International Journal of Algebra and Computation
Vol. 13, No. 1 (2003) 61–68

World Scientific
www.worldscientific.com

BREADTH-FIRST SEARCH AND THE ANDREWS–CURTIS CONJECTURE

GEORGE HAVAS* and COLIN RAMSAY†

Centre for Discrete Mathematics and Computing
School of Information Technology and Electrical Engineering
The University of Queensland, Queensland 4072, Australia
**havas@itee.uq.edu.au*
†*cram@itee.uq.edu.au*

Received 26 June 2001

Communicated by M. F. Newman

Andrews and Curtis conjectured in 1965 that every balanced presentation of the trivial group can be transformed into a standard presentation by a finite sequence of elementary transformations. Recent computational work by Miasnikov and Myasnikov on this problem has been based on genetic algorithms. We show that a computational attack based on a breadth-first search of the tree of equivalent presentations is also viable, and seems to outperform that based on genetic algorithms. It allows us to extract shorter proofs (in some cases, provably shortest) and to consider the length thirteen case for two generators. We prove that, up to equivalence, there is a unique minimum potential counterexample.

Keywords: Andrews–Curtis conjecture; trivial group; group presentations; computer generated proofs.

AMS Mathematics Subject Classification: primary 20F05, 68Q40, 20-04; secondary 57M05

1. Introduction

There are infinitely many finite presentations of the trivial group. The problem of determining whether or not presentations are of the trivial group is, in general, unsolvable. A presentation $\langle x_1, \ldots, x_n \,|\, r_1, \ldots, r_m \rangle$ is called *balanced* if $m = n$. The *standard* n-generator presentation of the trivial group is $\langle x_1, \ldots, x_n \,|\, x_1, \ldots, x_n \rangle$. One method of proving that a balanced presentation represents the trivial group is to reduce it to a standard presentation.

Consider the following transformations on the n-tuple of words $(r_1, \ldots, r_n)$:

AC1) replace r_i by $r_i r_j$, for some $j \neq i$;
AC2) replace r_i by r_i^{-1};
AC3) replace r_i by $g^{-1} r_i g$, where g is a generator or its inverse.

Each of these *AC-moves* leaves r_k fixed for all $k \neq i$. Normally, conjugation by an arbitrary word is allowed in AC3, but we adopt a more restrictive definition to facilitate counting AC-moves; arbitrary conjugation is simply a succession of AC3 moves. Two presentations are said to be *AC-equivalent* if one can be transformed to the other by a sequence of AC-moves. Andrews and Curtis [2] conjectured that every balanced presentation of the trivial group is AC-equivalent to the standard presentation. A good group-theoretic survey of this problem is given by Burns and Macedońska [3], while the topological implications are discussed by Wright [11].

Although many presentations are known to satisfy the conjecture, little theoretical progress has been made on the general case. Miasnikov [8] describes a genetic algorithm designed to produce sequences of AC-moves which prove AC-equivalence. An alternative approach (which has so far had limited success) is based on coset enumeration and is described in [7]. Miasnikov and Myasnikov [9] have shown that the conjecture holds for all balanced presentations of the trivial group on two generators with the total length of the relators at most twelve. They also note that the presentation $\langle x, y \,|\, x^3 = y^4, xyx = yxy\rangle$ is currently a minimal potential counterexample.

Miasnikov notes that the number of sequences of AC-moves grows exponentially, and that exhaustive enumeration does not seem profitable. In this paper we show a contrary result, and that an appropriately written search routine easily copes with two-generator balanced presentations of length twelve; in fact, we readily produce shorter proofs than the ones given in [8]. In Sec. 2 we discuss our technique and illustrate its effectiveness with a simple example, while in Sec. 3 we revisit five difficult examples of lengths eleven and twelve, and extract shorter proofs. In Sec. 4 we consider the length thirteen presentations on two generators and show that the potential counterexample given above is unique, up to AC-equivalence.

Our test machine was a SUN Enterprise 450 system running SunOS 5.8 and clocked at 400 MHz. Our programming was done in C, and we used the `gcc` compiler, with level 3 optimization. The AC-move enumeration software (known as ACME), along with many utility programmes and sample input and output files, is available via the second author's web page at `http://www.itee.uq.edu./~cram`; the utilities and examples allow all the work discussed here to be reproduced. Throughout, we adopt the convention of using upper-case letters to denote inverses so that, for example, $A = a^{-1}$. Although ACME can handle larger values of n, we restrict ourselves here to the $n = 2$ case.

2. Technique

If we have a balanced presentation on n generators, then there are $3n^2$ AC-moves that can be applied to it. So there are $(3n^2)^k$ sequences of length k comprising AC-moves. However, we are really interested in the resulting presentations, not in the sequences of AC-moves. The count $(3n^2)^k$ is only an upper bound on the possible number of different presentations after k moves, and it is unduly pessimistic. To

see this, consider the following two remarks. Firstly, any AC-move can be undone by one or more AC-moves, so that a given presentation is repeatedly regenerated after its first appearance. Secondly, although AC-moves do not, in general, commute, it is often possible to rearrange a sequence of moves without affecting the final presentation; that is, different sequences of moves give the same presentation. The number of different presentations still grows exponentially with k, but, on average, each presentation has fewer than $3n^2$ children which are new.

It is not possible, in general, to extract a proof of AC-equivalence without allowing the length of the presentation to increase. Although any length increase needed is not known *a priori*, in many cases it turns out to be modest. Putting an upper bound on the presentation lengths during a search has a dramatic impact on the size of the search space — it makes it finite. Unconstrained searches, except in very small cases, typically fail by overflowing the available memory. A constrained search is the key to our success; although not exhaustive, it allows much longer sequences of moves to be investigated. Further, if such a search completes with no result, we can put a lower bound on any length increase necessary to extract a proof, assuming that a proof exists.

We have written a simple utility, ACME, which does a breadth-first search, in terms of AC-moves, through the tree of AC-equivalent presentations. Each presentation is recorded when it is first encountered, and is ignored if seen again. If requested, only presentations satisfying certain length constraints are recorded. We have found that limiting the total length of the presentation works well, as does limiting the lengths of the individual relators. We normally stop the search as soon as any presentation of total relator length n is found. This will be some permutation of the n generators or their inverses; if required, the search can be continued until all possible $n!2^n$ presentations of length n are derived. (If any one of these is derivable, then they all are.) In some cases it is much faster to stop when some length-reducing sequence of moves is found, and to restart the process with the shorter presentation.

To illustrate the efficacy of our technique, consider the simple example of length ten $\langle a, b \,|\, a^b = a^2, b^a = b^2 \rangle$ discussed in [8]. If we take the relators in the forms $BabAA$ and $AbaBB$ and attempt this with no constraints, then the number of presentations at levels 0–8 of the tree are 1, 12, 101, 782, 5906, 44048, 326200, 2406914 and 17724982, respectively. The initial presentation has $3 \times 2^2 = 12$ children, but thereafter the average numbers of children are 8.42, 7.74, 7.55, 7.46, 7.41, 7.38 and 7.36. We did not have enough memory to go further, but the output does tell us that no sequence of eight or less AC-moves can reduce the presentation as given to any presentation of length two.

If we now limit our search to presentations both of whose relators have length at most five, then we generate all eight presentations of length two, and complete the tree, in 1.5 seconds. Four of the proofs require ten moves, four require eleven moves, and the entire search tree contains only 31240 presentations. The shortest proofs which generate presentations with the generators uninverted have eleven moves ([8]

takes thirteen), and one example is: $(BabAA, AbaBB) \to (ABabA, AbaBB) \to (AABab, AbaBB) \to (AABab, BAbaB) \to (AB, BAbaB) \to (BA, BAbaB) \to (BA, bABab) \to (BA, bAB) \to (BA, baB) \to (BA, a) \to (B, a) \to (b, a)$.

Note that in this proof the first three moves can be regarded as *nonessential*, in the sense that they simply transform the relators to their cyclically permuted or inverted equivalents. These moves can be eliminated by starting with the presentation in the form $(AABab, BAbaB)$. Similarly, the final move can be regarded as nonessential. Pruning these moves leaves a proof of length seven ([8] takes ten). This proof is short enough so that we can conduct a search of the unconstrained tree. If this is done no proof with fewer AC-moves is found. In fact, exhaustive testing using all possible cyclically permuted or inverted forms of the relators (100 possibilities in all) establishes that a proof of (essential) length seven is the shortest possible.

3. Some Length 11 and 12 Examples Revisited

The family of presentations $\langle a, b \,|\, a^m = b^{m+1}, aba = bab\rangle$, $m \geq 2$, described by Akbulut and Kirby [1], is a well-known series of potential counterexamples to the Andrews–Curtis conjecture. The $m = 3$ case is the example quoted in the introductory remarks, while the $m = 2$ case, of length eleven, was solved by Miasnikov. He gave a proof chain of eighteen essential moves; this proof requires the presentation to grow to length twenty-five. Investigations using ACME establish that no proof is possible if the intermediate presentations are limited in length to at most twelve. If length thirteen is allowed, then a proof in nineteen essential moves is possible. If the length is allowed to grow to fifteen, then a proof in fourteen essential moves is possible: $(aBBBa, BAbabA) \to (aBBBa, BAbaBBa) \to (aBBBa, AbaBBaB) \to (aBBaBBaB, AbaBBaB) \to (aBBaBBaB, bAbbABa) \to (aBBaBBaB, abAb\,bAB) \to (aBBaBBaB, BabAbbA) \to (aBBaBBaB, BaBaB) \to (BBaBBaBa, BaBaB) \to (BBaBBaBa, bAbAb) \to (BBaBBaBa, AbAbb) \to (BBa, AbAbb) \to (BBa, Ab) \to (B, Ab) \to (B, A)$.

We cannot prove that there is no shorter proof directly, since the unconstrained search space is too large. However, we can do so by combining a search starting from $(aaBBB, abaBAB)$ and its equivalents with a *reverse* search starting from (a, b) and its equivalents. (For a reverse search, move AC1 is changed to "replace r_i by $r_i r_j^{-1}$".) When the two searches reach a common presentation, then we have a shortest proof. An unconstrained search based on this idea established that no proof is possible in less than fourteen essential moves; the search took 37 minutes of CPU time and used about 2 gigabytes of memory.

Miasnikov also considers four presentations in the family $\langle x, y \,|\, Xy^m x = y^{m+1}, x = w\rangle$, where $m \geq 1$ and w is a word in x and y with exponent sum zero in x, due to Miller and Schupp [10]. The four length twelve examples with $w = y^{\pm 1} x y^{\pm 1} X$ seem to be particularly difficult. He gives a proof in seventeen essential moves for one of them, including two Whitehead automorphisms which are not in AC1–AC3. Whitehead automorphisms in a successful reduction can be

replaced by AC-moves, albeit at the expense of (possibly) increasing the length of the proof. For the other three presentations, Miasnikov uses his genetic algorithm to extract proofs that they are AC-equivalent to the length eleven presentation considered earlier. Using ACME, we were able to extract direct proofs for all four examples. Two of the proofs had fourteen essential moves and the other two had twenty-three.

4. The Length 13 Presentations

We used ACME to verify the results reported by Miasnikov and Myasnikov and to extend them to the length thirteen case. Our investigations establish the following result.

Theorem. *Let G be a group defined by a balanced presentation on two generators, with the sum of the relator lengths at most thirteen. Then*:

(i) *if G has trivial abelianization, G is trivial or is isomorphic to $L_2(5)$, the unique perfect group of order* 120;
(ii) *if G is trivial, its presentation is AC-equivalent to the standard presentation or to the presentation $\langle x, y \,|\, x^3 = y^4, xyx = yxy\rangle$.*

The enumeration procedure we use for presentations is different from that of Miasnikov and Myasnikov, and yields smaller lists. We start by fixing on a set of generators $\{a, b\}$ and an ordering $a < A < b < B$, and a total relator length l. We now do the following steps.

(A) Enumerate all balanced presentations on the generator set where the sum of the relator lengths is l, and the presentations are in *canonic* form. By this we mean that: each generator (or its inverse) is in at least one of the relators; the relators are distinct and nonempty; the relators are freely and cyclically reduced; each relator is the lexicographic minimum of all its cyclic permutations and their inverses; the relators are in order (length plus lexicographic).
(B) Calculate the abelianized relators $(a^k b^m, a^n b^p)$ of each presentation, and discard those with $kp - mn \neq \pm 1$; that is, with nontrivial abelian quotient.
(C) Discard all presentations where one relator, perhaps cycled or inverted, is contained in another.
(D) Discard all presentations where one relator can be written as the relation $x_i = x_j^m$, where $x_i \in \{a, A\}$ and $x_j \in \{b, B\}$ (or vice versa).

It is straightforward to see that the list after step (A) contains at least one copy of each presentation of interest, and that the presentations subsequently discarded are either nontrivial, or AC-equivalent to the standard presentation or to a presentation from a list for a smaller l. The procedure eliminates all presentations with relators of length four or less, so the smallest l which yields a nonempty list is $l = 10$. The list sizes for $l = 10$, 11, 12 and 13 are 20, 104, 204 and 1400 respectively.

Note that, in general, lists of presentations can also be pruned by retaining only a single representative of each orbit under the action of the *extended* symmetric group of order $2^n n!$ generated by the permutations $(x_i, x_j)(X_i, X_j)$, for all $i \neq j$, and (x_i, X_i), for all i. It is clear that two distinct presentations in the same orbit are either both AC-equivalent to the standard presentation or both not AC-equivalent to it. For $n = 2$, this pruning decreases the list sizes by a factor of up to eight; for example, the list for $l = 10$ reduces from twenty to four. Our lists are short enough, and ACME is fast enough, to render this pruning unnecessary for the $l \leq 13$ cases we considered.

To prove Part (i) of our theorem the lists were processed by a coset enumeration programme (we used ACE, see [6]). All the enumerations are easy, with all the groups being trivial apart from four in the $l = 12$ list and twenty in the $l = 13$ list which have order 120. Under the action of the extended symmetric group these twenty-four fall into six orbits, and representatives are: $(a^2 = bab, b^4 = aba)$; $(a^2 = bab, a^3 = b^5)$; $(a^2 = bab, b^3 = abAba)$; $(aba = bAb, a^5 = b^2)$; $(aba = bab, a^3 = baab)$; $(aba = bab, a^2 = baBab)$.

To prove Part (ii) of our theorem the lists, minus the twenty-four presentations for nontrivial groups from Part (i), were processed by ACME. All presentations in the lists for $l < 13$ are easily reduced to the standard presentation. For the $l = 13$ case, an increase in relator length to eighteen allowed all but twelve presentations to be reduced to the standard one. The twelve remaining fall into two orbits under the action of the extended symmetric group, and representatives are: $(a^3 = b^4, aba = bab)$; $(ab = bbaa, ba = aaabb)$. We were unable to find any length decreasing chains of moves for these presentations, and an exhaustive search (using a machine with $\approx$8 gigabytes of memory) established that no such chain is possible if each relator is allowed to grow in length by at most five. However, ACME was able to show that these twelve presentations are pairwise AC-equivalent; for example, the first representative can be transformed to the second in thirty-two essential moves. The following proof, which seems way beyond what can be readily found by hand, is representative of what ACME can achieve.

$$
\begin{aligned}
&(aBBBBaa, AbabAB) \to (baBBBBaaB, AbabAB) \to (baBBBBaaB,\\
&\quad AbaBBBaaB) \to (bAAbbbbAB, AbaBBBaaB) \to (AbAAbbbbABa,\\
&\quad AbaBBBaaB) \to (AbAAbaaB, AbaBBBaaB) \to (AbAAbaaB,\\
&\quad AAbaBBBaaBa) \to (AbAAbaaB, AAbaBBaaB) \to (AbAAbaaB,\\
&\quad AAAbaBBaaBa) \to (AbAAbaaB, AAAbaBaaB) \to (AbAAbaaB,\\
&\quad AAAAbaBaaBa) \to (AbAAbaaB, AAAAbaaaB) \to (BAbAAbaa,\\
&\quad AAAAbaaaB) \to (ABAbAAbaaa, AAAAbaaaB) \to (bABAbAAbaaaB,\\
&\quad AAAAbaaaB) \to (bABAbAAbaaaB, bAAAABaaaa) \to (bABAbaa,\\
&\quad bAAAABaaaa) \to (AbABAbaaa, bAAAABaaaa) \to (bAbABAbaaaB,
\end{aligned}
$$

$$bAAABaaaa) \to (bAbABaaa, bAAABaaaa) \to (AbAbABaaaa,$$
$$bAAABaaaa) \to (AbAbABaaaa, AAAAbaaaB) \to (AbAbaaB,$$
$$AAAAbaaaB) \to (BAbAbaa, AAAAbaaaB) \to (ABAbAbaaa,$$
$$AAAAbaaaB) \to (bABAbAbaaaB, AAAAbaaaB) \to (bABAbAbaaaB,$$
$$bAAABaaaa) \to (bABAbaaa, bAAABaaaa) \to (bbABAbaaaB,$$
$$bAAABaaaa) \to (bbABaaa, bAAABaaaa) \to (AAAbaBB,$$
$$bAAABaaaa) \to (AAAbaBB, abAAABaaa) \to (AAAbaBB, abAABB)\,.$$

5. Concluding Remarks

We have shown that an intelligently constrained brute-force enumeration of AC-move sequences is a useful technique for investigating the Andrews–Curtis conjecture. It produced better results than those reported for an approach based on genetic algorithms, and enabled us to extend those results. However, we have been unable to make any progress on the smallest potential counterexample, beyond showing its uniqueness up to AC-equivalence. Although ACME is capable of investigating long chains of AC-moves, it still has an exponentially growing search-space. Within this search-space there is no biasing either for or against any particular move. The fitness function used in a genetic algorithm is an attempt to favour moves which make progress towards a solution. Proofs extracted using ACME exhibit no obvious pattern, and the length of a presentation during a proof can increase and decrease many times. This goes some way towards explaining the disappointing performance of genetic algorithms, but gives little guidance as to how to modify either approach to improve their performance.

Our current version of ACME is rudimentary, and the technique is capable of considerable elaboration. The primary constraint is memory usage, and future versions will include a more space and time efficient data storage method (hash tables as opposed to trees). A more dramatic space saving, at the expense of running time, could be achieved by storing only a representative of each equivalence class of presentations instead of the full class.

As well as allowing us to demonstrate AC-equivalence, an enumeration of AC-moves allows us to investigate questions regarding the shortest or most elegant proofs. For example, as well as being a shortest proof, our proof for $(aaBBB, abaBAB)$ in Sec. 3 is not without elegance; there is no obvious motivation for many of the moves, and yet the final four length-reducing moves are very efficient and result in complete cancellation of the appended relators. Questions regarding the length and elegance of proofs are difficult in general. It is interesting to note that Hilbert apparently considered including as a 24th problem that of "finding criteria for finding simplest proofs". This problem was ultimately omitted since "simplicity is an extremely *complicated* notion" [5].

Another approach to proving AC-equivalence is introduced in [3] and amplified in [4]. This is based on "elementary M-transformations". The authors of [4] introduce the notion of *recalcitrance* which is in effect the least number of elementary M-transformations required to prove AC-equivalence. Elementary M-transformations do not seem easily amenable to computer investigation, in part because at each stage there are an infinite number of these to be considered. It is interesting to note that the presentation $\langle a, b \,|\, a^2 = b^3, aba = bab \rangle$ which requires fourteen essential AC-moves (see Sec. 3) is shown to have recalcitrance 2 in [4]. In this sense elementary M-transformations seem much more efficient than AC-moves for establishing AC-equivalence by hand.

Acknowledgments

This research was partially supported by the Australian Research Council.

References

1. S. Akbulut and R. Kirby, A potential smooth counterexample in dimension 4 to the Poincaré conjecture, the Schoenflies conjecture and the Andrews–Curtis conjecture, *Topology* **24** (1985), 375–390.
2. J. J. Andrews and M. L. Curtis, Free groups and handlebodies, *Proc. Amer. Math. Soc.* **16** (1965), 192–195.
3. R. G. Burns and O. Macedońska, Balanced presentations of the trivial group, *Bull. London Math. Soc.* **25** (1993), 513–526.
4. R. G. Burns, W. N. Herfort, S.-M. Kam, O. Macedońska and P. A. Zalesskii, Recalcitrance in groups, *Bull. Austral. Math. Soc.* **60** (1999), 245–251.
5. I. Grattan-Guinness, Response to Moore's letter, *Notices Amer. Math. Soc.* **48** (2001), 167.
6. G. Havas and C. Ramsay, Coset enumeration: ACE version 3.001 (2001). Available as `http://www.itee.uq.edu.au/~havas/ace3001.tar.gz`.
7. G. Havas and C. Ramsay, Andrews–Curtis and Todd–Coxeter proof words, *Group St. Andrews 2001 in Oxford*, London Mathematical Society Lecture Note Series, Cambridge University Press (to appear).
8. A. D. Miasnikov, Genetic algorithms and the Andrews–Curtis conjecture, *Int. J. Algebra Comput.* **9** (1999), 671–686.
9. A. D. Miasnikov and A. G. Myasnikov, Balanced presentations of the trivial group on two generators and the Andrews–Curtis conjecture, eds. W. M. Kantor and A. Seress, *Groups and Computation III*, Ohio State Math. Res. Inst. Publ. 8, de Gruyter, **8**, Berlin, 2001, pp. 257–263.
10. C. F. Miller III and P. E. Schupp, Some presentations of the trivial group. ed. R. H. Gilman, *Groups, Languages and Geometry*, *Contemp. Math.* Vol. 250, Amer. Math. Soc., Providence, RI, 1999, pp. 113–115.
11. P. Wright, Group presentations and formal deformations, *Trans. Amer. Math. Soc.* **208** (1975), 161–169.