Nielsen equivalence in a class of random groups

(joint work with Ilya Kapovich)

Richard Weidmann

CAU Kiel

17.4.2014

Theorem 1

For every $n \geq 2$ there exists a torsion-free one-ended word-hyperbolic group G of rank n admitting generating n-tuples $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ such that the $(2n-1)$-tuples

$$(a_1, \ldots, a_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}}) \text{ and } (b_1, \ldots, b_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}})$$

are not Nielsen-equivalent in G.

The groups are constructed using a probabilistic construction.

Theorem 1

For every $n \geq 2$ there exists a torsion-free one-ended word-hyperbolic group G of rank n admitting generating n-tuples $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ such that the $(2n-1)$-tuples

$$(a_1, \ldots, a_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}}) \text{ and } (b_1, \ldots, b_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}})$$

are not Nielsen-equivalent in G.

The groups are constructed using a probabilistic construction.

More precisely: We consider groups given by presentations of type

$$G = \langle a_1, \ldots, a_n, b_1, \ldots, b_n | a_i = u_i(\underline{b}), b_i = v_i(\underline{a}), \text{ for } i = 1, \ldots, n \rangle.$$

where the $u_i(\underline{b})$ are reduced words in the $b_i^{\pm 1}$ and the $v_i(\underline{a})$ are reduced words in the $a_i^{\pm 1}$ such that

$$|v_1| = \cdots = |v_n| = |v_n| = \cdots = |u_n| = N$$

for some $N \in \mathbb{N}$.

It is trivial that $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ are generating tuples.

We show that as N tends to infinity the probability that such a group satisfies the conclusion of the Theorem tends to 1 if the u_i and v_i are chosen at random.

More precisely: We consider groups given by presentations of type

$$G = \langle a_1, \ldots, a_n, b_1, \ldots, b_n | a_i = u_i(\underline{b}), b_i = v_i(\underline{a}), \text{ for } i = 1, \ldots, n \rangle.$$

where the $u_i(\underline{b})$ are reduced words in the $b_i^{\pm 1}$ and the $v_i(\underline{a})$ are reduced words in the $a_i^{\pm 1}$ such that

$$|v_1| = \cdots = |v_n| = |v_n| = \cdots = |u_n| = N$$

for some $N \in \mathbb{N}$.

It is trivial that $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ are generating tuples.

We show that as N tends to infinity the probability that such a group satisfies the conclusion of the Theorem tends to 1 if the u_i and v_i are chosen at random.

More precisely: We consider groups given by presentations of type

$$G = \langle a_1, \ldots, a_n, b_1, \ldots, b_n | a_i = u_i(\underline{b}), b_i = v_i(\underline{a}), \text{ for } i = 1, \ldots, n\rangle.$$

where the $u_i(\underline{b})$ are reduced words in the $b_i^{\pm 1}$ and the $v_i(\underline{a})$ are reduced words in the $a_i^{\pm 1}$ such that

$$|v_1| = \cdots = |v_n| = |v_n| = \cdots = |u_n| = N$$

for some $N \in \mathbb{N}$.

It is trivial that $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ are generating tuples.

We show that as N tends to infinity the probability that such a group satisfies the conclusion of the Theorem tends to 1 if the u_i and v_i are chosen at random.

Let G be a group and $n \in \mathbb{N}$. Consider the set G^n of n-tuples of elements of G.

We say that $T = (g_1, \ldots, g_n)$ and $T' = (g_1', \ldots, g_n')$ are elementary equivalent if one of the following holds:

1. $g_i' = g_{\sigma(i)}$ for all i and some $\sigma \in S_n$.
2. $g_i' = g_i^{-1}$ for some i and $g_j' = g_j$ for $j \neq i$.
3. $g_i' = g_i g_j$ for some $i \neq j$ and $g_k' = g_k$ for $k \neq i$.

We say that T and T' are Nielsen equivalent and write $T \sim T'$ if there exist

$$T = T_0, T_1, \ldots, T_m = T'$$

such that T_{i-1} and T_i elementary equivalent for $1 \leq i \leq m$.

Let G be a group and $n \in \mathbb{N}$. Consider the set G^n of n-tuples of elements of G.

We say that $T = (g_1, \ldots, g_n)$ and $T' = (g_1', \ldots, g_n')$ are elementary equivalent if one of the following holds:

1. $g_i' = g_{\sigma(i)}$ for all i and some $\sigma \in S_n$.
2. $g_i' = g_i^{-1}$ for some i and $g_j' = g_j$ for $j \neq i$.
3. $g_i' = g_i g_j$ for some $i \neq j$ and $g_k' = g_k$ for $k \neq i$.

We say that T and T' are Nielsen equivalent and write $T \sim T'$ if there exist

$$T = T_0, T_1, \ldots, T_m = T'$$

such that T_{i-1} and T_i elementary equivalent for $1 \leq i \leq m$.

Let G be a group and $n \in \mathbb{N}$. Consider the set G^n of n-tuples of elements of G.

We say that $T = (g_1, \ldots, g_n)$ and $T' = (g_1', \ldots, g_n')$ are elementary equivalent if one of the following holds:

1. $g_i' = g_{\sigma(i)}$ for all i and some $\sigma \in S_n$.
2. $g_i' = g_i^{-1}$ for some i and $g_j' = g_j$ for $j \neq i$.
3. $g_i' = g_i g_j$ for some $i \neq j$ and $g_k' = g_k$ for $k \neq i$.

We say that T and T' are Nielsen equivalent and write $T \sim T'$ if there exist

$$T = T_0, T_1, \ldots, T_m = T'$$

such that T_{i-1} and T_i elementary equivalent for $1 \leq i \leq m$.

Alternative definition of Nielsen equivalence:

Can identify elements of G^n with elements of $\text{Hom}(F_n, G)$ via the bijection

$$G^n \to \text{Hom}(F_n, G),\ \ T \mapsto \phi_T$$

where for any $T = (g_1, \ldots, g_n)$ the homomorphism $\phi_T : F_n \to G$ is given by $\phi_T(x_i) = g_i$ for $1 \leq i \leq n$. Note $F_n := F(x_1, \ldots, x_n)$.

Fact: Let $T, T' \in G^n$. Then $T \sim T'$ iff $\phi_T = \phi_{T'} \circ \alpha$ for some $\alpha \in \text{Aut}(F_n)$.

Thus Nielsen equivalence classes of n-tuples correspond to $\text{Aut}(F_n)$-orbits of $\text{Hom}(F_n, G)$ under the natural right action of $\text{Aut}(F_n)$.

Alternative definition of Nielsen equivalence:

Can identify elements of G^n with elements of $\mathrm{Hom}(F_n, G)$ via the bijection

$$G^n \to \mathrm{Hom}(F_n, G), \ \ T \mapsto \phi_T$$

where for any $T = (g_1, \ldots, g_n)$ the homomorphism $\phi_T : F_n \to G$ is given by $\phi_T(x_i) = g_i$ for $1 \leq i \leq n$. Note $F_n := F(x_1, \ldots, x_n)$.

Fact: Let $T, T' \in G^n$. Then $T \sim T'$ iff $\phi_T = \phi_{T'} \circ \alpha$ for some $\alpha \in \mathrm{Aut}(F_n)$.

Thus Nielsen equivalence classes of n-tuples correspond to $\mathrm{Aut}(F_n)$-orbits of $\mathrm{Hom}(F_n, G)$ under the natural right action of $\mathrm{Aut}(F_n)$.

Above fact is a reformulation of the classical result of Nielsen that states that $\mathrm{Aut}(F_n)$ is generated by automorhisms of the following types (now called Nielsen automorphisms):

1. $F_n \to F_n$, $x_i \mapsto x_{\sigma(i)}$ for $1 \leq i \leq n$ and some $\sigma \in S_n$.
2. $F_n \to F_n$, $x_i \mapsto x_i^{-1}$ for some i and $x_j \mapsto x_j$ for $j \neq i$.
3. $F_n \to F_n$, $x_i \mapsto x_i x_j$ for some $i \neq j$ and $x_k \mapsto x_k$ for $k \neq i$.

We will be mostly interested in Nielsen equivalence classes of generating tuples, i.e. in $\mathrm{Aut}(F_n)$-orbits of $\mathrm{Epi}(F_n, G)$.

Above fact is a reformulation of the classical result of Nielsen that states that $\mathrm{Aut}(F_n)$ is generated by automorhisms of the following types (now called Nielsen automorphisms):

1. $F_n \to F_n$, $x_i \mapsto x_{\sigma(i)}$ for $1 \leq i \leq n$ and some $\sigma \in S_n$.
2. $F_n \to F_n$, $x_i \mapsto x_i^{-1}$ for some i and $x_j \mapsto x_j$ for $j \neq i$.
3. $F_n \to F_n$, $x_i \mapsto x_i x_j$ for some $i \neq j$ and $x_k \mapsto x_k$ for $k \neq i$.

We will be mostly interested in Nielsen equivalence classes of generating tuples, i.e. in $\mathrm{Aut}(F_n)$-orbits of $\mathrm{Epi}(F_n, G)$.

Fix a group G. There are a number of natural problems:

1. Is there an algorithm that decides whether two given (generating) n-tuples of G are Nielsen equivalent?
2. Are there at most finitely many Nielsen classes of generating n-tuples of G for given n?
3. Classify all Nielsen-classes of generating n-tuples of G for given n.

Problems are usually very hard and often undecidable.

First problem is at least as hard as the generalized word problem as

$$(g_1, \ldots, g_n, 1) \sim (g_1, \ldots, g_n, h) \iff h \in \langle g_1, \ldots, g_n \rangle.$$

Thus the Rips construction shows that Nielsen-equivalence is undecidable even in small cancellation groups.

Fix a group G. There are a number of natural problems:

1. Is there an algorithm that decides whether two given (generating) n-tuples of G are Nielsen equivalent?
2. Are there at most finitely many Nielsen classes of generating n-tuples of G for given n?
3. Classify all Nielsen-classes of generating n-tuples of G for given n.

Problems are usually very hard and often undecidable.

First problem is at least as hard as the generalized word problem as

$$(g_1, \ldots, g_n, 1) \sim (g_1, \ldots, g_n, h) \iff h \in \langle g_1, \ldots, g_n \rangle.$$

Thus the Rips construction shows that Nielsen-equivalence is undecidable even in small cancellation groups.

Fix a group G. There are a number of natural problems:

1. Is there an algorithm that decides whether two given (generating) n-tuples of G are Nielsen equivalent?
2. Are there at most finitely many Nielsen classes of generating n-tuples of G for given n?
3. Classify all Nielsen-classes of generating n-tuples of G for given n.

Problems are usually very hard and often undecidable.

First problem is at least as hard as the generalized word problem as

$$(g_1, \ldots, g_n, 1) \sim (g_1, \ldots, g_n, h) \iff h \in \langle g_1, \ldots, g_n \rangle.$$

Thus the Rips construction shows that Nielsen-equivalence is undecidable even in small cancellation groups.

Fix a group G. There are a number of natural problems:

1. Is there an algorithm that decides whether two given (generating) n-tuples of G are Nielsen equivalent?
2. Are there at most finitely many Nielsen classes of generating n-tuples of G for given n?
3. Classify all Nielsen-classes of generating n-tuples of G for given n.

Problems are usually very hard and often undecidable.

First problem is at least as hard as the generalized word problem as

$$(g_1, \ldots, g_n, 1) \sim (g_1, \ldots, g_n, h) \iff h \in \langle g_1, \ldots, g_n \rangle.$$

Thus the Rips construction shows that Nielsen-equivalence is undecidable even in small cancellation groups.

Nielsen equivalence plays an important role in finite groups, in particular in relation to the product replacement algorithm. We will only focus on infinite groups.

Some positive results:

1. Nielsen: An generating n-tuple of F_k is Nielsen-quivalent to $(x_1, \ldots, x_k, \underbrace{1, \ldots, 1}_{n-k \text{ times}})$.
2. Grushko: Any generating tuple of $A * B$ is Nielsen equivalent to a tuple $(g_1, \ldots, g_n)$ with $g_i \in A \cup B$ for $1 \leq i \leq n$.
3. An analogue of Nielsen's result for surface groups due to Zieschang and Louder.

Many related results. All proofs are similar. They rely on replacing a given generating tuple with a reduced one by cancellation/folding methods. No need to distinguish Nielsen classes.

Nielsen equivalence plays an important role in finite groups, in particular in relation to the product replacement algorithm. We will only focus on infinite groups.

Some positive results:

1. Nielsen: An generating n-tuple of F_k is Nielsen-quivalent to $(x_1, \ldots, x_k, \underbrace{1, \ldots, 1}_{n-k \text{ times}})$.
2. Grushko: Any generating tuple of $A * B$ is Nielsen equivalent to a tuple $(g_1, \ldots, g_n)$ with $g_i \in A \cup B$ for $1 \leq i \leq n$.
3. An analogue of Nielsen's result for surface groups due to Zieschang and Louder.

Many related results. All proofs are similar. They rely on replacing a given generating tuple with a reduced one by cancellation/folding methods. No need to distinguish Nielsen classes.

Nielsen equivalence plays an important role in finite groups, in particular in relation to the product replacement algorithm. We will only focus on infinite groups.

Some positive results:

1. Nielsen: An generating n-tuple of F_k is Nielsen-quivalent to $(x_1, \ldots, x_k, \underbrace{1, \ldots, 1}_{n-k \text{ times}})$.
2. Grushko: Any generating tuple of $A * B$ is Nielsen equivalent to a tuple $(g_1, \ldots, g_n)$ with $g_i \in A \cup B$ for $1 \leq i \leq n$.
3. An analogue of Nielsen's result for surface groups due to Zieschang and Louder.

Many related results. All proofs are similar. They rely on replacing a given generating tuple with a reduced one by cancellation/folding methods. No need to distinguish Nielsen classes.

Nielsen equivalence plays an important role in finite groups, in particular in relation to the product replacement algorithm. We will only focus on infinite groups.

Some positive results:

1. Nielsen: An generating n-tuple of F_k is Nielsen-quivalent to $(x_1, \ldots, x_k, \underbrace{1, \ldots, 1}_{n-k \text{ times}})$.
2. Grushko: Any generating tuple of $A * B$ is Nielsen equivalent to a tuple $(g_1, \ldots, g_n)$ with $g_i \in A \cup B$ for $1 \leq i \leq n$.
3. An analogue of Nielsen's result for surface groups due to Zieschang and Louder.

Many related results. All proofs are similar. They rely on replacing a given generating tuple with a reduced one by cancellation/folding methods. No need to distinguish Nielsen classes.

Nielsen equivalence plays an important role in finite groups, in particular in relation to the product replacement algorithm. We will only focus on infinite groups.

Some positive results:

1. Nielsen: An generating n-tuple of F_k is Nielsen-quivalent to $(x_1, \ldots, x_k, \underbrace{1, \ldots, 1}_{n-k \text{ times}})$.
2. Grushko: Any generating tuple of $A * B$ is Nielsen equivalent to a tuple $(g_1, \ldots, g_n)$ with $g_i \in A \cup B$ for $1 \leq i \leq n$.
3. An analogue of Nielsen's result for surface groups due to Zieschang and Louder.

Many related results. All proofs are similar. They rely on replacing a given generating tuple with a reduced one by cancellation/folding methods. No need to distinguish Nielsen classes.

Distinguishing Nielsen classes is difficult. An exception is the case $n = 2$, i.e. the case of pairs of elements. There is test provided the conjugace problem is solvable:

If $(g_1, g_2) \sim (h_1, h_2)$ then $[g_1, g_2]$ is conjugate to $[h_1, h_2]^{\pm 1}$.

No such test for $n \geq 3$.

However there are a number of ways to distinguish classes in very specific situations, see work of Zieschang, Rost, Rosenberger, Noskov, Lustig-Moriah, Evans, ...

Distinguishing Nielsen classes is difficult. An exception is the case $n = 2$, i.e. the case of pairs of elements. There is test provided the conjugace problem is solvable:

If $(g_1, g_2) \sim (h_1, h_2)$ then $[g_1, g_2]$ is conjugate to $[h_1, h_2]^{\pm 1}$.

No such test for $n \geq 3$.

However there are a number of ways to distinguish classes in very specific situations, see work of Zieschang, Rost, Rosenberger, Noskov, Lustig-Moriah, Evans, ...

Distinguishing Nielsen classes is difficult. An exception is the case $n = 2$, i.e. the case of pairs of elements. There is test provided the conjugace problem is solvable:

If $(g_1, g_2) \sim (h_1, h_2)$ then $[g_1, g_2]$ is conjugate to $[h_1, h_2]^{\pm 1}$.

No such test for $n \geq 3$.

However there are a number of ways to distinguish classes in very specific situations, see work of Zieschang, Rost, Rosenberger, Noskov, Lustig-Moriah, Evans, ...

Distinguishing Nielsen classes is difficult. An exception is the case $n = 2$, i.e. the case of pairs of elements. There is test provided the conjugace problem is solvable:

If $(g_1, g_2) \sim (h_1, h_2)$ then $[g_1, g_2]$ is conjugate to $[h_1, h_2]^{\pm 1}$.

No such test for $n \geq 3$.

However there are a number of ways to distinguish classes in very specific situations, see work of Zieschang, Rost, Rosenberger, Noskov, Lustig-Moriah, Evans, ...

If $T = (g_1, \ldots, g_n) \in G^n$ then we call the $(n+k)$-tuple

$$(g_1, \ldots, g_n, \underbrace{1, \ldots, 1}_{k \text{ times}})$$

the k-th stabilisation of T.

Often the first stabilisation of two generating tuples are Nielsen equivalent, even if the tuples aren't. Moreover the following is trivial:

If $(g_1, \ldots, g_n)$ and $(g_1', \ldots, g_n')$ are generating n-tuples of G then

$$(g_1, \ldots, g_n, \underbrace{1, \ldots, 1}_{n \text{ times}}) \sim (g_1', \ldots, g_n', \underbrace{1, \ldots, 1}_{n \text{ times}}).$$

Work of Evans implies that in general many stabilisations are needed to make two generating tuples Nielsen equivalent:

If $T = (g_1, \ldots, g_n) \in G^n$ then we call the $(n+k)$-tuple

$$(g_1, \ldots, g_n, \underbrace{1, \ldots, 1}_{k \text{ times}})$$

the k-th stabilisation of T.

Often the first stabilisation of two generating tuples are Nielsen equivalent, even if the tuples aren't. Moreover the following is trivial:

If $(g_1, \ldots, g_n)$ and $(g_1', \ldots, g_n')$ are generating n-tuples of G then

$$(g_1, \ldots, g_n, \underbrace{1, \ldots, 1}_{n \text{ times}}) \sim (g_1', \ldots, g_n', \underbrace{1, \ldots, 1}_{n \text{ times}}).$$

Work of Evans implies that in general many stabilisations are needed to make two generating tuples Nielsen equivalent:

If $T = (g_1, \ldots, g_n) \in G^n$ then we call the $(n+k)$-tuple

$$(g_1, \ldots, g_n, \underbrace{1, \ldots, 1}_{k \text{ times}})$$

the k-th stabilisation of T.

Often the first stabilisation of two generating tuples are Nielsen equivalent, even if the tuples aren't. Moreover the following is trivial:

If $(g_1, \ldots, g_n)$ and $(g_1', \ldots, g_n')$ are generating n-tuples of G then

$$(g_1, \ldots, g_n, \underbrace{1, \ldots, 1}_{n \text{ times}}) \sim (g_1', \ldots, g_n', \underbrace{1, \ldots, 1}_{n \text{ times}}).$$

Work of Evans implies that in general many stabilisations are needed to make two generating tuples Nielsen equivalent:

If $T = (g_1, \ldots, g_n) \in G^n$ then we call the $(n+k)$-tuple

$$(g_1, \ldots, g_n, \underbrace{1, \ldots, 1}_{k \text{ times}})$$

the k-th stabilisation of T.

Often the first stabilisation of two generating tuples are Nielsen equivalent, even if the tuples aren't. Moreover the following is trivial:

If $(g_1, \ldots, g_n)$ and $(g_1', \ldots, g_n')$ are generating n-tuples of G then

$$(g_1, \ldots, g_n, \underbrace{1, \ldots, 1}_{n \text{ times}}) \sim (g_1', \ldots, g_n', \underbrace{1, \ldots, 1}_{n \text{ times}}).$$

Work of Evans implies that in general many stabilisations are needed to make two generating tuples Nielsen equivalent:

Theorem 2 (Evans)

For every $k \geq 1$ there exists a $(2^k + k + 1)$ generated metabelian group G and generating 2^{k+1}-tuples T and T' such that the k-th stabilisations of T and T' are Nielsen-inequivalent.

Note there is still a large gap between the trivial upper bound on the number of stabilisations needed to make two generating tuples equivalent and the number given by Evans.

The result presented in this talk shows that the trivial upper bound is in fact the best possible.

While Evans' methods are algebraic/homological, ours are combinatorial/geometric with a dose of randomness.

Theorem 2 (Evans)

For every $k \geq 1$ there exists a $(2^k + k + 1)$ generated metabelian group G and generating 2^{k+1}-tuples T and T' such that the k-th stabilisations of T and T' are Nielsen-inequivalent.

Note there is still a large gap between the trivial upper bound on the number of stabilisations needed to make two generating tuples equivalent and the number given by Evans.

The result presented in this talk shows that the trivial upper bound is in fact the best possible.

While Evans' methods are algebraic/homological, ours are combinatorial/geometric with a dose of randomness.

Theorem 2 (Evans)

For every $k \geq 1$ there exists a $(2^k + k + 1)$ generated metabelian group G and generating 2^{k+1}-tuples T and T' such that the k-th stabilisations of T and T' are Nielsen-inequivalent.

Note there is still a large gap between the trivial upper bound on the number of stabilisations needed to make two generating tuples equivalent and the number given by Evans.

The result presented in this talk shows that the trivial upper bound is in fact the best possible.

While Evans' methods are algebraic/homological, ours are combinatorial/geometric with a dose of randomness.

Theorem 2 (Evans)

For every $k \geq 1$ there exists a $(2^k + k + 1)$ generated metabelian group G and generating 2^{k+1}-tuples T and T' such that the k-th stabilisations of T and T' are Nielsen-inequivalent.

Note there is still a large gap between the trivial upper bound on the number of stabilisations needed to make two generating tuples equivalent and the number given by Evans.

The result presented in this talk shows that the trivial upper bound is in fact the best possible.

While Evans' methods are algebraic/homological, ours are combinatorial/geometric with a dose of randomness.

Recall that we consider groups of type

$$G = \langle a_1, \ldots, a_n, b_1, \ldots, b_n | a_i = u_i(\underline{b}), b_i = v_i(\underline{a}), \text{ for } i = 1, \ldots, n \rangle.$$

where the u_i and v_i are long random words of the same length N.

We may assume that the presentation is a $C'(1/\lambda)$ small cancellation presentation for λ arbitrarily large. Thus for any $\alpha < 1$ we may assume that any reduced word w in the $a_i^{\pm 1}$ and $b_i^{\pm 1}$ that represents the trivial element contains a subword of a cyclic conjugate of some defining relator (or its inverse) of length at least $\alpha \cdot N$.

It follows in particular that any word w in the $a_i^{\pm 1}$ such that $w =_G b_i$ for some i contains a long subword of some $v_j(\underline{a})$.

Recall that we consider groups of type

$$G = \langle a_1, \dots, a_n, b_1, \dots, b_n | a_i = u_i(\underline{b}), b_i = v_i(\underline{a}), \text{ for } i = 1, \dots, n \rangle.$$

where the u_i and v_i are long random words of the same length N.

We may assume that the presentation is a $C'(1/\lambda)$ small cancellation presentation for λ arbitrarily large. Thus for any $\alpha < 1$ we may assume that any reduced word w in the $a_i^{\pm 1}$ and $b_i^{\pm 1}$ that represents the trivial element contains a subword of a cyclic conjugate of some defining relator (or its inverse) of length at least $\alpha \cdot N$.

It follows in particular that any word w in the $a_i^{\pm 1}$ such that $w =_G b_i$ for some i contains a long subword of some $v_j(\underline{a})$.

Recall that we consider groups of type

$$G = \langle a_1, \dots, a_n, b_1, \dots, b_n | a_i = u_i(\underline{b}), b_i = v_i(\underline{a}), \text{ for } i = 1, \dots, n \rangle.$$

where the u_i and v_i are long random words of the same length N.

We may assume that the presentation is a $C'(1/\lambda)$ small cancellation presentation for λ arbitrarily large. Thus for any $\alpha < 1$ we may assume that any reduced word w in the $a_i^{\pm 1}$ and $b_i^{\pm 1}$ that represents the trivial element contains a subword of a cyclic conjugate of some defining relator (or its inverse) of length at least $\alpha \cdot N$.

It follows in particular that any word w in the $a_i^{\pm 1}$ such that $w =_G b_i$ for some i contains a long subword of some $v_j(\underline{a})$.

We will first argue that (generically) $(a_1, \ldots, a_n) \not\sim (b_1, \ldots, b_n)$. The proof of the Main Theorem follows a similar strategy.

The proof is by contradiction.
Thus we assume that $(a_1, \ldots, a_n) \sim (b_1, \ldots, b_n)$.

Thus there exists a basis $(w_1, \ldots, w_n)$ of the free group $F(A) = F(a_1, \ldots, a_n)$ such that $w_i =_G b_i$ for $1 \le i \le n$. Note that the w_i are just the images of the a_i under some automorphism of $F(A)$.

If Γ is the wedge of n (subdivided) circuits with labels $w_1, \ldots, w_n$ and R_n the rose with n loop edges and labels $a_1, \ldots, a_n$ then the label preserving map is π_1-surjective.

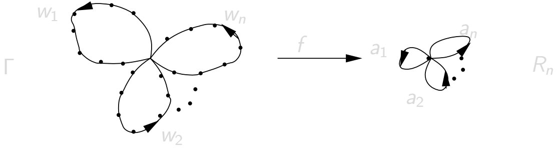

We will first argue that (generically) $(a_1,\ldots,a_n)\not\sim(b_1,\ldots,b_n)$. The proof of the Main Theorem follows a similar strategy.

The proof is by contradiction.
Thus we assume that $(a_1,\ldots,a_n)\sim(b_1,\ldots,b_n)$.

Thus there exists a basis $(w_1,\ldots,w_n)$ of the free group $F(A)=F(a_1,\ldots,a_n)$ such that $w_i=_G b_i$ for $1\le i\le n$. Note that the w_i are just the images of the a_i under some automorphism of $F(A)$.

If Γ is the wedge of n (subdivided) circuits with labels $w_1,\ldots,w_n$ and R_n the rose with n loop edges and labels $a_1,\ldots,a_n$ then the label preserving map is π_1-surjective.

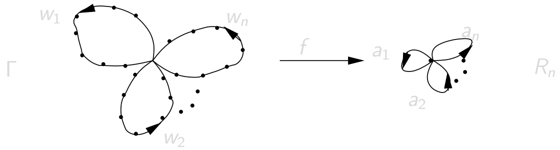

We will first argue that (generically) $(a_1, \ldots, a_n) \not\sim (b_1, \ldots, b_n)$. The proof of the Main Theorem follows a similar strategy.

The proof is by contradiction.
Thus we assume that $(a_1, \ldots, a_n) \sim (b_1, \ldots, b_n)$.

Thus there exists a basis $(w_1, \ldots, w_n)$ of the free group $F(A) = F(a_1, \ldots, a_n)$ such that $w_i =_G b_i$ for $1 \leq i \leq n$. Note that the w_i are just the images of the a_i under some automorphism of $F(A)$.

If Γ is the wedge of n (subdivided) circuits with labels $w_1, \ldots, w_n$ and R_n the rose with n loop edges and labels $a_1, \ldots, a_n$ then the label preserving map is π_1-surjective.

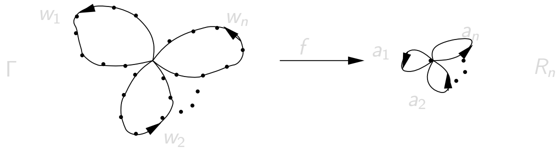

We will first argue that (generically) $(a_1, \ldots, a_n) \not\sim (b_1, \ldots, b_n)$. The proof of the Main Theorem follows a similar strategy.

The proof is by contradiction.
Thus we assume that $(a_1, \ldots, a_n) \sim (b_1, \ldots, b_n)$.

Thus there exists a basis $(w_1, \ldots, w_n)$ of the free group $F(A) = F(a_1, \ldots, a_n)$ such that $w_i =_G b_i$ for $1 \leq i \leq n$. Note that the w_i are just the images of the a_i under some automorphism of $F(A)$.

If Γ is the wedge of n (subdivided) circuits with labels $w_1, \ldots, w_n$ and R_n the rose with n loop edges and labels $a_1, \ldots, a_n$ then the label preserving map is π_1-surjective.

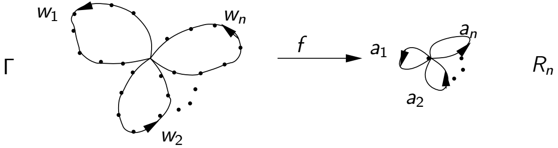

It is a fundamental observation of Stallings (see also Dicks and others) that f factors as a product of folds.

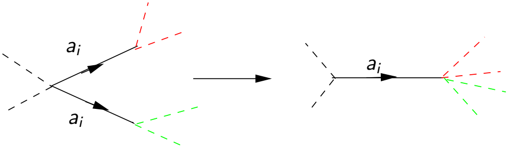

Any w_i must be readable in any graph that occurs in the sequence.
However not everything is readable in the graph that occurs just before R_n.

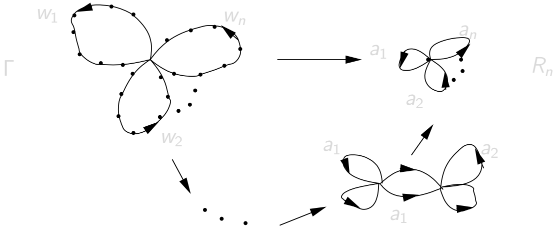

Figure : $a_1^{-1}a_2$ is not readable in the graph that appears just before R_n

It is a fundamental observation of Stallings (see also Dicks and others) that f factors as a product of folds.

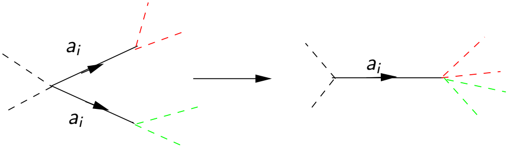

Any w_i must be readable in any graph that occurs in the sequence. However not everything is readable in the graph that occurs just before R_n.

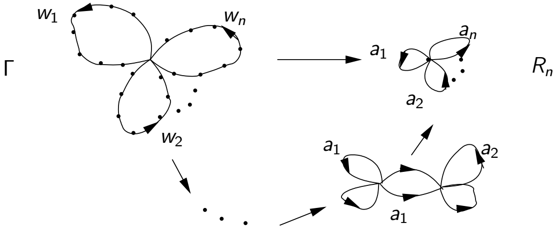

Figure : $a_1^{-1}a_2$ is not readable in the graph that appears just before R_n

Thus we have the following:

1. w_i contains a subword w of length at least $N/2$ of some $v_j(\underline{a})$.
2. For any graph that occurs in a folding sequence just before R_n a word of length 2 is not readable.
3. w_i is readable in any graph occuring in the folding sequence.

This gives a constradiction as w contains all subwords of length 2 with probabilty tending to 1 as N tends to infinity.

Thus for large N, we have $(a_1, \ldots, a_n) \not\sim (b_1, \ldots, b_n)$ with high probability.

Remark: An alternative argument for (2) can be given using the Whitehead graph.

Thus we have the following:

1. w_i contains a subword w of length at least $N/2$ of some $v_j(\underline{a})$.
2. For any graph that occurs in a folding sequence just before R_n a word of length 2 is not readable.
3. w_i is readable in any graph occuring in the folding sequence.

This gives a constradiction as w contains all subwords of length 2 with probabilty tending to 1 as N tends to infinity.

Thus for large N, we have $(a_1, \ldots, a_n) \not\sim (b_1, \ldots, b_n)$ with high probability.

Remark: An alternative argument for (2) can be given using the Whitehead graph.

Thus we have the following:

1. w_i contains a subword w of length at least $N/2$ of some $v_j(\underline{a})$.
2. For any graph that occurs in a folding sequence just before R_n a word of length 2 is not readable.
3. w_i is readable in any graph occuring in the folding sequence.

This gives a constradiction as w contains all subwords of length 2 with probabilty tending to 1 as N tends to infinity.

Thus for large N, we have $(a_1, \ldots, a_n) \not\sim (b_1, \ldots, b_n)$ with high probability.

Remark: An alternative argument for (2) can be given using the Whitehead graph.

Thus we have the following:

1. w_i contains a subword w of length at least $N/2$ of some $v_j(\underline{a})$.
2. For any graph that occurs in a folding sequence just before R_n a word of length 2 is not readable.
3. w_i is readable in any graph occuring in the folding sequence.

This gives a constradiction as w contains all subwords of length 2 with probabilty tending to 1 as N tends to infinity.

Thus for large N, we have $(a_1, \ldots, a_n) \not\sim (b_1, \ldots, b_n)$ with high probability.

Remark: An alternative argument for (2) can be given using the Whitehead graph.

The proof of the general case follows the same strategy. Thus assume that

$$(a_1, \ldots, a_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}}) \sim (b_1, \ldots, b_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}}).$$

It follows that $(a_1, \ldots, a_n, 1, \ldots, 1)$ is in $F(A)$ Nielsen equivalent to $(w_1, \ldots, w_{2n-1})$ such that $w_i =_G b_i$ for $1 \leq i \leq n$.

We can now construct Γ as before as the wedge of $2n-1$ circuits with labels $w_1, \ldots, w_{2n-1}$ and get π_1-surjective map to R_n which factors as a product of folds.

However at some point in the folding sequence a copy of R_n can occur as a subgraph which means everything is readable. We inspect the last graph in the folding sequence before this occurs.

The proof of the general case follows the same strategy. Thus assume that

$$(a_1, \dots, a_n, \underbrace{1, \dots, 1}_{n-1 \text{ times}}) \sim (b_1, \dots, b_n, \underbrace{1, \dots, 1}_{n-1 \text{ times}}).$$

It follows that $(a_1, \dots, a_n, 1, \dots, 1)$ is in $F(A)$ Nielsen equivalent to $(w_1, \dots, w_{2n-1})$ such that $w_i =_G b_i$ for $1 \leq i \leq n$.

We can now construct Γ as before as the wedge of $2n-1$ circuits with labels $w_1, \dots, w_{2n-1}$ and get π_1-surjective map to R_n which factors as a product of folds.

However at some point in the folding sequence a copy of R_n can occur as a subgraph which means everything is readable. We inspect the last graph in the folding sequence before this occurs.

The proof of the general case follows the same strategy. Thus assume that

$$(a_1, \ldots, a_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}}) \sim (b_1, \ldots, b_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}}).$$

It follows that $(a_1, \ldots, a_n, 1, \ldots, 1)$ is in $F(A)$ Nielsen equivalent to $(w_1, \ldots, w_{2n-1})$ such that $w_i =_G b_i$ for $1 \leq i \leq n$.

We can now construct Γ as before as the wedge of $2n-1$ circuits with labels $w_1, \ldots, w_{2n-1}$ and get π_1-surjective map to R_n which factors as a product of folds.

However at some point in the folding sequence a copy of R_n can occur as a subgraph which means everything is readable. We inspect the last graph in the folding sequence before this occurs.

The proof of the general case follows the same strategy. Thus assume that

$$(a_1, \ldots, a_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}}) \sim (b_1, \ldots, b_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}}).$$

It follows that $(a_1, \ldots, a_n, 1, \ldots, 1)$ is in $F(A)$ Nielsen equivalent to $(w_1, \ldots, w_{2n-1})$ such that $w_i =_G b_i$ for $1 \leq i \leq n$.

We can now construct Γ as before as the wedge of $2n-1$ circuits with labels $w_1, \ldots, w_{2n-1}$ and get π_1-surjective map to R_n which factors as a product of folds.

However at some point in the folding sequence a copy of R_n can occur as a subgraph which means everything is readable. We inspect the last graph in the folding sequence before this occurs.

Let Δ be the last graph in the folding sequence that does not contain a copy of R_n. Thus Δ contains a subgraph Ψ (fat edges) that folds onto R_n with a single fold.

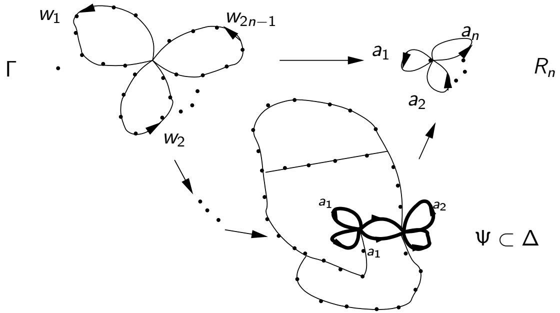

We will now explain how the w_i (or subwords that are also subwords of the v_j) can be read in Δ. Note that any word can be read in such a graph as we have no control over the part of Δ that is the complement of Ψ.

Let Δ be the last graph in the folding sequence that does not contain a copy of R_n. Thus Δ contains a subgraph Ψ (fat edges) that folds onto R_n with a single fold.

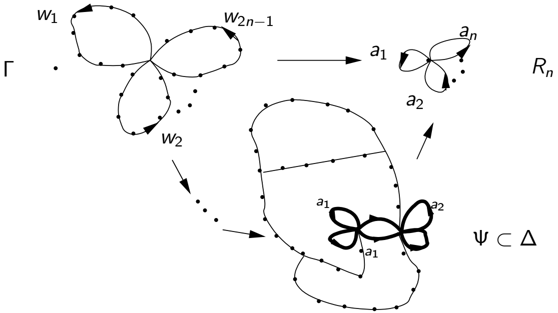

We will now explain how the w_i (or subwords that are also subwords of the v_j) can be read in Δ. Note that any word can be read in such a graph as we have no control over the part of Δ that is the complement of Ψ.

Note that the following hold for the finite labeled graph Δ:

1. $b(\Delta) \leq 2n - 1$
2. Δ does not contain a subgraph that covers R_n.

We call such a graph *tame*.

Theorem 3

Let Δ be tame. Then there is a word that is not readable in Δ.

Note that the following hold for the finite labeled graph Δ:

1. $b(\Delta) \leq 2n - 1$
2. Δ does not contain a subgraph that covers R_n.

We call such a graph *tame*.

Theorem 3

Let Δ be tame. Then there is a word that is not readable in Δ.

Note however that any word is readable in some tame graph. We will see that generically such words can only be read in a very controlled way.

Definition 4

Let $\alpha \in [0, 1]$. We say that a path γ in some graph Γ is α-injective if γ crosses at least $\alpha \cdot |\gamma|$ distinct topological edges.

Remark: 1-injective means no edge (pair) is travelled twice.

Theorem 5

Let $\alpha \in [0, 1)$. The set Ω of all reduced words in the $a_i^{\pm 1}$ contains a generic subset S such that the following holds:
Let $s \in S$ and Γ be a tame connected core graph. Then any path in Γ that reads s is α-injective.

Note however that any word is readable in some tame graph. We will see that generically such words can only be read in a very controlled way.

Definition 4

Let $\alpha \in [0,1]$. We say that a path γ in some graph Γ is α-injective if γ crosses at least $\alpha \cdot |\gamma|$ distinct topological edges.

Remark: 1-injective means no edge (pair) is travelled twice.

Theorem 5

Let $\alpha \in [0,1)$. The set Ω of all reduced words in the $a_i^{\pm 1}$ contains a generic subset S such that the following holds:
Let $s \in S$ and Γ be a tame connected core graph. Then any path in Γ that reads s is α-injective.

Note however that any word is readable in some tame graph. We will see that generically such words can only be read in a very controlled way.

Definition 4

Let $\alpha \in [0, 1]$. We say that a path γ in some graph Γ is α-injective if γ crosses at least $\alpha \cdot |\gamma|$ distinct topological edges.

Remark: 1-injective means no edge (pair) is travelled twice.

Theorem 5

Let $\alpha \in [0, 1)$. The set Ω of all reduced words in the $a_i^{\pm 1}$ contains a generic subset S such that the following holds:
Let $s \in S$ and Γ be a tame connected core graph. Then any path in Γ that reads s is α-injective.

This allows us to immediately deal with one more special case, namely the case that $w_i = v_i$ for $1 \leq i \leq n$. Recall that we have the following map:

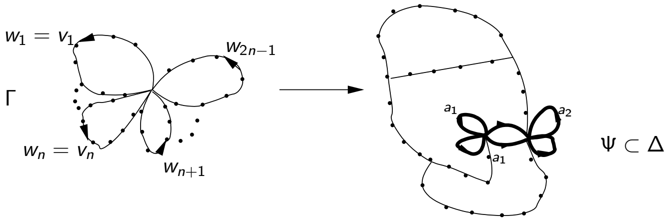

The above theorem implies that the $w_i = v_i$ are read by essentially injective paths. Moreover they are distinct generic words, thus for each v_i there is an arc in Δ outside Ψ that is only travelled once and not travelled by the other v_j. This implies that $b(\Delta) \geq 2n$ as $b(\Psi) = n$, a contradiction.

This allows us to immediately deal with one more special case, namely the case that $w_i = v_i$ for $1 \leq i \leq n$. Recall that we have the following map:

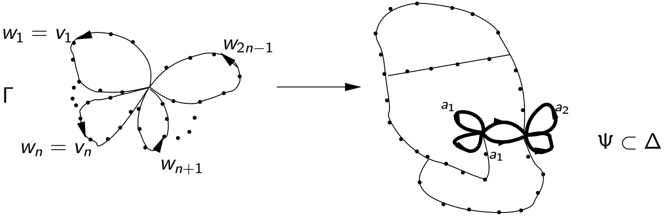

The above theorem implies that the $w_i = v_i$ are read by essentially injective paths. Moreover they are distinct generic words, thus for each v_i there is an arc in Δ outside Ψ that is only travelled once and not travelled by the other v_j. This implies that $b(\Delta) \geq 2n$ as $b(\Psi) = n$, a contradiction.

The remaining case is the case where $w_i \neq v_i$ for some $1 \leq i \leq n$. In this case we will see that we can find another graph Γ where the circuits are of shorter length.

This is achieved by surgery on Γ, a modification introduced by Arzhantseva and Olshankii. The length considered however is not the usual word length but significantly more subtle.

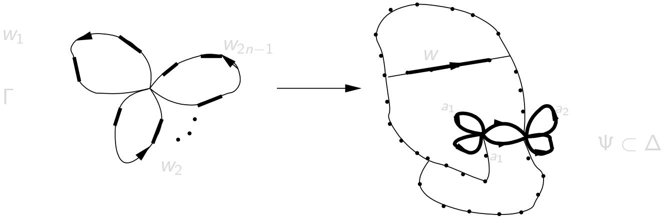

Figure : Replace arc with label w and all its preimages in Γ by arc with label $\bar{w}$ where $|\bar{w}| < |w|$ and $\bar{w} =_G w$.

The remaining case is the case where $w_i \neq v_i$ for some $1 \leq i \leq n$. In this case we will see that we can find another graph Γ where the circuits are of shorter length.

This is achieved by surgery on Γ, a modification introduced by Arzhantseva and Olshankii. The length considered however is not the usual word length but significantly more subtle.

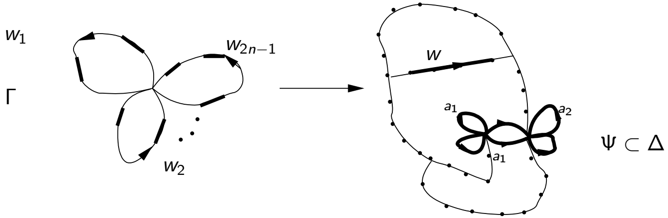

Figure : Replace arc with label w and all its preimages in Γ by arc with label $\bar{w}$ where $|\bar{w}| < |w|$ and $\bar{w} =_G w$.

This last case relies on the fact that the presentation

$$\langle a_1, \ldots, a_n | U_1, \ldots, U_n \rangle$$

obtained by Tietze transformation eliminating the b_i is again (an arbitrarily good) small cancellation group and that any reduced word representing b_i that is distinct from v_i must contain almost all of some cyclic conjugate of some U_j which must be read α-injectively in Δ for some α close to 1.