Solving word problems via generalisations of small cancellation

A Track Record

A.1 The Research Team

Prof. Stephen Linton is a Professor of Computer Science at the University of St Andrews. He has worked in computational algebra since 1986 and has coordinated the development of the GAP system [6] since its transfer from Aachen in 1997 (more recently, in cooperation with three other centres). He has been Director of the Centre for Interdisciplinary Research in Computational Algebra (CIRCA) since its inception in 2000. His work includes new algorithms for algebraic problems, most relevantly work on generalisations of the Todd-Coxeter algorithm, novel techniques for combining these algorithms into the GAP system, and their application to research problems in mathematics and computer science. He is an editor of the international journal *Applicable Algebra and Error Correcting Codes*. He has been principal investigator on five large EPSRC grants, including EP/C523229 "Multidisciplinary Critical Mass in Computational Algebra and Applications" (£1.1m) and EP/G055181 "High Performance Computational Algebra" (£1.5m across four sites). He is the Coordinator of the EU FP6 SCIEnce project (RII3-CT-026133, 3.2m Euros), developing symbolic computation software as research infrastructure. He won Best Paper award at the prestigious IEEE Visualisation conference in 2003 for a paper applying Computational Group Theory to the analysis of a family of algorithms in computer graphics.

Dr. Max Neunhöffer is currently a Senior Research Fellow in the School of Mathematics and Statistics at St Andrews, working on EP/C523229 "Multidisciplinary Critical Mass in Computational Algebra", and will advance to a Lectureship in Mathematics in September 2010. He has been at St Andrews since 2007, and before that worked for 10 years at Lehrstuhl D für Mathematik at RWTH Aachen, Germany. He has been involved in the development of GAP since the 1990s, both developing the core system and writing package code. His main mathematical area is the representation theory of groups and finite-dimensional algebras, with a particular emphasis on computational methods. He is highly experienced in connecting abstract mathematical theory with actual computer computations to make both areas benefit from one other: the project will benefit substantially from this special expertise. Max Neunhöffer has 14 mathematical publications and many refereed software publications, including 8 GAP packages.

Richard Parker is a freelance computer consultant who has maintained a lifelong involvement in computational algebra and representation theory, occasionally as a contract researcher at the universities of Cambridge and St Andrews, but more often as an independent researcher. Mathematically, he is best known as an author of the landmark *Atlas of Finite Groups* (over 2000 citations on Google Scholar). His contributions to the Atlas (and its successors, the *Atlas of Brauer Characters* and the on-line *Atlas of Group Representations*) were mainly computational, initially checking and correcting the character tables, later pioneering the solution of representation-theoretic problems by computing with matrices over finite fields. To this end he developed a body of algorithms (and several implementations of them) known as the "MeatAxe", which was used to find the modular character tables of many finite groups (including the long-standing "last character of M_{24}" from Prof. G.D. James' PhD thesis). His work was also central to the existence proof of the sporadic group J_4, in the process originating many of the methods which are now central to the "Matrix Group Recognition Project". After a long string of results on modular character tables of particular groups (mainly in the *Modular Atlas Project*) Richard Parker started working with the character tables themselves on the computer, requiring the development of new techniques in exact integer linear algebra ("the Integral MeatAxe") now used in GAP and elsewhere. His other work has ranged widely across algebra and combinatorics, including determining the covering radius of the Leech lattice and introducing (with Prof. D.J. Benson) the concept of Brauer species, which generalise Brauer characters. More recently, he has been working for several years on the preliminary investigations which allow us to formulate this project.

Dr. Colva Roney-Dougal is a Lecturer in Pure Mathematics at St Andrews, appointed in September 2005. Previously she was a Postdoctoral Research Fellow in the School of Computer Science at St Andrews, working with Prof. Linton and Prof. Ian Gent, and before that she was a Postdoctoral Researcher with the Computational Algebra group at the University of Sydney. She has 25 publications, including many on matrix group algorithms, and in particular she has developed many group-theoretic algorithms which exploit the geometry (in the sense of Aschbacher) preserved by the matrix group. Another key theme of her research is primitive permutation groups, where she has classified the primitive permutation groups of degree up to 4095, using a mixture of theoretical and computational techniques. She also has many publications in constraint satisfaction, an area of theoretical computer science which studies search algorithms. Since 2008 she has been an editorial advisor for the *London Mathematical Society Journal of Computation and Mathematics*. She has appeared on national radio programmes ("In Our Time" on Radio 4 and "Next generation thinkers" on Radio 3) which popularise mathematics. She was recently awarded a £24K grant by the International Centre for Mathematical Sciences in Edinburgh to run a workshop on the Matrix Group Recognition project, and also received a Newly Appointed

Lecturers grant from the Nuffield Foundation, along with numerous smaller grants (<£6K).

A.2 Expertise and Collaborations at Host Organisations

Our research is based in the Centre for Interdisciplinary Research in Computational Algebra (CIRCA), a thriving interdisciplinary research centre, which brings together researchers in algebra, combinatorics, and algorithms. The Centre was established in 2000 to foster new and exciting collaborative research between the Schools of Mathematics and Statistics and of Computer Science. It has dedicated accommodation comprising offices, computer area, library and social area. This provides a friendly, informal environment for interactions among the staff, students and visitors. For more formal discussions of daily business and academic issues, there are regular meetings of the Centre. More information about CIRCA, including its technical reports series, can be obtained from `www-circa.mcs.st-and.ac.uk`.

Through contacts in the two Schools that CIRCA spans (Mathematics and Statistics, and Computer Science) we have access to expertise in a wide range of areas relevant to this proposal; we summarise a few here:

- The algebra group in St Andrews has been interested in combinatorial group theory since at least the 1970s, when pioneering work was carried out by Professors Campbell and Robertson. Although retired, they are both still extremely active mathematically, and are in the department most days.
- Developing from this, the algebra group has long been interested in questions of decidability in groups and semigroups, for example Professor Ruskuc and Dr Quick recently were awarded £350K by EPSRC for a grant entitled "Automata, Languages, Decidability in Algebra", which has led to the appointment of Dr Collin Bleak as a Postdoctoral Research Fellow for three years (commencing June 2010).
- There is also extensive experience in working with finite state automata, both in the sense of automatic groups and also uses of automata in combinatorics.
- A major theme in the Computer Science branch of CIRCA is Constraint Satisfaction, which essentially is the study of search techniques; we expect search to play an important role when developing complete sets of rules for assigning curvature.
- CIRCA has an outstanding track record of major open-source software projects. For example the recent "HPC GAP" EPSRC grant of £1.5m (Principal Investigator: Professor Linton) seeks to develop GAP to exploit the full range of parallel computers, from multicore desktops to national supercomputers.

A.3 National and International Collaborations

CIRCA is extremely well-connected in a wide range of research communities, including abstract algebra, computational algebra, combinatorics, constraint programming, symbolic computation and parallel computing. These connections are maintained by a strong network of joint projects, visits, staff who have studied or worked in other centres, and collaboration in conference organisation, software development and maintenance, and other activities. It would not be feasible, or useful to list them all here. Instead we mention only our main formal collaborations relevant to this project: with the other Centres of GAP development in Aachen, Braunschweig and Fort Collins; with the Centre for Algebra of the University of Lisbon – a centre of expertise in semigroup theory; with Professor Derek Holt at Warwick, one of the founders of the theory of automatic groups; and with Dr George Havas, at Brisbane, a world expert on coset enumeration.

A.4 Selected Bibliography

We finish with a short selection of our recent and relevant publications:

1. Carlson, J. F., Neunhöffer, M., Roney-Dougal, C.M. A polynomial-time reduction algorithm for groups of semilinear or subfield class. *J. Algebra* 322 (2009) 613–637.
2. Holmes, P. E., Linton, S. A., O'Brien, E. A., Ryba, A. J. E., Wilson, R. A. Constructive membership in black-box groups. *J. Group Theory* 11 (2008), no. 6, 747–763.
3. Holt, D.F., Roney-Dougal, C.M. Constructing maximal subgroups of classical groups. *LMS J. Comput. Math.* 8 (2005), 46–79.
4. Holt, D.F., Roney-Dougal, C.M. Constructing maximal subgroups of orthogonal groups. *LMS J. Comput. Math.* To appear.
5. Linton, S.A. Double coset enumeration. *J. Symbolic Comput.* 12 (1991), no. 4-5, 415–426.
6. Linton, S.A. Generalisations of the Todd-Coxeter algorithm. Computational algebra and number theory (Sydney, 1992), 29–51, Kluwer Acad. Publ., Dordrecht, 1995.
7. Linton, S.A., Parker, R.A., Walsh, P.G., Wilson R.A. Computer construction of the Monster, *J. Group Theory* 1 (1998), 307-337.
8. Neunhöffer, M., Scherotzke, S. Formulas for primitive idempotents in Frobenius algebras and an application to decomposition maps, *Representation Theory* 12 (2008), 170–185.
9. Wilson, R.A., *et al* (includes Linton and Parker). ATLAS of Finite Group Representations - Version 3, `http://brauer.maths.qmul.ac.uk/Atlas/v3/`

B Background

It is well known that the word problem for finitely-presented groups is algorithmically unsolvable in general. Nevertheless, practical procedures that attempt to solve it in particular instances have been a central element in computational group theory since its inception in the 1930s. Most available methods are variations on three themes: coset enumeration, following Todd and Coxeter; string rewriting, following Knuth and Bendix; and the construction of automatic structures, following Epstein, Holt and others [3]. Other approaches investigate finitely-presented groups by way of quotient groups of various kinds, but do not provide a solution for the word problem. If the word problem can be solved for such groups, then it becomes possible to test whether the quotient is in fact isomorphic to the original group, enabling one (for example) to upgrade the nilpotent quotient algorithm to a method which will prove or disprove the nilpotency of the input group.

A key tool will be *van Kampen diagrams*, which encode derivations of new relators from given ones. Briefly, every rule $w = 1$ that holds in the group presentation $\langle X|R\rangle$ is the boundary of a van Kampen diagram, which consists of vertices, edges and regions (faces), where the edges are labelled with elements of the free group on X. The product of the labels around each region is freely reduced without cancellation, and for internal regions is equal to a cyclically-reduced conjugate of a relator (or of the inverse of a relator).

The word problem is solvable for a group G with given finite presentation if and only if G satisfies a recursive *isoperimetric inequality*: every equality $w = 1$ in G has a van Kampen diagram proving it with at most $f(|w|)$ regions, where f is a recursive function [7]. Although the word problem is solvable for such groups, there do not always exist effective algorithms to solve it in reasonable time. We will be generalising methods that apply to groups satisfying linear and quadratic isoperimetric inequalities, and seeking fast, efficient solutions to the word problem. The (minimal) exponent of the isoperimetric inequality depends only on the isomorphism type of G, and any group satisfying a sub-quadratic isoperimetric inequality in fact satisfies a linear one (see [10]). A key result is that groups satisfying linear isoperimetric inequalities are *hyperbolic* or *negatively curved* (see [9]), while "non-positively curved" groups satisfy a quadratic isoperimetric inequality. There is, however, no effective algorithm to decide if the group defined by a given finite presentation falls into either of these classes.

Hyperbolic groups are automatic (although not all non-positively curved groups are), but the automatic structures may be very large or hard to find. So, if the given presentation defines a hyperbolic group, the algorithm of [3] will, in principle, show this, but in practice it may require enormous time or resource. It is not known whether hyperbolic groups all have generating sets on which, even in principle, the Knuth-Bendix string rewriting algorithm will terminate, but there are known examples where it will not terminate with the given generators.

An alternative approach to proving that a finitely-presented group satisfies a linear or quadratic isoperimetric inequality is to check whether the presentation satisfies a *small cancellation condition.* This approach derives from seminal work of Lyndon and Schupp [16]. Small cancellation conditions ensure that all regions of any nontrivial van Kampen diagram have a boundary with at least n edges – namely condition $C'(1/(n-1))$ (metric) or $C(n)$ (non-metric) – and such that all vertices have degree at least some fixed k: condition $T(k)$. If these conditions hold for certain combinations of n and k then the presentation defines a group which satisfies a given quadratic isoperimetric inequality, and therefore has solvable word problem. More general conditions such as $V(6)$ and W look at larger subdiagrams: see for example [11, 20].

Our main aim is to massively extend small cancellation in a number of directions. Firstly we intend to define much more flexible sets of conditions, involving large subdiagrams, that still imply global curvature properties of the group. Indeed, we expect our software to automatically develop and validate such sets of conditions on a *per group* basis, using our idea of *curvature redistribution*, described in the second paragraph of Section D.2 below. Secondly, if the conditions are not met, we will develop ways of locating key obstacles in the form of large van Kampen diagrams with short boundary, or "Short Theorems with Long Proofs" (abbreviated STLP hereafter). These may be of independent interest, or may be added as extra relations to give a possibly more tractable presentation of the same group. This addition of relations is actually just one example of our third new idea, which is that, if a presentation is proving intractable, it may be possible to transform or embed it into a more tractable one, possibly of a significantly different kind (for instance, a quotient of a free product of finite and free groups). Finally, we observe that the possibility of computationally exploiting van Kampen diagrams and their geometry is not limited to the context of finitely-presented groups. See [12] for an introduction to *small overlap monoids* (and recent work in [14]); small cancellation theory also makes sense for HNN extensions of groups [16], cancellation semigroups and many other settings. More speculatively, we aim ultimately to consider algebraic structures that are not given by generators and relations, but do have some notion of element length that could give rise to geometry, generalising the work of [13].

Once we have shown that a group satisfies an isoperimetric inequality and so has solvable word problem, there remains the question of finding an *efficient* solution. An algorithm solving the word problem is *Dehn-type* if there exists a finite set of pairs (w_i, w_i') of words with $|w_i| \geq |w_i'|$ and $w_i(w_i')^{-1} = 1$ in G, such that if $w = 1$ in G then w_i is a subword of w for some i. If $|w_i| > |w_i'|$ for all i then this is Dehn's algorithm and executes in time linear in $|w|$, given the pairs (w_i, w_i'). Condition $C'(1/6)$ implies the applicability of Dehn's algorithm using the set of pairs (w, w'), where ww'^{-1} is obtained from a relator by cycling and cancellation, but other conditions are in general less helpful. The literature often gives no algorithm more efficient than simply checking all van Kampen diagrams with at most the number of regions given by the isoperimetric inequality, which is completely infeasible for larger problems. We plan, at least initially, to focus on *VR systems* [15]: λ-confluent, but not necessarily terminating, string rewriting systems with rules $w_i \to w_i'$ in which $|w_i| \geq |w_i'|$. Such systems exist, for example for groups satisfying $C'(1/3)$–$T(4)$ to which Dehn's algorithm need not apply.

Our preliminary investigations show that this is a large and worthwhile area for investigation, with the potential to transform our view of algorithms for finitely-presented groups. By explicitly considering whether the relators of such groups can be fitted together to make a highly curved area, we go directly to the heart of the problems addressed less directly by existing algorithms. For example, our curvature conditions will often be sufficient to prove that the input group is infinite, and may prove (or disprove) hyperbolicity. Conversely, our STLPs can be used as input for coset enumeration algorithms, and should help them to run much faster.

C Research Hypothesis and Objectives

We have the following main aims:

AIM 1 Explore, through theory, algorithms and software, the idea of redistributing curvature in minimal van Kampen diagrams as a route to solving word problems.

AIM 2 Extend the range of algebraic structures and representations to which both the current and our new methods of small cancellation apply, for example to cancellation semigroups and to matrix groups.

AIM 3 Develop practical algorithms and releasable software based on our work in pursuit of the other aims.

These aims are supported by the following objectives:

1.1. Develop algorithms and supporting theory which use the idea of curvature redistribution to analyse groups and show that their minimal van Kampen diagrams cannot have large areas of positive curvature.

1.2. Explore the relationship between the curvature properties of van Kampen diagrams and the types of word problem solvers that may be applied to the corresponding algebraic object.

1.3. Investigate ways of mapping algebraic and combinatorial structures which do not have nice geometric properties into ones that do, for example by modifying the input presentation or viewing a group as a quotient of a free product of free and finite groups.

2.1. Create a uniform and computational small cancellation theory for a wide class of algebraic structures, including groups, semigroups, monoids and cancellation semigroups.

2.2. Generalise the notion of "diagrams" from the classical van Kampen diagrams to algebraic structures that are not represented by generators and relations, but do have a suitable notion of element length.

3.1 Develop and release implementations of practical use for computational algebraists and others.

The computational extension of small cancellation theory is an entirely **novel** concept, and will dramatically extend both the existing theory of finitely-presented groups and related structures, and the availability of practical algorithms to work with them. In-depth investigation of this concept is **timely** for at least two reasons. Firstly, preliminary work by the investigators, over several years, has defined the problems and concepts sufficiently to permit a focused and structured investigation. For example, we have recently completed a prototype program (see Subsection D.2). Secondly, the recent award of the Abel prize to Gromov has sparked an upsurge of interest in geometric group theory: we are well-positioned to take advantage of this.

D Programme and Methodology

D.1 Methodology

Our methodology will closely integrate theoretical and practical investigations. Theoretical developments suggest algorithmic approaches, which we can test quickly using rapidly-developed prototypes in GAP, or explore more thoroughly (and on larger examples) using highly efficient C implementations. Experiments will suggest both algorithmic improvements and further theoretical insights. Algorithms (for instance to search effectively for specific kinds of van Kampen diagrams up to some size limit) are also of independent interest and, in these cases, we will investigate both practical performance and provable asymptotic complexity.

Once our software reaches the point of being useful for research problems, we will make it available to the community for use, testing and feedback under a free software license. We will also take advantage of existing free software, where applicable, minimising the amount of routine support programming required. For instance, we can see applications for the GNU linear programming toolkit in optimising curvature redistribution schemes.

One of our goals is to compute with very large presentations, with perhaps millions of generators. No currently available computational tools can deal with such presentations, however they can arise naturally in topology, for example in the computation of the fundamental group of a finite simplicial complex. They also frequently arise when applying the Reidemeister–Schreier method to obtain a presentation of a finite index subgroup. To quantify some of our goals, we briefly recall the asymptotic properties of these groups. Fix integers n and l, fix $\delta \in [0, 1]$, and pick uniformly at random a set of $(2n-1)^{\delta l}$ cyclically reduced relators of length l on n generators. Letting $l \to \infty$, Gromov proved that if $\delta < 1/2$ then the resulting presentation almost surely defines a hyperbolic group, whilst if $\delta > 1/2$ then the presentation almost surely defines the trivial group or C_2. Similar statements apply when the length of the relators is fixed, but their number tends to infinity. For $\delta < 1/5$ Ollivier [18] has shown that such presentations almost surely satisfy small cancellation conditions. For tiny δ, all sufficiently large presentations are deficient, and so present infinite groups for which Knuth–Bendix is efficient. We expect our techniques to work for random presentations with larger δ than this, and one measure of our success will be how much larger.

We will use a broad and heterogeneous set of benchmarks, including:

1. Random presentations with up to 10^7 generators, for values of δ such that the groups are almost surely hyperbolic.
2. Large but not complicated presentations of the trivial group.
3. One-relator quotients of the modular group $C_2 * C_3$ (see [4] and the references therein).
4. Hand-constructed examples that have arisen in other contexts.
5. Monoid and semigroup variants of the above types of presentation, as appropriate.

The performance of our programs on these examples will be one measure of our success, and, when examined in detail, also a source of insight which will inform subsequent developments.

D.2 Programme of Work

The heart of our work in this project is to develop, prove correct and study, software that attempts to do one of two things to an algebraic structure. Either it tries to construct a solver for the associated word problem, and prove its correctness, or it tries to identify "interesting" STLPs: words that are equal to the identity in the given structure, but cannot be shown to be so by any simple word problem solver. We call this "the main program" although it might in fact consist of a collection of related programs. The currently envisaged architecture of the system essentially divides this process into two phases: two dimensional and one dimensional.

The two-dimensional phase searches for a proof that the given structure is non-positively curved, thereby showing in particular that the word problem is solvable. These proofs will be of a particular form, generalising the basic proofs of classical small cancellation theory and following ultimately from Euler's formula $V - E + F = 2$. We consider each internal edge in a (minimal) van Kampen diagram to be initially endowed with one unit of *divergence* (negative curvature); our proof must give rules for the allocation of this divergence to the vertices and regions, such that all vertices and regions far enough from the boundary receive at least one unit. This shows that the contribution of the interior of the diagram to the left-hand side of Euler's equation cannot be positive, showing that all of the curvature of any (minimal) van Kampen diagram must be found near the boundary. Initially, at least, we will consider finite sets of rules, each of which applies to all edges whose neighbourhoods meet certain conditions, and allocates curvature from the edge to various vertices and regions within a bounded distance. To prove the validity of such a set of rules, we must find an unavoidable set of vertex and region configurations that can be checked. Very considerable care is needed to avoid double-counting. We say that a presentation for which such a curvature redistribution scheme exists satisfies a *non-homogeneous small cancellation condition.*

As a simple example, suppose we have a group presentation satisfying $C'(1/6)$ (and $T(3)$, which we assume throughout). Then the rule we use is that each edge allocates $1/6$ unit of divergence to the region on each side, and $1/3$ unit to the vertex at each end. A generic vertex thus receives at least one unit of divergence, and a generic region at least $7/6$ units, so the group is, in fact, negatively curved. More complex conditions such as W and $V(6)$ (see [11, 20]) can be validated by sets of rules in which edges behave differently depending on the nature of their neighbourhoods.

If this stage of the procedure fails, then it must be because we have been able to construct large subdiagrams of overall positive curvature, which serve as starting points for building STLPs.

If two-dimensional analysis succeeds, we now move to the one-dimensional phase, considering possible boundary configurations, aided by the fact that the boundary of any van Kampen diagram must contribute positive curvature. The normal aim of this analysis is to build up a collection of rewrite rules such that, in any relator proved by a van Kampen diagram, one of the rewrite rules applies to the relator, and either shortens it, or preserves its length while reducing the area of the diagram. This gives a concrete and usable word problem solver for the structure (some additional, finite, technical checks on short words are needed to deal with pathological conditions, such as a single region being on the boundary in two widely-separated places).

If this stage of the procedure fails, it can only be because we have actually constructed a long positively-curved boundary, which we can then try to complete to an STLP.

Prototype program: We have already developed a proof-of-concept implementation of an algorithm to modify a presentation and hence find a word-problem solver. This algorithm applies only to the case where all relators have length 3: any presentation can be modified to make this the case. The prototype tries to modify the presentation such that every vertex in a minimal van Kampen diagram can be shown to have degree six or more, a particularly straightforward approach. It seems that such a program can be very fast with large presentations, producing a word-problem solver for presentations with a million generators and two million random relators in minutes.

Based on this, and other, work we have developed a fairly detailed understanding of the operation of this proposed main program, and ideas for many of the necessary algorithms and data structures.

D.2.1 Work Programme for Aim 1

The first two linked workpackages cover the development and study of an initial version of the main program, targeted at finitely-presented groups and addressing Objectives 1.1 and 1.2. WP 1.1 will focus on the software itself and the supporting search algorithms and data structures, and WP 1.2 on related theoretical developments: clarifying the relationship between non-homogeneous small cancellation conditions and the applicability of different types of word problem solvers such as VR systems; and the relationship between the class of such groups and automatic groups. VR systems naturally generalise Dehn's algorithm, while classical work of Gersten and Short [8] shows that small cancellation groups are automatic. This latter result has recently been generalised to certain non-homogeneous small cancellation conditions [20], but it is possible that not all such conditions imply automaticity.

Classical small cancellation theory deals with *reduced* van Kampen diagrams in which no two adjacent regions are labelled with (conjugates of) mutually inverse relations. Clearly, from the perspective of solving the word problem, this notion can be strengthened to *minimal* van Kampen diagrams – those with the smallest possible number of regions for a given boundary (or even the smallest weight, where different types of region might have different non-negative weights). This allows us to exclude from consideration any subdiagram whose boundary can be filled with fewer regions (resp. regions of less total weight). Equivalently, we can forbid diagrams which contain more than half of any *sphere* (a reduced van Kampen diagram with no boundary), reducing the search spaces of several of our algorithms. It is thus very useful to obtain a theoretical and a practical understanding of the set of spheres allowed by a presentation. In the literature the classical theoretical way to understand spheres is via second homotopy groups and crossed modules: see [16] for an introduction; the research question here is to develop a practical, algorithmic approach. Two further linked workpackages will investigate these theoretical questions (WP 1.3) and further develop the main program to efficiently discover and take advantage of these spheres (WP 1.4). This will build on WP 1.1, and in particular on techniques for finding STLPs if one- or two-dimensional analysis fails.

Our prototype (which works only with relators of length 3) addresses Objective 1.3 in two ways. Firstly, it eliminates vertices of degree two by deducing equalities between generators. Secondly, it introduces new generators and relations. In doing so it determines a set of spheres which define positively-curved subdiagrams that cannot occur in minimal proofs, so that if the process terminates then all minimal proofs satisfy $T(6)$. The next workpackage (WP 1.5) is to further investigate the behaviour of these processes, and determine when (if ever) termination results can be proved.

A final workpackage WP 1.6 covers a range of more speculative ideas. Final decisions about the division of effort between these tasks will depend on current developments at that point.

- **Conjugacy and Isomorphism:** In conjugacy, similar notions of curvature apply but the diagrams are now annular rather than simply-connected. The isomorphism problem is considerably harder, however there are a range of results applying to hyperbolic groups [1, 2, 19].
- **Unbounded curvature redistribution schemes:** In certain groups, it may be necessary to redistribute curvature arbitrarily far. To handle this, we will need to extend our finite sets of rules to infinite – initially regular – languages of rules. This may further enlarge the class of groups that we are considering.
- **Torsion:** Torsion elements in $C'(1/6)$ small cancellation groups are "visible" in the presentation. We will determine when this generalises to non-homogeneous small cancellation.
- **Generalised word problem:** This has been helpful to other algorithmic approaches.
- **Constant of hyperbolicity:** It would be extremely useful to be able to bound the thickness of geodesic triangles in groups that are known to be hyperbolic.

D.2.2 Work programme for Aim 2

The investigations described in the previous section are concerned with small cancellation quotients of free groups. However, there are many other algebraic structures for which a theory of small cancellation quotients has been developed: free products (with or without amalgamation), HNN extensions, free semigroups and free monoids, for example. In some, but not all, of these structures a corresponding algorithm to solve the word problem is known.

The key ideas of our main program described above do not appear to depend on the fact that we are working in a free group – they are essentially combinatorial, and deal with fitting regions together like a jigsaw. Our proof-of-concept implementation works in free products of free groups and cyclic groups of order two, and can easily be generalised to work in arbitrary free products of finite and free groups, possibly even further. A non small-cancellation quotient of a free group can become small cancellation when viewed as a quotient of such free products, tying together Objectives 1.3 and 2.1. The main question to be addressed in this part of the project is to characterise the most general settings in which different versions of our algorithms can be applied.

We will begin, in WP 2.1, by extending our theory and software to work in free products with amalgamation, and HNN extensions in place of the free group. We eventually aim to develop a notion of *relative non-homogeneous small cancellation*, which mimics much of the theory of relatively hyperbolic groups [5, 9], and in particular small cancellation over relatively hyperbolic groups. Following that, we explore the case of monoids and semigroups (WP 2.2). We will start with the notion of small overlap monoids (see traditional and recent work e.g. [12, 14]) and extend it, both through curvature redistribution, and to more general structures such as cancellation semigroups.

A final, adventurous workpackage (WP 2.3) addresses Objective 2.2. We will explore the hypothesis that many classes of objects that are equipped with a notion of the "length" of a group element, such as Coxeter groups, should be suitable for a definition of a diagram, even though elements are not represented as words: the work of Hurley [13] shows that in an abstract setting a weak form of HNN extension suffices to determine a length and normal form. The appropriate interpretation of the word problem is then to determine whether a given group element lies in the normal closure of a given set of group elements. If it does, a suitably-defined diagram will exist that expresses it as a product of their conjugates. We will investigate what conditions are needed to ensure useful geometric properties of such diagrams, and determine the computational and algebraic consequences of those conditions.

Intriguingly, groups acting on lattices are equipped with a natural notion of length. Furthermore, various groups related to the Monster sporadic group can be constructed as quotients of the group of automorphisms of a certain 16-dimensional Lorentzian lattice. We will seek to develop tools to enable computation with such quotients in this way.

D.2.3 Work Programme for Aim 3

Our third aim in this project is to develop and release usable implementations of the algorithms we develop. The work associated with this, above and beyond developing software for direct use within the project, forms workpackage 3.1. We will release our software under a Free and Open license such as the GPL. Our experience with GAP shows that this is an effective way to get rapid uptake, feedback and contributions from the academic community. It will also allow us to integrate the software with GAP (and through GAP with the new SAGE system). Although we expect to implement the key sections of the "main program" in C, for performance, we will provide a GAP interface, making the software usable as a standard GAP package for convenience.

D.3 Management

All four investigators are highly experienced, and have been collaborating and publishing together for many years on both large and small research projects. We therefore anticipate that there will be minimal need for formal management structures. However the Principal Investigator will hold overall management responsibility, and will keep track of the project's progress relative to our objectives and workplan.

Regular meetings of all four investigators will be held, supplemented by meetings of subsets of the investigators on some workpackages. We will all collaborate on all areas of the project, but will also each have key responsibilities. Parker and Roney-Dougal, with input from Linton, will be primarily responsible for the theoretical side of the project, whilst Parker and Neunhöffer will take the lead on implementation, again assisted by Linton. To maximise the impact of our work, we will rapidly turn our progress into journal articles and publicly-accessible code: we will use version-control software so that all of us can edit both code and papers simultaneously.

E Relevance to Academic Beneficiaries

The academic benefits of this research are of two kinds: increased understanding and new questions arising from our theoretical and algorithmic results; and the availability of our software for use by others in their research.

On the theoretical side, a uniform understanding and generalisation of small cancellation properties will be of great interest to geometrical group theorists and should produce new connections between their work and that of monoid and semigroup theorists and others. A deeper understanding of what can happen in the van Kampen diagrams of non-positively curved groups (or groups which "just" fail to be non-positively curved), together with computational tools for exploring examples, may be expected to spark new lines of inquiry in geometric group theory. The curvature rules that we construct can be used for a range of tasks, including (but not limited to) proofs of hyperbolicity, proofs that the group has infinite order, and the finding of free subgroups. If we find STLPs, then they can be used to speed up coset enumeration and quotient algorithms; conversely if we construct a word-problem solver then this can show that a given quotient is isomorphic to the input group.

On the practical side, our programs might help anyone with presentations of algebraic structures that they wish to investigate. By releasing our algorithms as a GAP package we will ensure that many users will benefit, even if they are not specialists. For example, by removing the need to apply Tietze transformations to reduce the size of the output presentation, we anticipate that our methods will improve the efficacy of the Reidemeister–Schreier algorithm. We specifically hope to support the work of George Havas on quotients of the modular group, and to address the very large presentations that arise from some topological problems, supporting the relatively new area of computational topology, as practised, for example, by Graham Ellis's group at NUI Galway.

F Dissemination and Exploitation

Dissemination: Results will be disseminated by means of publications in major journals, presentations at both major international conferences and more specialist fora, and via new software. Our algorithms will also be distributed to the academic and industrial communities via the standard GAP package distribution mechanisms, which have been shown to be effective.

Dr. Roney-Dougal will build on her record of public communication, and will produce at least one popular article and one radio programme aimed at the general public on our discoveries: many of the key ideas are eminently suitable for conveying the excitement of pure mathematics research.

Exploitation: The University of St Andrews has a Research and Enterprise Service whose remit specifically includes technology transfer and licencing of research results. Should commercially exploitable results arise from the project, Research and Enterprise Services will be approached to help with the technology transfer.

References

[1] Dahmani, F., Groves, D. The isomorphism problem for toral relatively hyperbolic groups. *Publ. Math. Inst. Hautes Études Sci.* 107 (2008) 211–290.

[2] Dahmani, F., Guirardel, V. The isomorphism problem for all hyperbolic groups. Preprint: `arxiv.org/abs/1002.2590`.

[3] Epstein, D., Cannon, J., Holt, D., Levy, Silvio V., Paterson, M., Thurston, W. *Word processing in groups.* Jones and Bartlett Publishers 1992.

[4] Conder, M., Havas, G., Newman, M. On one-relator quotients of the modular group. In *Proceedings of Groups St Andrews in Bath, 2009.* CUP, to appear. Preprint: `www.itee.uq.edu.au/~havas/orqmg/`

[5] Farb, B. Relatively hyperbolic groups. *Geom. Funct. Anal.* 8 (1998) 810–840.

[6] The GAP Group, GAP Version 4.4.12, 2008. See `www.gap-system.org`.

[7] Gersten, S. Isoperimetric and isodiametric functions of finite presentations. In *Geometric group theory, Vol. 1.* CUP (1993) 79–96.

[8] Gersten, S., Short, H. Small cancellation theory and automatic groups. *Invent. Math.* 102 (1990) 305–334.

[9] Gromov, M. Hyperbolic groups. In *Essays in group theory.* Springer (1987) 75–263.

[10] Gromov, M. Infinite groups as geometric objects. In *Proc. ICM.* PWN: Warsaw (1984) 385–392.

[11] Huck, G., Rosebrock, S. Applications of diagrams to decision problems. In *Two-dimensional homotopy and combinatorial group theory* CUP (1993) 189–218.

[12] Higgins, P. *Techniques of semigroup theory.* OUP 1992.

[13] Hurley, B. Small cancellation theory over groups equipped with an integer-valued length function. In *Word problems, II.* North-Holland (1980) 157–214.

[14] Kambites, M. Small overlap monoids. I. The word problem. *J. Algebra* 321 (2009) 2187–2205.

[15] Kemmerich, S. *Unendliche Reduktionssysteme.* PhD Thesis, RWTH Aachen, 1983.

[16] Lyndon, R., Schupp, P. *Combinatorial group theory.* Springer-Verlag 1977.

[17] Ollivier, Y. Critical densities for random quotients of hyperbolic groups. *C. R. Math. Acad. Sci. Paris* 336 (2003) 391–394.

[18] Ollivier, Y. Some small cancellation properties of random groups. *Internat. J. Algebra Comput.* 17 (2007) 37–51.

[19] Sela, Z. The isomorphism problem for hyperbolic groups. I. *Ann. of Math. (2)* 141 (1995) 217–283.

[20] Weiss, U. On biautomaticity of non-homogenous small-cancellation groups. *Internat. J. Algebra Comput.* 17 (2007) 797–820.