_{>0}}$.

The following lemma is clear.

Lemma 1. *Let (A, i) be an edge-indexed graph. Let (A_0, i_0) be the ultimate quotient of (A, i). Then*

(1) *For each $s \in \mathbb{Z}_{>0}$ we have $d_s \leqslant d_{s+1}$.*
(2) *For all $s \in \mathbb{Z}_{>0}$*

$$d_s \leqslant diam(A_0, i_0).$$

(3)

$$\lim_{s \longrightarrow \infty} d_s = diam(A_0, i_0).$$

It follows from the Lemma that if $d_s < \infty$ for each s, then (A_0, i_0) is finite and

$$diam(A_0) = \max_{u, v \in VA_0} d(u, v) = \lim_{s \longrightarrow \infty} d_s.$$

We say that a tree X is *non-uniform if* X is not the universal covering of a finite connected graph. The following gives a sufficient condition for the universal covering tree of an edge-indexed graph (A, i) to be non-uniform.

Corollary 1. *Let (A, i) be an infinite edge-indexed graph.*

(a) *If*

$$\lim_{s \longrightarrow \infty} d_s = \infty$$

then (A_0, i_0) is infinite.

(b) *If there exists an* R_n-*class* C, $n \geqslant 1$, *of the BT algorithm for* (A, i) *such that*

$$\max_{v \in VA} d(c, C) = \infty$$

then (A_0, i_0) *is infinite*.

Proof. For (a), if $\lim_{s \longrightarrow \infty} d_s = \infty$ then $diam(A_0) = \infty$ and thus (A_0, i_0) is infinite. For (b), observe that in this case $d_n = \infty$. □

Condition (b) in Corollary 1 is sufficient but not necessary as is demonstrated by Example 3 below. As an application, if (A, i) has finite volume then its ultimate quotient also has finite volume, and hence if infinite, automatically satisfies the BT criterion for non-discreteness. It follows immediately that if the sequence $(d_s)_{s \geqslant 0}$ diverges, and if (A, i) has finite volume, then the automorphism group of the universal covering tree of (A, i) is not discrete.

Corollary 2. *Let* (A, i) *be an infinite edge-indexed graph and let* $n \geqslant 1$. *If any* R_n-*class of the BT algorithm for* (A, i) *is finite*, *then* (A_0, i_0) *is infinite*.

Proof. Since (A, i) is infinite and any R_n-class C is finite, there are vertices arbitrarily far away from C and therefore there exists $s \in \mathbb{Z}_{>0}$ such that $\max_{v \in VA, C \subseteq VA} d(v, C) = \infty$. □

Finally, we use the results in this section to give a lower bound on the minimum number of generators of a uniform tree lattice. For a detailed discussion of tree lattices and related notions, we refer the reader to [3].

Corollary 3. *Let* (A, i) *be an unimodular edge-indexed graph with universal covering tree* $X = \widetilde{(A, i)}$. *Let* (A_0, i_0) *be the ultimate quotient of* (A, i) *and suppose that* (A_0, i_0) *is finite*. *Let* $d = diam(A_0, i_0)$. *Then there is a uniform* X- *lattice* Γ *with the same quotient graph as* $G = Aut(X)$ *and with at least* d *generators*.

Proof. Since (A_0, i_0) is finite and automatically unimodular, by [3] there is a uniform X-lattice Γ with

$$I(\Gamma \backslash\backslash X) = I(G \backslash\backslash X) = (A_0, i_0),$$

where $I(\Gamma \backslash\backslash X)$ and $I(G \backslash\backslash X)$ denote the edge-indexed graphs for quotients of X by Γ and $G = Aut(X)$, respectively. Then

$$\begin{aligned} d = diam(A_0, i_0) = diam(\Gamma \backslash X) &\leqslant \text{number of vertices of } \Gamma \backslash X \\ &\leqslant \text{number of generators of } \Gamma. \end{aligned}$$ □

5. Examples

In this section we give a number of examples.

Example 1. A graph of diameter 2 (see also ([3], p. 211)).

A=

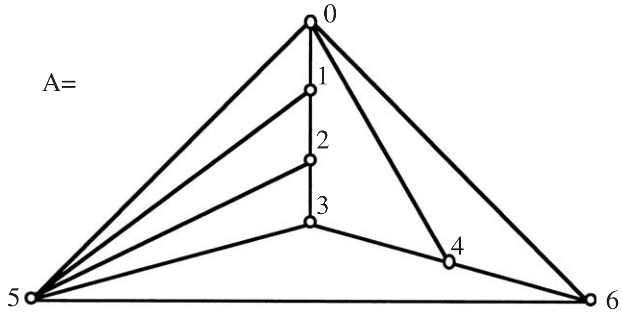

As in ([3], p. 211) we have

R_1**-classes:** {0}, {1, 2, 3, 4, 6}, {5}.
R_2**-classes:** {0}, {1, 6}, {2, 3}, {4}, {5}.
R_3**-classes:** {0}, {1}, {2}, {3}, {4}, {5}, {6}.

We have $d_1 = 1$ since

0 is adjacent to 6,
5 is adjacent to 0,
5 is adjacent to 6

and $d_2 = 2$ since {2, 3} has distance 2 from {0} and all other distances are smaller. Moreover $d_n = 2, n \geqslant 3$ since even though 2 and 3 are separated in $R_n, n \geqslant 3$, {3} has distance 2 from {0}.

Example 2. An infinite edge-indexed graph with finite ultimate quotient.
Let (A, i)=

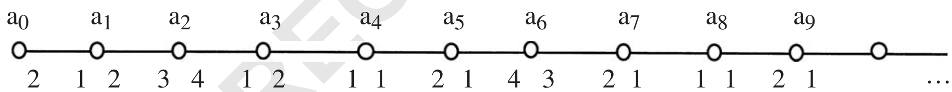

We claim that (A_0, i_0)=

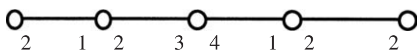

and so (A_0, i_0) has diameter 4.
We have

R_1**-classes:** $\{a_0, a_4, a_8, \ldots\}$, $\{a_1, a_3, a_5, \ldots\}$, $\{a_2, a_6, a_{10}, \ldots\}$.
R_2**-classes:** $\{a_0, a_4, a_8, \ldots\}$, $\{a_2, a_6, a_{10}, \ldots\}$, $\{a_3, a_5, a_{11}, a_{13}, \ldots\}$, $\{a_1, a_7, a_9, a_{15}, a_{17}, \ldots\}$.
R_3**-classes:** $\{a_0, a_8, a_{16}, \ldots\}$, $\{a_4, a_{12}, a_{20}, \ldots\}$, $\{a_1, a_7, a_9, \ldots\}$, $\{a_2, a_6, a_{10}, \ldots\}$, $\{a_3, a_5, a_{11}, a_{13}, \ldots\}$

and $R_n = R_3$, $n > 3$. Thus the ultimate quotient is (A_0, i_0) as above. Moreover

$d_1 = 2$ since a_2 is not adjacent to a_0 or a_4 and all other distances are smaller,
$d_2 = 3$ since a_0 has distance 3 from a_3 and all other distances are smaller,
$d_3 = 4$ since a_0 has distance 4 from a_4 and all other distances are smaller.

Example 3. Fix $n \in \mathbb{Z}_{>0}$. Here $d_s < \infty$ for $s < n$ but $d_n = \infty$ and the BT algorithm does not terminate.

Consider a semi-infinite ray (A, i), with $VA = \{a_0, a_1, \ldots\}$, $EA = \{e_0, \overline{e_0}, e_1, \ldots\}$. Choose $k \in \mathbb{Z}_{\geqslant 0}$ and define indices for all $j \in \mathbb{Z}_{\geqslant 0}$ as follows:

$$i(e_0) := 2.$$

If k divides j and $j \neq 0$:

$$i(e_j) := 2, \quad i(\overline{e_{j-1}}) := 2,$$

otherwise

$$i(e_j) := 1, \quad i(\overline{e_{j-1}}) = 1.$$

Then the R_n-classes for (A, i) in the BT algorithm are

R_1 classes:

$$\{a_0, a_1, \ldots, a_{k-1}, a_{k+1}, \ldots\}, \\ \cup \\ \{a_k, a_{2k}, \ldots\}.$$

R_2 classes:

$$\{a_0, a_1, \ldots, a_{k-2}, a_{k+2}, \ldots\}, \\ \cup \\ \{a_k, a_{2k}, \ldots\}, \\ \cup \\ \{a_{k-1}, a_{k+1}, a_{2k-1}, a_{2k+1}, \ldots\}. \\ \vdots$$

$R_{(k/2)-1}$ classes:

$$\{a_0, a_1, \ldots, a_{k/2}\}, \\ \cup \\ \{a_k, a_{2k}, \ldots\}, \\ \cup \\ \{a_{k-1}, a_{k+1}, a_{2k-1}, a_{2k+1}, \ldots\}, \\ \cup \\ \{a_{k-2}, a_{k+2}, a_{2k-2}, a_{2k+2}, \ldots\}, \\ \cup \\ \{a_{(k/2)+1}, a_{(3k/2)-1}, a_{(3k/2)+1}, a_{(5k/2)-1}, \ldots\}.$$

Then $d_s = k$ for $s < (k/2) - 1$, but the first $R_{(k/2)-1}$ class is finite and so by Corollary 5, $d_{(k/2)-1} = \infty$. Hence the ultimate quotient of (A, i) is infinite. Moreover

$$p^{-1}(v) \quad \not\supseteq \quad \cap_{C \in R_k} N(d(x, C), C)$$

for $k < n$.

Example 4. $d_s < \infty$ for every $s \geqslant 1$ but $\lim_{s \to \infty} d_s = \infty$.

We define a tree (A, i) recursively as follows. Start with a semi-infinite ray (A', i'), and let $VA' = \{a_0, a_1, \ldots\}$. Let

$T_0 :=$

and attach a terminal vertex of T_0 to a_k, k odd. Let $a_k \in VA'$, $k \geqslant 1$. If $2 \mid k$ but $2^n \nmid k$ for any $n > 1$ then attach the following tree denoted T_1 to vertex a_k at the bold vertex:

$T_1 :=$

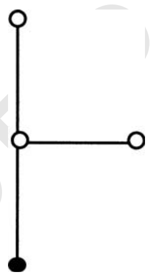

If $2^j \mid k$ but $2^{j+n} \nmid k$ for any $n > 1$ then attach the following tree denoted T_j to vertex a_k at the bold vertex:

$T_j :=$

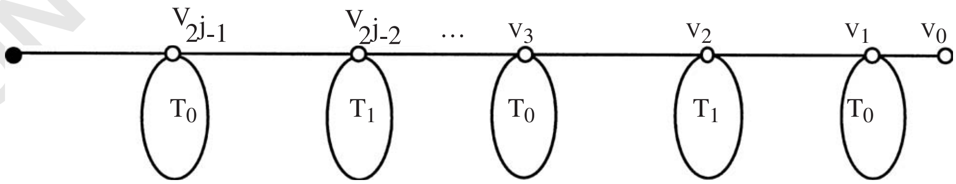

Then

$(A, i) :=$

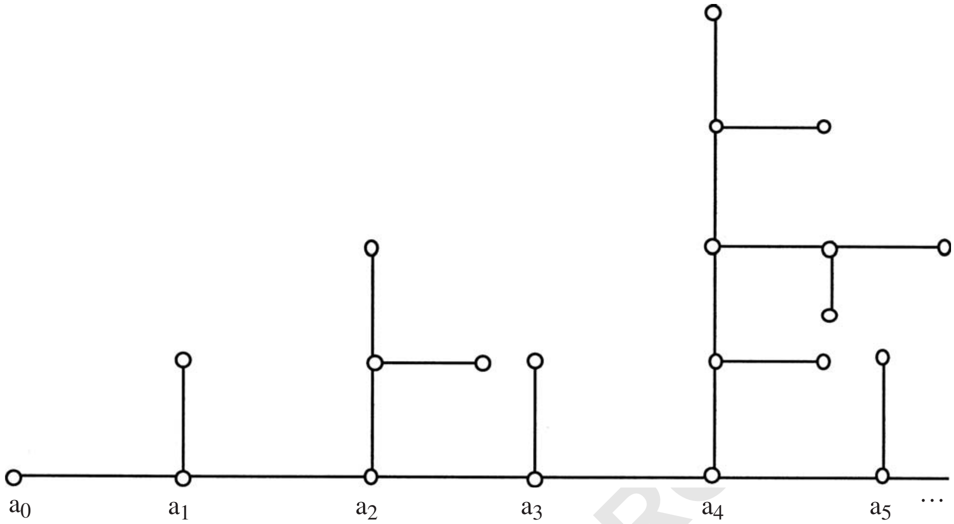

It can be shown that $d_s < \infty$ for every $s \geqslant 1$, in fact

$$d_s < 2^{\lceil \log_2 s \rceil + 4}.$$

But $\lim_{s \longrightarrow \infty} d_s = \infty$. The verification of these facts is routine but lengthy and is left to the reader.

6. Automata, regular languages and state minimization

In this section we recall the basic properties of finite state automata, regular languages and the state minimizing algorithm. Our reference for this section is [4].

An *alphabet* Ω is a finite set, and an element $\omega \in \Omega$ is called a *letter* of Ω. A finite sequence of letters is called a *string*. We use Ω^* to denote the set of all strings over Ω. A *language* over Ω is a subset of Ω^*.

Let Ω be an alphabet. A quintuple $\mathscr{F} = (S, \Omega, \mu, F, s_0)$ is a (deterministic) *finite state automaton* if the following conditions are satisfied:

(1) S is a finite set called the *state set*, F is a subset of S and $s_0 \in S$. The elements of S are called *states*, $s_0 \in S$ is called the *initial state*, and the elements of F are called *final states*.
(2) μ is a map from $S \times \Omega$ to S.

Let $\mathscr{F} = (S, \Omega, \mu, F, s_0)$ be a finite state automaton. For a string $w = u_1, u_2, \ldots, u_n \in \Omega^*$ and a state $s \in S$ we define $\mu^* : S \times \Omega^* \longrightarrow S$ by

$$\mu^*(s, w) = \mu(\ldots \mu(\mu(s, a_1), a_2) \ldots, a_n).$$

We say that a string w is *accepted* by $\mathscr{F}$ if $\mu^*(s_0, w) \in F$. We let $L(\mathscr{F})$ denote the set of strings accepted by $\mathscr{F}$.

A language $\mathscr{L}$ over an alphabet Ω is called *regular* if there is a finite state automaton $\mathscr{F}$ over Ω such $\mathscr{L} = L(\mathscr{F})$.

To each finite state automaton $\mathscr{F}$ is associated a finite directed graph Σ. The vertices of Σ represent the states of $\mathscr{F}$ and the edges of Σ are indexed by symbols which represent the transitions between states. The final states of $\mathscr{F}$ are represented by distinguished vertices of the graph Σ.

A state s of a finite state automaton $\mathscr{F}$ over Ω that cannot be reached by any string of Ω^* from the initial state is called an *inaccessible state*. A state that is not a final state and from which no final state can be reached by any string of Ω^* is called a *failure state*.

We can obtain a *minimized* automaton with the same language as $\mathscr{F}$ in two steps. We first create a 'reduced' automaton by removing all inaccessible states and amalgamating all failure states into a single failure state, thus removing obvious redundancies. Second, we run the *state minimizing algorithm* to reduce the states. This is described as follows. Let $\mathscr{F} = (S, \Omega, \mu, F, s_0)$ be a finite state automaton. Let U_i, $i \geqslant 0$ denote the following equivalence relation on S:

$$U_0 = S,$$

$$U_1 = \{\text{final states}\} \cup \{\text{non final states}\} \subset U_0.$$

For each $i \geqslant 1$, $r, s \in S$ we obtain U_{i+1} from U_i by

$$r \equiv_{U_{i+1}} s \text{ if and only if } r \equiv_{U_i} s$$

and for each $w \in \Omega$

$$\mu(w, r) \equiv_{U_i} \mu(w, s).$$

For some $i \geqslant 1$ we have

$$U_i = U_{i+1} = U_{i+2} \ldots$$

Let U_∞ denote the stable equivalence relation with corresponding *minimized* automaton

$$\mathscr{F}_{\min} = (S_{\min}, \Omega, \mu_{\min}, F_{\min}, s_0),$$

where $S_{\min} = S/U_\infty$, $F_{\min} = F/U_\infty$, and for $[s] \in S_{\min}$, $s \in [s]$ we have

$$\mu_{\min}(\omega, [s]) = [\mu(\omega, s)].$$

Thus $L(\mathscr{F}) = L(\mathscr{F}_{\min})$, and $\mathscr{F}_{\min}$ is the unique smallest automaton with this property. If $\mathscr{F} = (S, \Omega, \mu, F, s_0)$ is any automaton (with S not necessarily finite) and $\mathscr{F}_{\min} = (S_{\min}, \Omega, \mu_{\min}, F_{\min}, s_0)$ is the corresponding minimized automaton, the Myhill–Nerode Theorem [4] implies that the language $L(\mathscr{F})$ is regular if and only if $S_{\min}$ is a finite set.

7. The edge-indexed graph of an automaton

Let $\mathscr{F}$ be a finite state automaton. In this section we associate a (non-unique) edge-indexed graph to $\mathscr{F}$. We recall that $\mathscr{F}$ has an associated finite directed graph Σ whose vertices represent the states of $\mathscr{F}$ and whose edges are indexed by symbols which represent the transitions between states. The graph Σ contains a few distinguished vertices which represent the final states of $\mathscr{F}$. Let $S_1, S_2, \ldots, S_n$ denote the symbols that index the edges of Σ. Choose bijections

$$\begin{aligned} f : \{S_1, S_2, \ldots, S_n\} &\longrightarrow \{2, 2^2, \ldots, 2^n\}, \\ S_j &\mapsto 2^j, \\ g : \{S_1, S_2, \ldots, S_n\} &\longrightarrow \{2^{n+1}, 2^{n+2}, \ldots, 2^{n+n}\}, \\ S_j &\mapsto 2^{n+j}. \end{aligned}$$

We build a directed edge-indexed graph $(A, i)_{\mathscr{F}}$ as follows. We take $A = \Sigma$. Let $e \in E\Sigma$. If e has as its initial point a non-final state, say S_j, we choose $i(e) = f(S_j) = 2^j$. If the initial point of e is a final state, we choose $i(e) = g(S_j) = 2^{n+j}$.

Theorem 2. *Let $\mathscr{F}$ be a finite state automaton with edge-indexed graph $(A, i)_{\mathscr{F}}$, and let $\mathscr{F}_0$ be the minimized automaton. Let $I(\mathscr{F}_0)$ be the edge-indexed graph of the minimized automaton. Let $(A_0, i_0)_{\mathscr{F}}$ be the BT ultimate quotient of the edge-indexed graph $(A, i)_{\mathscr{F}}$. Then*

$$I(\mathscr{F}_0) \cong (A_0, i_0)_{\mathscr{F}},$$

where $\cong$ is a graph isomorphism preserving edge-indices.

Proof. We claim that at each step of each algorithm, the equivalence classes coincide. We prove this by induction on the step of the given algorithm, denoted n.

For the BT algorithm, the R_1-classes (step 1) are determined by degrees of the vertices. For the edge-indexed graph of the finite state automaton $\mathscr{F}$, there are only two types of vertices, namely, those of degree $2^{k+1}-1$ and those of degree $2^{2k+1}-1$ for each $k=0, 1, \ldots$. The vertex v has degree $2^{k+1}-1$ if and only if it corresponds to a non-final state of $\mathscr{F}$ and v has degree $2^{2k+1}-1$ if and only if it corresponds to a final state of $\mathscr{F}$. Hence the claim is true for $n = 1$.

Now let $n = s \geqslant 1$. Then a pair (u, v) of vertices belongs to the R_{s+1}-classes of the BT algorithm if and only if using notation as in Section 3 we have

$$i(u, x) = i(v, x)$$

and x belongs to an equivalence class at the previous step, step s, of the BT algorithm. But for each vertex $x \in VA$, the pair (u, v) belongs to the R_s equivalence classes of the state minimizing algorithm if and only if u transitions into x at step s by the same symbols as v does. That is, $(u, v) \in R_{s+1}$. □

8. A lower bound on the minimum number of states

Let (A, i) be an edge-indexed graph. Let $s \in \mathbb{Z}_{>0}$ be a step of the BT algorithm applied to (A, i). Let $v \in VA$ and let C be an equivalence class of vertices that occurs at step s. As in Section 5 we define

$$d_s := \max_{v \in VA, C} d(v, C),$$

where

$$d(v, C) = \min_{x \in C} d(x, v)$$

and $d(x, v)$ denotes the length of the shortest reduced path in A between x and v in VA.

Let (A_0, i_0) be the ultimate quotient of (A, i). Then by Theorem 4 the sequence $(d_s)_{s \in \mathbb{Z}_{>0}}$ is strictly increasing, for all $s \in \mathbb{Z}_{>0} d_s \leqslant diam(A_0, i_0)$, and $\lim_{s \longrightarrow \infty} d_s = diam(A_0, i_0)$.

Let $\mathscr{F}$ denote a finite state automaton with corresponding edge-indexed graph $(A, i)_{\mathscr{F}}$. Let $(A_0, i_0)_{\mathscr{F}} = I(\mathscr{F}_0)$ denote the edge-indexed graph of the minimized automaton $\mathscr{F}_0$. Then

$$diam(A_0, i_0)_{\mathscr{F}} \leqslant \text{no. of vertices in } (A_0, i_0)_{\mathscr{F}} = \text{ min no. of states of } \mathscr{F}.$$

It follows that if $(A_0, i_0)_{\mathscr{F}}$ is infinite, then the minimized automaton $\mathscr{F}_0$ has infinitely many states.

The following is a corollary of the results in Sections 8 and 9.

Corollary 4. *Let $\mathscr{F}$ denote a finite state automaton with corresponding edge-indexed graph $(A, i) = (A, i)_{\mathscr{F}}$. Let $(A_0, i_0) = (A_0, i_0)_{\mathscr{F}}$ denote the edge-indexed graph of the minimized automaton $\mathscr{F}_0$. Consider the minimization algorithm of Bass and Tits applied to (A, i). The following conditions are equivalent*:

(1) $d_1 = \infty$.
(2) *$\mathscr{F}$ contains states arbitrarily far from an accept state*.

Under the above conditions the ultimate quotient (A_0, i_0) is infinite and $L(\mathscr{F})$ is not regular.

9. Languages for group theoretic decision problems

In this section we investigate the languages of certain group theoretic decision problems for the automaton of a Cayley graph of a finitely generated group. Using the edge-indexed graph of an automaton as a tool, we give a new proof that the language for the word problem in a finitely generated group is regular if and only if the group is finite, and a new proof that the language of the membership problem for finite subgroups is not regular.

We will make use of the Myhill–Nerode Theorem [4] which provides a necessary and sufficient condition for a language to be regular.

Given a language L, define an equivalence relation R_L on strings by $x R_L y$ if there is no distinguishing extension z with the property that exactly one of the strings xz and yz is in

L. The Myhill–Nerode Theorem states that the number of states in the smallest automaton accepting L is equal to the number of equivalence classes in R_L.

A consequence of the Myhill–Nerode Theorem is that a language L is regular if and only if the number of equivalence classes of R_L is finite. An immediate corollary is that if a language defines an infinite set of equivalence classes, it is not regular.

Let G be a finitely generated group and fix a finite generating set X for G. We assume that X is symmetric, that is $X = X^{-1}$. The *word problem language*, denoted $WP(G)$, for G is the following:

$$WP(G) = \{w \mid w \in (X \cup X^{-1})^*,\ w =_G 1\},$$

where $(X \cup X^{-1})^*$ consists of all possible words in X and X^{-1}.

In [1], Anisimov and Seifert show that a group has a regular word problem if and only if it is finite. In this case $WP(G)$ is the language of a finite state automaton $\mathscr{F}$.

The automaton of a finitely generated group G is obtained readily from its Cayley graph $\Gamma(G, X)$ which we now describe. The vertices of $\Gamma(G, X)$ are the elements of G. Vertices g and h of $\Gamma(G, X)$ are connected by an oriented edge e if $gs =_G h$ for some $s \in X$, and we label the edge e with the generator s. Thus the edge e^{-1} is labeled with the generator s^{-1}. We can view $\Gamma(G, X)$ as an automaton F^G where the vertices correspond to states and the edges to transition functions between states. The unique initial and final states correspond to the identity group element 1_G. Let $(A, i)_G$ denote the edge-indexed graph associated to this automaton, and let $(A_0, i_0)_G$ denote its BT ultimate quotient. We give a short proof of the AS theorem.

Theorem 3. *The language $WP(G)$ is regular if and only if $(A, i)_G$ is finite if and only if G is finite.*

Proof. Suppose first that $(A, i)_G$ is finite. Then $(A_0, i_0)_G$ is finite and so the minimization F_0^G of F^G has finitely many states, thus the language $WP(G)$ is regular.

Conversely, the language $WP(G)$ is regular if and only if the minimization F_0^G of F_G is finite if and only if $(A_0, i_0)_G = I(F_0^G)$ is finite. We claim that this implies that $(A, i)_G$ is finite. Assume for the sake of contradiction that $(A, i)_G$ is infinite. Then then R_1-classes of the BT algorithm for $(A, i)_G$ are

$$\{1_G\} \cup G - \{1_G\},$$

since the vertex corresponding to $\{1_G\}$ represents a final state and so its (indexed) degree is

$$deg_{(A,i)_G}(1_G) = 2^{n+1} + 2^{n+2} + \cdots + 2^{n+n} = 2^n(2^1 + 2^2 + \cdots + 2^n).$$

For $v \neq 1_G$,

$$deg_{(A,i)_G}(v) = 2^1 + 2^2 + \cdots + 2^n,$$

hence $deg_{(A,i)_G}(1_G) = 2^n deg_{(A,i)_G} \neq deg_{(A,i)_G}(v)$ for $v \neq 1_G$, so the equivalence class of 1_G is finite. By Corollary 5 if any R_n-class of the BT algorithm for $(A, i)_G$ is finite, the ultimate quotient $(A_0, i_0)_G$ is infinite, this is a contradiction, as $(A_0, i_0)_G$ is finite. Thus $(A, i)_G$ is finite. □

Theorem 4. *Let G be a finitely generated group with finite generating set X, closed under inversion. Let H be a subgroup of G. Let $L(H)$ be the set of all words in X that represent elements of H. Then $L(H)$ is regular if and only if the index of H in G is finite.*

Proof. Let A be the Cayley graph of G with respect to X. We can construct an infinite automaton for $L(H)$ by letting the identity be the start state and setting every vertex corresponding to an element of H to be an accept state. Now let (A_H, i_H) be the corresponding edge-indexed graph. Let $(g_1, \ldots, g_n, \ldots)$ be a set of words in X representing right coset representatives of H of minimal length. We will make use of the Schreier graph of right cosets, where the vertices correspond to right cosets and edges correspond to right multiplication by generators.

If $[G : H]$ is finite, the Schreier graph is finite, yielding an automaton for $L(H)$ by setting the identity coset to be the unique start and end state. Now assume $[G : H]$ is not finite. Let $C \in R_1$ be the class in the first step of the equivalence relation of the BT algorithm for (A_H, i_H), and consider the set of coset representatives $\{g_1, \ldots\}$. Since the g_i are of minimal length, $d(g_i, C)$ is the length of the word g_i. However, since there are infinitely many g_i, they cannot have length bounded below by a constant. Thus $d_1 = \infty$, so by Corollary 7, $L(H)$ is not regular. □

Acknowledgements

The authors would like to thank the University of Hong Kong whose hospitality is greatly appreciated.

References

[1] A.V. Anisimov, F.D. Seifert, Zur algebraischen Charakteristik der durch kontext-freie Sprachen definierten Gruppen, Elektron, Informationsverarb, Kyber-net, vol. 11, 1975, pp. 695–702.
[2] H. Bass, A. Lubotzky, Tree Lattices, Progress in Mathematics, vol. 176, Birkhauser, Boston, 2000.
[3] H. Bass, J. Tits, A discreteness criterion for certain tree automorphism groups, Appendix [3] in Tree Lattices, in: Progress in Mathematics, vol. 176, Birkhauser, Boston, 2000.
[4] J. Hopcroft, R. Motwani, J. Ullman, Introduction to Automata Theory, Languages, and Computation, Addison-Wesley, Reading, 2000.