

SYMBOLIC COLLECTION USING DEEP THOUGHT

C. R. LEEDHAM-GREEN AND LEONARD H. SOICHER

Abstract

We describe the “Deep Thought” algorithm, which can, among other things, take a commutator presentation for a finitely generated torsion-free nilpotent group G , and produce explicit polynomials for the multiplication of elements of G . These polynomials were first shown to exist by Philip Hall, and allow for “symbolic collection” in finitely generated nilpotent groups. We discuss various practical issues in calculations in such groups, including the construction of a hybrid collector, making use of both the polynomials and ordinary collection from the left.

1. Introduction

Let G be a finitely generated torsion-free nilpotent group. Then G has a central series

$$G = G_1 > G_2 > \dots > G_{n+1} = \{1\},$$

such that, for each $1 \leq r \leq n$, the central factor G_r/G_{r+1} is infinite cyclic, generated by $G_{r+1}a_r$, for some $a_r \in G_r$. Given these a_1, \dots, a_n , each element $x \in G$ has a unique expression of the form

$$x = a_1^{x_1} \dots a_n^{x_n},$$

with $x_1, \dots, x_n \in \mathbb{Z}$. We call (x_1, \dots, x_n) the *vector of exponents* for x .

Philip Hall [3, Theorem 6.5] showed that there are rational polynomials f_1, \dots, f_n (in $2n$ variables) and g_1, \dots, g_n (in $n + 1$ variables), which describe multiplication and powering in G . More precisely, suppose $t \in \mathbb{Z}$, and that (x_1, \dots, x_n) , (y_1, \dots, y_n) , (z_1, \dots, z_n) , (w_1, \dots, w_n) are the respective vectors of exponents for $x, y, xy, x^t \in G$. Then, for $1 \leq r \leq n$,

$$z_r = f_r(x_1, \dots, x_n, y_1, \dots, y_n) \quad \text{and} \quad w_r = g_r(x_1, \dots, x_n, t).$$

Now we know that

$$a_j a_i = a_i a_j a_{j+1}^{c_{i,j,j+1}} \dots a_n^{c_{i,j,n}},$$

for certain integers $c_{i,j,k}$ ($1 \leq i < j < k \leq n$). One thing we do in this paper is to show how to compute polynomials in variables corresponding to the $x_i, y_i, c_{i,j,k}$, such that these polynomials have the properties described above for Hall’s polynomials f_r , when restricted to specific $c_{i,j,k}$. These polynomials are very complicated indeed, and in practice the construction and evaluation of the polynomials goes much faster when many of the $c_{i,j,k}$ are zero.

We acknowledge our great debt to the work of Philip Hall, who showed in [3] the existence of the polynomials of the form our “Deep Thought” algorithm calculates in the

Received 21st February 1997, revised 15th January 1998; published 1st June 1998.

1991 Mathematics Subject Classification 20F12, 20F18, 20F05

© 1998, C. R. Leedham-Green and Leonard H. Soicher

case of finitely generated torsion-free nilpotent groups. Charles Sims [8, pp. 441–445] has already described a method of constructing these polynomials of Hall (using interpolation and symbolic algebra), although we have not yet compared his approach to ours. An even more important influence is Philip Hall’s earlier paper [2]. There, he introduces the process of collection, and the reader will see a very strong resemblance between our methods and the work of Hall on the calculation of $(xy)^m$ in the free group $F(x, y)$ modulo some term of its lower central series.

The natural setting for what we do is actually monoids, not groups. In the next section we discuss monoids M satisfying certain relations like those of a finitely generated nilpotent group, and we will develop the Deep Thought algorithm for such monoids.

In Section 7 we show that the polynomials produced by Deep Thought can be used for the multiplication of xy even when our monoid M is a finitely generated nilpotent group, and we allow the x_i , y_i , and $c_{i,j,k}$ to be arbitrary integers. We also show how to invert elements of such a group M .

In Section 8 we discuss efficiency issues related to the application of Deep Thought.

Deep Thought provides a form of symbolic collection for finitely generated nilpotent groups, and in Section 9 we discuss how to combine Deep Thought with ordinary collection from the left [4], to be able to compute effectively and efficiently in arbitrary finitely generated nilpotent groups.

2. The monoid case

Let M be a monoid which is generated by elements

$$a_1, \dots, a_n,$$

such that the the following relations hold:

$$a_j a_i = a_i a_j a_{j+1}^{c_{i,j,j+1}} \dots a_n^{c_{i,j,n}} \quad (1 \leq i < j \leq n), \quad (1)$$

for certain non-negative integers $c_{i,j,k}$ ($1 \leq i < j < k \leq n$). The relations (1) allow us to rewrite an arbitrary word w in a_1, \dots, a_n to a reduced word v , such that $v = w$ in M . A *reduced* word is a word of the form

$$a_1^{x_1} \dots a_n^{x_n},$$

where x_1, \dots, x_n are non-negative integers. The rewriting process we use is called *collection to the left*, introduced by Philip Hall [2], and detailed in Figure 1. (Note that since we have no power relations or negative powers to cancel generators during the collection process, collection to the left is probably the best strategy.)

Let

$$x = a_1^{x_1} \dots a_n^{x_n}, \quad y = a_1^{y_1} \dots a_n^{y_n}$$

be reduced words. We can use collection to the left to determine a reduced word

$$z = a_1^{z_1} \dots a_n^{z_n},$$

such that $xy = z$ in M . We shall show how each exponent z_r ($1 \leq r \leq n$) is the value of a certain polynomial f_r evaluated on the x_i , y_i , and $c_{i,j,k}$, by describing the algorithm, Deep Thought, which calculates these polynomials, given the relations (1).

while w is not a reduced word **do**

Let i^* be the least i such that w contains a subword of the form $a_j a_i$ with $i < j$.

Among the subwords of the form $a_j a_{i^*}$ with $i^* < j$, let $w^* = a_j a_{i^*}$ be the leftmost such subword.

Replace w^* in w by $a_{i^*} a_j a_{j^*+1}^{c_{j^*,j^*,j^*+1}} \cdots a_n^{c_{i^*,j^*,n}}$.

od;

Figure 1: Collection to the left, of w in M .

3. Letters occurring in a collection

Throughout this Section we keep the notation of Section 2.

In order to analyse the process of collection (to the left) of xy in M , we shall attach a unique label to each specific a_i which is introduced in the collection process. These labels come from the set of “letters”, which are defined later. Each letter labelling a generator in the final reduced word is either a letter labelling a generator from the original left-hand word x , or a letter labelling a generator from the original right-hand word y , or is a letter labelling a generator from a word of the form $a_{j+1}^{c_{i,j,j+1}} \cdots a_n^{c_{i,j,n}}$, which is introduced when an $a_j a_i$ (with $i < j$) is collected. We emphasize that in a collection to the left in M , each generator (and its label) which is introduced will appear in the final collected reduced word, since there are no power relations or inverses to cancel generators.

3.1. Atoms and non-atoms

We now (recursively) define the concept of a letter, and show how letters are used to label uniquely the generators which are introduced in the collection of xy in M . (As usual, $x = a_1^{x_1} \cdots a_n^{x_n}$ and $y = a_1^{y_1} \cdots a_n^{y_n}$.)

A *letter* is either an atom or a non-atom. An *atom* α is a 3-tuple

$$(\text{side}(\alpha), \text{num}(\alpha), \text{pos}(\alpha)),$$

such that $\text{side}(\alpha) \in \{\mathbb{L}, \mathbb{R}\}$, and $\text{num}(\alpha)$ and $\text{pos}(\alpha)$ are positive integers. We use the notation

$$\text{num}(\alpha)_{\text{pos}(\alpha)}^{\text{side}(\alpha)}$$

for an atom α .

If $x_i \geq j > 0$ then the atom $i_j^{\mathbb{L}}$ is used to label the j -th from the left a_i in x (the Left-hand word). Similarly, if $y_i \geq j > 0$ then $i_j^{\mathbb{R}}$ is used to label the j -th from the left a_i in y (the Right-hand word).

If a letter is not an atom then it is a non-atom. A *non-atom* α is a 4-tuple

$$(\text{left}(\alpha), \text{right}(\alpha), \text{num}(\alpha), \text{pos}(\alpha)),$$

such that $\text{left}(\alpha)$ and $\text{right}(\alpha)$ are letters, and $\text{num}(\alpha)$ and $\text{pos}(\alpha)$ are positive integers. We use the notation

$$[\text{left}(\alpha), \text{right}(\alpha); \text{num}(\alpha)_{\text{pos}(\alpha)}]$$

for a non-atom α . We also require that this recursion terminates, so that all letters can be

described, using the above notation, as finite expressions involving only atoms. (Note that a letter is just a finite binary tree, with certain information attached to the nodes, and this is how we view a letter from a computational point of view. The leaves of these trees are atoms.)

Now suppose $i < j < k$ and $c_{i,j,k} \geq r > 0$, and that in a collection step we replace a specific $a_j a_i$ by $a_i a_j a_{j+1}^{c_{i,j,j+1}} \dots a_n^{c_{i,j,n}}$. We may assume (inductively) that these specific a_j, a_i have respective labels δ, γ . They retain these labels, and we label the r -th from the left a_k which is introduced in this step by the letter $[\delta, \gamma; k_r]$.

We have thus defined letters, and shown how to label uniquely each specific generator introduced in the collection of xy . Indeed, we can now think of the collection process as collecting letters rather than generators. We say that a letter *occurs* in the collection of xy if it labels a generator introduced in this collection. Thus, a letter occurs in the collection of xy if and only if it labels a generator in the final collected reduced word.

Example 1. Define the monoid N to be generated by a_1, a_2, a_3, a_4 , subject (only) to the relations

$$a_2 a_1 = a_1 a_2 a_3^2, \quad a_3 a_1 = a_1 a_3 a_4, \quad a_3 a_2 = a_2 a_3 a_4^5, \quad \text{and } a_4 a_i = a_i a_4, \quad (i = 1, 2, 3).$$

Let $x = a_1 a_2^3 a_3$ and $y = a_1^3 a_3 a_4$. Examples of atoms which occur in the collection of xy in N are 3_1^L and 1_2^R . An example of a non-atom which occurs in the collection of xy is $[[2_2^L, 1_1^R; 3_1], 1_3^R; 4_1]$.

4. Relations on letters

We now define various relations on the set of letters, including the important equivalence relation, \sim . Given monoid relations of the form (1), the Deep Thought algorithm first determines the \sim -classes which can have elements occurring in a collection to the left. For each such equivalence class, Deep Thought also determines a polynomial which evaluates to the number of letters in that class that occur in a specified collection.

First we make clear what we mean for two letters α, β to be *equal* (written $\alpha = \beta$). Simply, they must both be atoms or both non-atoms, and must be equal as 3-tuples or 4-tuples, respectively. The conditions for letters α, β to be *almost equal* (written $\alpha \approx \beta$) are the same as for them being equal, except that we allow the possibility that $\text{pos}(\alpha) \neq \text{pos}(\beta)$. Any letter is both equal and almost equal to itself.

Example 2. $[[2_1^L, 1_1^R; 3_1], 1_2^R; 4_5]$ is almost equal to $[[2_1^L, 1_1^R; 3_1], 1_2^R; 4_3]$, but is not equal to it. On the other hand, $[[2_1^L, 1_1^R; 3_1], 1_2^R; 4_5]$ is not almost equal to $[[2_1^L, 1_1^R; 3_2], 1_2^R; 4_5]$. Also note that $1_2^L \approx 1_1^L$, but $1_2^L \not\approx 1_2^R$.

A letter α determines a finite sequence $\text{Seq}(\alpha)$ of (not necessarily distinct) *subletters* of α , as follows. If α is an atom then $\text{Seq}(\alpha) = (\alpha)$; otherwise

$$\text{Seq}(\alpha) = \text{Seq}(\text{left}(\alpha)) \text{ concatenate } \text{Seq}(\text{right}(\alpha)) \text{ concatenate } (\alpha).$$

(In binary tree language, $\text{Seq}(\alpha)$ is a postorder transversal of α .) Let $m = \text{Length}(\text{Seq}(\alpha))$. For $1 \leq i \leq m$, we denote the i -th element of $\text{Seq}(\alpha)$ by $\text{Seq}(\alpha, i)$, and let

$$\text{Sub}(\alpha) = \{\text{Seq}(\alpha, i) \mid \Delta \leq i \leq m\}$$

denote the set of subletters of α .

The first condition that must be satisfied to have $\alpha \sim \beta$ is that α and β have the same *structure*, which roughly means that α and β are exactly the same if we (recursively) ignore all pos attributes. More precisely, α and β have the same structure if the following are satisfied:

- α, β must both be atoms or both be non-atoms.
- If α, β are atoms they must be almost equal.
- If α, β are non-atoms, then $\text{num}(\alpha) = \text{num}(\beta)$, $\text{left}(\alpha)$ has the same structure as $\text{left}(\beta)$, and $\text{right}(\alpha)$ has the same structure as $\text{right}(\beta)$.

Note that if α and β have the same structure, we have $\text{Length}(\text{Seq}(\alpha)) = \text{Length}(\text{Seq}(\beta))$.

We are now in a position to define the relation \sim on letters. We have $\alpha \sim \beta$ if each of the following is satisfied:

- α, β have the same structure.
- For $1 \leq i < j \leq \text{Length}(\text{Seq}(\alpha))$:
 - $\text{Seq}(\alpha, i) \approx \text{Seq}(\alpha, j)$ if and only if $\text{Seq}(\beta, i) \approx \text{Seq}(\beta, j)$.
 - If $\text{Seq}(\alpha, i) \approx \text{Seq}(\alpha, j)$ then $\text{pos}(\text{Seq}(\beta, i))$ and $\text{pos}(\text{Seq}(\beta, j))$ must be in exactly the same relation ($<$, $=$, or $>$) as $\text{pos}(\text{Seq}(\alpha, i))$ and $\text{pos}(\text{Seq}(\alpha, j))$.

Example 3. We have

$$[[[2_4^L, 1_3^R; 3_1], 1_4^R; 4_1] \sim [[2_2^L, 1_1^R; 3_1], 1_3^R; 4_5].$$

As a more complicated example, we have

$$\begin{aligned} & [[[[2_2^L, 1_1^R; 4_2], 1_2^R; 6_1], [[2_2^L, 1_1^R; 4_3], 1_3^R; 5_1]; 12_2] \\ & \sim [[[[2_2^L, 1_1^R; 4_1], 1_2^R; 6_1], [[2_2^L, 1_1^R; 4_3], 1_3^R; 5_1]; 12_2] \\ & \not\sim [[[[2_2^L, 1_1^R; 4_2], 1_2^R; 6_1], [[2_2^L, 1_1^R; 4_2], 1_3^R; 5_1]; 12_2]. \end{aligned}$$

Lemma 1. *The relation \sim is an equivalence relation on the set of letters.*

Proof This follows from the fact that both “has the same structure as” and \approx are equivalence relations on the set of letters. \square

For each \sim -class of letters, we define a well-ordering \leq on that class. Suppose $\alpha \sim \beta$. If α and β are both atoms then define $\alpha \leq \beta$ if $\text{pos}(\alpha) \leq \text{pos}(\beta)$. Otherwise, $\alpha \leq \beta$ if in lexicographic order we have

$$(\text{left}(\alpha), \text{right}(\alpha), \text{pos}(\alpha)) \leq (\text{left}(\beta), \text{right}(\beta), \text{pos}(\beta)).$$

(Note that if α and β are non-atoms with $\alpha \sim \beta$, then $\text{left}(\alpha) \sim \text{left}(\beta)$ and $\text{right}(\alpha) \sim \text{right}(\beta)$.)

Lemma 2. *A letter α is least in its \sim -class if and only if for each \approx -class A of $\text{Sub}(\alpha)$, we have $\{\text{pos}(\rho) \mid \rho \in A\} = \{1, \dots, |A|\}$.*

We leave the straightforward proof as an exercise in understanding the many definitions in this section.

Example 4. The least letter in the \sim -class of

$$[[[[2_2^L, 1_1^R; 4_1], 1_2^R; 6_1], [[2_2^L, 1_1^R; 4_3], 1_3^R; 5_1]; 12_2]$$

is

$$[[[[2_1^L, 1_1^R; 4_1], 1_2^R; 6_1], [[2_1^L, 1_1^R; 4_2], 1_3^R; 5_1]; 12_1].$$

```

earlier := function( $\alpha$ ,  $\beta$ )
#
# Suppose that  $\alpha$  and  $\beta$  are letters occurring in a collection (to the left),
# such that  $\alpha$ ,  $\beta$  are not both atoms, and if both  $\alpha$ ,  $\beta$  are
# non-atoms then  $\text{left}(\alpha) \neq \text{left}(\beta)$  or  $\text{right}(\alpha) \neq \text{right}(\beta)$ .
# Then this boolean function returns true if and only if a
# generator is labelled by  $\alpha$  strictly earlier in the collection
# than a generator is labelled by  $\beta$ .
#
if  $\alpha$  is an atom then return true; fi;
if  $\beta$  is an atom then return false; fi;
if  $\text{right}(\alpha) = \text{right}(\beta)$  then return left_of( $\text{left}(\beta)$ ,  $\text{left}(\alpha)$ ); fi;
if  $\text{num}(\text{right}(\alpha)) = \text{num}(\text{right}(\beta))$  then return left_of( $\text{right}(\alpha)$ ,  $\text{right}(\beta)$ ); fi;
return  $\text{num}(\text{right}(\alpha)) < \text{num}(\text{right}(\beta))$ ;
end;

```

Figure 2: The function “earlier”.

5. The function “left_of”

One basic step in the Deep Thought algorithm is the following. Given two letters α , β , which occur in a collection (to the left), with $\text{num}(\alpha) = j > i = \text{num}(\beta)$, we must decide if at the first instance in the collection that α and β both label generators, the generator labelled by α is to the left of the generator labelled by β in the word being collected. If this is so, then all the letters $[\alpha, \beta; k_r]$ will occur for $1 \leq r \leq c_{i,j,k}$.

This decision is made by the function “left_of”, detailed in Figure 3, operating in a double recursion with the function “earlier”, detailed in Figure 2. These functions are given in an algorithmic language similar to GAP [7]. In particular, “#” denotes a comment until the end-of-line, “fi” denotes the end of an if-statement, “od” denotes the end of a while-statement, and “return” followed by an expression means to return the value of that expression as the function value and then terminate the execution of the function.

It not difficult to prove that $\text{earlier}(\alpha, \beta)$ and $\text{left_of}(\alpha, \beta)$ both return the correct results when given valid input, by considering the various cases the functions handle, and using induction on $\text{Length}(\text{Seq}(\alpha)) + \text{Length}(\text{Seq}(\beta))$.

On examination of the functions left_of and earlier , we see that the following (useful) lemma holds.

Lemma 3. *Suppose α is a non-atom, with $\text{left}(\alpha) \neq \text{right}(\alpha)$. Then if $\beta \sim \alpha$, we have $\text{left_of}(\text{left}(\beta), \text{right}(\beta))$ if and only if $\text{left_of}(\text{left}(\alpha), \text{right}(\alpha))$.*

(Note that $\text{left_of}(\text{left}(\alpha), \text{right}(\alpha))$ holds when the non-atom α occurs in a collection.)

6. The Deep Thought algorithm

We are now in a position to describe the basic Deep Thought algorithm.

```

left_of := function( $\alpha$ ,  $\beta$ )
#
# Suppose that  $\alpha$  and  $\beta$  are letters occurring in a collection
# (to the left), such that  $\alpha \neq \beta$ .
# Then this boolean function returns true if and only if
# at the first instance that letters  $\alpha$  and  $\beta$  both label generators,
# the generator labelled by  $\alpha$  is to the left of that labelled by  $\beta$ .
#
if both  $\alpha$ ,  $\beta$  are atoms then
    if side( $\alpha$ ) = L and side( $\beta$ ) = R then return true; fi;
    if side( $\alpha$ ) = R and side( $\beta$ ) = L then return false; fi;
    if num( $\alpha$ ) = num( $\beta$ ) then return pos( $\alpha$ ) < pos( $\beta$ ); fi;
    return num( $\alpha$ ) < num( $\beta$ );
fi;
if both  $\alpha$ ,  $\beta$  are non-atoms and right( $\alpha$ ) = right( $\beta$ ) and left( $\alpha$ ) = left( $\beta$ ) then
    if num( $\alpha$ ) = num( $\beta$ ) then return pos( $\alpha$ ) < pos( $\beta$ ); fi;
    return num( $\alpha$ ) < num( $\beta$ );
fi;
if earlier( $\alpha$ ,  $\beta$ ) then return not left_of( $\beta$ ,  $\alpha$ ); fi;
#
# At this point, we know that  $\beta$  appears earlier than  $\alpha$ , and so  $\alpha$  must be a non-atom.
#
if num( $\beta$ ) < num(right( $\alpha$ )) then return false; fi;
if  $\beta$  = right( $\alpha$ ) then return false; fi;
if num( $\beta$ ) = num(right( $\alpha$ )) then return left_of(right( $\alpha$ ),  $\beta$ ); fi;
#
# At this point, we have num( $\beta$ ) > num(right( $\alpha$ )).
#
if  $\beta$  = left( $\alpha$ ) then return false; fi;
return left_of(left( $\alpha$ ),  $\beta$ );
end;

```

Figure 3: The function “left_of”.

The Deep Thought algorithm takes as input the integer n , and relations of the form (1). The algorithm outputs polynomials f_1, \dots, f_n in indeterminates corresponding to $x_1, \dots, x_n, y_1, \dots, y_n$ and the non-zero $c_{i,j,k}$, such that if z_r is the value of f_r evaluated on $x_1, \dots, x_n, y_1, \dots, y_n$ and the non-zero $c_{i,j,k}$, then

$$a_1^{z_1} \cdots a_n^{z_n} = a_1^{x_1} \cdots a_n^{x_n} a_1^{y_1} \cdots a_n^{y_n},$$

in a monoid M whose generators a_1, \dots, a_n satisfy (1).

Remark The actual values of the non-zero $c_{i,j,k}$ are only required in the evaluation of the polynomials produced by Deep Thought. Indeed, Deep Thought produces polynomials which are valid for all monoids N generated by b_1, \dots, b_n satisfying relations

$$b_j b_i = b_i b_j b_{j+1}^{d_{i,j,j+1}} \cdots b_n^{d_{i,j,n}} \quad (1 \leq i < j \leq n),$$

as long as $d_{i,j,k} = 0$ whenever $c_{i,j,k} = 0$. The reason we only consider the non-zero $c_{i,j,k}$ in the algorithm (rather than to obtain more general formulae) is to keep the polynomials f_r of reasonable size for the given monoid relations.

Deep Thought starts by determining, for each $r = 1, \dots, n$, a set repr_r . When complete, repr_r is a set of representatives for all distinct \sim -classes of letters α , with $\text{num}(\alpha) = r$, such that α occurs in a collection of xy , for some values of the x_i, y_i , and non-zero $c_{i,j,k}$. Then, we will define a polynomial g_α in indeterminates corresponding to the x_i, y_i , and non-zero $c_{i,j,k}$, such that g_α evaluates to the number of letters occurring in the \sim -class of α in the actual collection corresponding to given values of the x_i, y_i , and non-zero $c_{i,j,k}$.

The procedure `set_reps` for determining the sets repr_r is detailed in Figure 4.

After `set_reps` has been executed, we see that each repr_r is a set of inequivalent letters α , with $\text{num}(\alpha) = r$, that α is least in its \sim -class, and that any letter β , with $\text{num}(\beta) = r$, occurring in the collection of xy is equivalent to some element of repr_r .

6.1. The polynomials g_α

Let non-negative integers $x_i, y_i, c_{i,j,k}$ be given, and let ρ be any letter. Define t_ρ by

$$t_\rho = \begin{cases} x_{\text{num}(\rho)} & \text{if } \rho \text{ an atom, side}(\rho) = \text{L} \\ y_{\text{num}(\rho)} & \text{if } \rho \text{ an atom, side}(\rho) = \text{R} \\ c_{\text{num}(\text{right}(\rho)), \text{num}(\text{left}(\rho)), \text{num}(\rho)} & \text{otherwise.} \end{cases}$$

So t_ρ depends only on the integers $x_i, y_i, c_{i,j,k}$ and the \approx -class A of ρ . We define t_A for this \approx -class A to be t_ρ .

Theorem 1. *Let $x = a_1^{x_1} \cdots a_n^{x_n}$, $y = a_1^{y_1} \cdots a_n^{y_n}$. Then for each $r = 1, \dots, n$, and each $\alpha \in \text{repr}_r$ calculated by the procedure `set_reps` with input n and relations (1), the number of elements in the \sim -class of α which occur in the collection (to the left) of xy is*

$$n_\alpha = \prod_{A \in \text{Sub}(\alpha)/\approx} \binom{t_A}{|A|}. \quad (2)$$

Proof For any letter α , define $N(\alpha)$ to be the set of all letters $\beta \sim \alpha$, such that $\text{pos}(\gamma) \leq t_\gamma$ for each $\gamma \in \text{Sub}(\beta)$. Note that for each $B \in \text{Sub}(\alpha)/\approx$ this restriction gives us exactly $\binom{t_B}{|B|}$ possibilities for $\{\text{pos}(\sigma) \mid \sigma \in B\}$, so $N(\alpha)$ has size n_α , where n_α is given by (2).

Now if β is a letter occurring in the collection of xy , we plainly must have $\max\{\text{pos}(\sigma) \mid \sigma \in B\} \leq t_B$, for each \approx -class B of $\text{Sub}(\beta)$. Therefore, $N(\alpha)$ contains the set $C(\alpha)$ of letters which occur in the collection of xy and are in the \sim -class of α . We complete the proof of the theorem by showing that $N(\alpha) \subseteq C(\alpha)$, when α is in one of the sets repr_r .

```

set_reps := procedure( $n, (c_{i,j,k})$ )
#
for  $r := 1, \dots, n$  do # initialize  $\text{reps}_r$  with representative atoms
     $\text{reps}_r := \{r_1^L, r_1^R\}$ ;
od;
for  $r := 3, \dots, n$  do # determine the representative non-atoms
    #
    # Loop invariant: the sets  $\text{reps}_1, \dots, \text{reps}_{r-1}$  are complete.
    #
    for each  $i, j$ , with  $1 \leq i < j < r$ , such that  $c_{i,j,j+1} = \dots = c_{i,j,r-1} = 0 \neq c_{i,j,r}$ , do
        #
        # Determine representatives for all letters which can possibly
        # occur when an  $a_j a_i$  is collected.
        #
        for each  $\alpha \in \text{reps}_i, \beta \in \text{reps}_j$  do
            #
            # Lemma 2 is useful in the tricky business of generating
            # efficiently the pairs  $\gamma, \delta$  to be looped over next.
            #
            for each pair  $\gamma, \delta$  with
                 $\gamma \sim \alpha, \delta \sim \beta$ , and  $[\delta, \gamma; r_1]$  the least letter in its  $\sim$ -class do
                    if  $\text{left\_of}(\delta, \gamma)$  then
                        #
                        # for each  $k$  such that  $c_{i,j,k} \neq 0$ ,  $[\delta, \gamma; k_1]$  is the least letter in its  $\sim$ -class
                        #
                        for each  $k \in \{r, \dots, n\}$  such that  $c_{i,j,k} \neq 0$  do Add  $[\delta, \gamma; k_1]$  to  $\text{reps}_k$ ; od;
                    fi;
                od;
            od;
        od;
    od;
od;
end;

```

Figure 4: The procedure “set_reps”.

Let $R = \cup_{r=1}^n \text{reps}_r$. For each letter $\alpha \in R$ and for each $\beta \in N(\alpha)$, we show that $\beta \in C(\alpha)$, by induction on $l = \text{Length}(\text{Seq}(\beta))$.

If $l = 1$ then β is an atom, and $\text{pos}(\beta) \leq t_\beta$, that is, $\text{pos}(\beta) \leq x_{\text{num}(\beta)}$ if $\text{side}(\beta) = \text{L}$, and $\text{pos}(\beta) \leq y_{\text{num}(\beta)}$ if $\text{side}(\beta) = \text{R}$. Thus the result holds for $l = 1$.

Now assume that $\beta \in N(\alpha)$ is a non-atom, and let λ be the least letter in the \sim -class of $\text{left}(\beta)$ and ρ be the least letter in the \sim -class of $\text{right}(\beta)$. Since $\beta \sim \alpha$, it follows that $\text{left}(\beta) \sim \text{left}(\alpha)$, and $\text{right}(\beta) \sim \text{right}(\alpha)$; so λ is also the least letter in the \sim -class of $\text{left}(\alpha)$ and ρ is the least letter in the \sim -class of $\text{right}(\alpha)$. Since $\alpha \in R$, from the set_reps construction of the sets reps_r , it follows that $\lambda \in R$ and $\rho \in R$. Also $\beta \in N(\alpha)$ implies $\text{left}(\beta) \in N(\lambda)$ and $\text{right}(\beta) \in N(\rho)$. Therefore, by our inductive hypothesis, $\text{left}(\beta)$ and $\text{right}(\beta)$ occur in the collection of xy . Since $\alpha \in R$, we know that $\text{left_of}(\text{left}(\alpha), \text{right}(\alpha))$ holds, but then $\text{left_of}(\text{left}(\beta), \text{right}(\beta))$ also holds by Lemma 3. But $\text{pos}(\beta) \leq t_\beta = c_{\text{num}(\text{right}(\beta)), \text{num}(\text{left}(\beta)), \text{num}(\beta)}$, so β occurs in the collection of xy , and the proof is complete. \square

We now define the polynomial g_α that evaluates to the number of letters equivalent under \sim to α occurring in the collection of xy . This definition is justified by the preceding theorem.

Let $X_i, Y_i, C_{i,j,k}$ be indeterminates, and let ρ be any letter. Define the indeterminate T_ρ by

$$T_\rho = \begin{cases} X_{\text{num}(\rho)} & \text{if } \rho \text{ an atom, side}(\rho) = \text{L} \\ Y_{\text{num}(\rho)} & \text{if } \rho \text{ an atom, side}(\rho) = \text{R} \\ C_{\text{num}(\text{right}(\rho)), \text{num}(\text{left}(\rho)), \text{num}(\rho)} & \text{otherwise.} \end{cases}$$

So T_ρ depends only on the indeterminates $X_i, Y_i, C_{i,j,k}$ and the \approx -class A of ρ . We define T_A for this \approx -class A to be T_ρ . Let T be an indeterminate, k a non-negative integer, and define the polynomial $\binom{T}{k}$ by

$$\binom{T}{k} = \prod_{i=1}^k \frac{T - i + 1}{i}.$$

Now, finally, define

$$g_\alpha = \prod_{A \in \text{Sub}(\alpha)/\approx} \binom{T_A}{|A|}.$$

For example, if $\delta = [[2_1^{\text{L}}, 1_1^{\text{R}}; 3_1], 1_2^{\text{R}}; 4_1]$, then $g_\delta = \binom{X_2}{1} \binom{Y_1}{2} \binom{C_{1,2,3}}{1} \binom{C_{1,3,4}}{1}$.

Note that g_α has total degree equal to $|\text{Sub}(\alpha)|$.

6.2. The output of Deep Thought

The final steps in the Deep Thought algorithm, after computing the sets reps_r , are to calculate

$$f_r = \sum_{\alpha \in \text{reps}_r} g_\alpha,$$

for $1 \leq r \leq n$, and to output these f_r .

Remark To evaluate an f_r on given integers $x_i, y_i, c_{i,j,k}$, we simply substitute for each indeterminate T_A in the expression for f_r , the corresponding value t_A . Note that f_r is a rational polynomial which takes integer values when evaluated on integers.

7. The group case and negative exponents

Let G be an arbitrary finitely generated nilpotent group. Then for some n , G contains a generating sequence a_1, \dots, a_n , such that these generators satisfy relations of the form (1), for certain (not necessarily non-negative) integers $c_{i,j,k}$ ($1 \leq i < j < k \leq n$). Each element of G can be written in the form $a_1^{x_1} \cdots a_n^{x_n}$, where x_1, \dots, x_n are integers.

We may run the Deep Thought algorithm with input n and the relations (1), even if some of the $c_{i,j,k}$ are negative, since the algorithm only cares whether a given $c_{i,j,k}$ is zero or not. Suppose this run produces sets $\text{reps}_1, \dots, \text{reps}_n$, and polynomials f_1, \dots, f_n as output. Then we have the following:

Theorem 2. *Let $x_1, \dots, x_n, y_1, \dots, y_n$ be arbitrary integers, and for $1 \leq r \leq n$, let z_r be the integer obtained by evaluating f_r , on the given x_i, y_i , and non-zero $c_{i,j,k}$. If $x = a_1^{x_1} \cdots a_n^{x_n}$, and $y = a_1^{y_1} \cdots a_n^{y_n}$, and $z = a_1^{z_1} \cdots a_n^{z_n}$, then $xy = z$ in G .*

Proof Since G is a polycyclic group, it is residually finite by a theorem of K. Hirsch (see [6, 5.4.17]), and so $xy = z$ in G if and only if $\bar{x}\bar{y} = \bar{z}$ in every finite quotient \bar{G} of G . Thus, it suffices to prove the theorem under the assumption that G is finite.

We now assume that G is finite, but $xy \neq z$ in G . Then $xy \neq z$ in G if we add to each exponent x_r, y_r , and z_r of x, y , and z some multiple of $|G|$. Note that this would be a consequence of adding to each x_i, y_i , and non-zero $c_{i,j,k}$ a multiple of

$$\max\{|A| \mid A \in \text{Sub}(\alpha) / \approx, \alpha \in \cup_{r=1}^n \text{reps}_r\} \times |G|,$$

and recalculating each z_r . But then this would imply that $xy \neq z$ in G even if each x_i, y_i , and $c_{i,j,k}$ is non-negative, which would contradict the fact that Deep Thought works when all such exponents are non-negative. \square

7.1. Inverting elements of G

We now show how to invert elements of G using Deep Thought polynomials. More generally, we show how to solve for y in the equation $xy = z$ in G , where now $x = a_1^{x_1} \cdots a_n^{x_n}$ and $z = a_1^{z_1} \cdots a_n^{z_n}$ are given, and we wish to determine $y = a_1^{y_1} \cdots a_n^{y_n}$ (of course, calculating $y = x^{-1}$ is just the special case $z_1 = \dots = z_n = 0$).

We first make a copy x' of x . Then, for $i := 1, \dots, n$, given the loop invariant that x' is of the form $a_1^{z_1} \cdots a_{i-1}^{z_{i-1}} a_i^{x'_i} \cdots a_n^{x_n}$, set $y_i := z_i - x'_i$, and use the Deep Thought polynomials to calculate $x' := x' a_i^{y_i}$.

7.2. Calculating normal forms of elements of G

We have shown that Deep Thought polynomials can be used to multiply and invert elements of an arbitrary finitely generated nilpotent group, and that Deep Thought handles negative exponents correctly. The only problem is that the result $w = a_1^{w_1} \cdots a_n^{w_n}$ of such a multiplication or inversion need not be uniquely determined by the group element represented by w . We get around this problem by using a consistent power-commutator presentation, and we show how to use the power relations of such a presentation and Deep Thought to convert a result into canonical (normal) form.

Let G be an arbitrary finitely generated nilpotent group. Then G has (for some n) a so-called consistent power-commutator presentation of the form

$$\langle a_1, \dots, a_n \mid a_i^{m_i} = a_{i+1}^{c_{i,i,i+1}} \cdots a_n^{c_{i,i,n}}, a_j a_i = a_i a_j a_{j+1}^{c_{i,j,j+1}} \cdots a_n^{c_{i,j,n}} \ (i < j) \rangle, \quad (3)$$

such that each m_i ($1 \leq i \leq n$) is a non-negative integer, and each element of G has a unique expression of the form $a_1^{x_1} \cdots a_n^{x_n}$, where x_1, \dots, x_n are integers, and if $m_i > 0$, then $0 \leq x_i < m_i$. (If $m_i = 0$ then we must have $c_{i,i,i+1} = \cdots = c_{i,i,n} = 0$ and the relation with $m_i = 0$ can effectively be ignored.) This unique expression for an element of G is referred to as the *normal form* of that element.

Now if each $m_i = 0$ in the consistent power-commutator presentation (3) for G , then Deep Thought polynomials can be used directly to multiply elements in G , such that the result is always in normal form. Otherwise, we may obtain a result $z = a_1^{z_1} \cdots a_n^{z_n}$, such that for some i with $m_i > 0$, we have $z_i < 0$ or $z_i \geq m_i$. Suppose this is the case, and let j be the least i with this property. Then we need to “normalize” $z^* = a_j^{z_j} \cdots a_n^{z_n}$, that is, replace it by a word in normal form representing the same element of G . Let $z_j = qm_j + r$, where q and r are integers and $0 \leq r < m_j$. We calculate

$$w := (a_{j+1}^{c_{j,j,j+1}} \cdots a_n^{c_{j,j,n}})^q a_{j+1}^{z_{j+1}} \cdots a_n^{z_n},$$

using Deep Thought polynomials, so that w has the form $a_{j+1}^{w_{j+1}} \cdots a_n^{w_n}$. We then (recursively) normalize w , and so the normal form for z^* is $a_j^r w$, and the normal form for z is

$$a_1^{z_1} \cdots a_{j-1}^{z_{j-1}} a_j^r w.$$

7.3. Determining the order of an element of G

Suppose (3) is a consistent power-commutator presentation for G , and x is an element of G in normal form. Then we can calculate the order $|x|$ of x as follows.

If x is the empty word then $|x| = 1$. Otherwise, $x = a_j^{x_j} \cdots a_n^{x_n}$, with $j \leq n$ and $x_j \neq 0$. If $m_j = 0$ then $|x| = \infty$. Otherwise let $m = m_j / \gcd(m_j, x_j)$. Now $|x| = \infty$ if and only if $|x^m| = \infty$, and if the order of x is finite then m divides $|x|$. We then calculate the normal form of $y = x^m$, which is $a_l^{y_l} \cdots a_n^{y_n}$, with $l > j$, and recursively determine $|y|$. If $|y| = \infty$ then $|x| = \infty$, otherwise $|x| = m|y|$.

7.4. On the degree of Deep Thought polynomials

Suppose (3) is a consistent power-commutator presentation for the group G , such that the central series defined by a_1, \dots, a_n refines the lower central series

$$G = \gamma_1(G) > \gamma_2(G) > \cdots > \gamma_{c+1}(G) = \{1\}$$

of G . Let α be an element of a set reps_r determined by Deep Thought with input (3).

It is not difficult to see that $\text{Seq}(\alpha)$ contains at most c elements which are atoms (since $[\gamma_i(G), \gamma_j(G)] \leq \gamma_{i+j}(G)$). It follows that g_α , considered as a polynomial in the X_i and Y_i only, has degree at most c , and the same holds for the Deep Thought polynomial f_r .

Suppose $\text{Seq}(\alpha)$ has exactly d elements which are atoms. Then

$$\text{Length}(\text{Seq}(\alpha)) = 2d - 1$$

(proof by induction on d). It follows that g_α has total degree (in the X_i, Y_i , and $C_{i,j,k}$) at most $2d - 1 \leq 2c - 1$, and the same holds for the Deep Thought polynomial f_r .

8. Efficiency issues in applying Deep Thought

Let G be a group with presentation (3), and suppose we are interested in efficient multiplication and inversion of elements in G , rather than in the polynomials f_r produced by Deep Thought.

```

for  $r := 1, \dots, n$  do
  if  $r = s$  then
     $\text{reps}_{rs} := \{r_1^L, r_1^R\}$ ;
  else
     $\text{reps}_{rs} := \{r_1^L\}$ ;
  fi;
od;
    
```

 Figure 5: Initial **for**-loop for the calculation of the sets reps_{rs} .

8.1. The polynomials f_{rs}

The first important observation is that we should not actually calculate and use the Deep Thought polynomials f_r for multiplying and inverting elements of G , but instead, closely related polynomials f_{rs} , described below, which can be calculated by a small variant of the Deep Thought algorithm.

Let $\text{reps}_1, \dots, \text{reps}_n$ be the sets of representative letters produced by Deep Thought, using input (3). For $1 \leq r, s \leq n$, define

$\text{reps}_{rs} =$

$\{\alpha \in \text{reps}_r \mid \text{if } \beta \in \text{Sub}(\alpha), \beta \text{ an atom with side}(\beta) = R, \text{ then num}(\beta) = s\}$,

and

$$f_{rs} = \sum_{\alpha \in \text{reps}_{rs}} g_\alpha.$$

Then for $1 \leq r, s \leq n$, we see that if z_r is the value of f_{rs} evaluated on x_1, \dots, x_n, y_s , and the (non-zero) $c_{i,j,k}$, then

$$a_1^{x_1} \cdots a_n^{x_n} a_s^{y_s} = a_1^{z_1} \cdots a_n^{z_n} \quad (4)$$

in the group G .

To determine $z = xy$, for arbitrary $y = a_1^{y_1} \cdots a_n^{y_n}$, we just set $z := x$, and then, for $s := 1, \dots, n$, we set $z := za_s^{y_s}$, using the f_{rs} .

Furthermore, efficient multiplication of the form (4) is exactly what we need when inverting elements of G and, as it will turn out, what we also need for hybrid collection in G , described in Section 9.

Even better, calculating all the sets reps_{rs} is computationally no harder (and is often easier) than calculating all the sets reps_r . To calculate the sets reps_{rs} , for a fixed s and for $1 \leq r \leq n$, we use an algorithm which is almost the same as that detailed in Figure 4. The first change is that the initial **for**-loop should be replaced by the code in Figure 5. Then, throughout the rest of the algorithm, every occurrence of reps_\square should be replaced by $\text{reps}_{\square s}$.

8.2. More on efficiency

The next observation is that if (in our fixed presentation (3)) some $c_{i,j,k} > 0$, then when calculating the sets reps_{rs} , we need not include a representative α with an \approx -class A of subletters with $T_A = C_{i,j,k}$ but $|A| > c_{i,j,k}$, because if $\beta \sim \alpha$ is a subletter of γ , then g_γ

while w is not in normal form **do**

Let w^* be the leftmost subword of w of the form $a_j a_i$ with $i < j$, or of the form $a_i^{m_i}$, where $m_i > 0$ in (3).

if $w^* = a_j a_i$, with $i < j$, **then** replace w^* in w by $a_i a_j a_{j+1}^{c_{i,j,j+1}} \cdots a_n^{c_{i,j,n}}$;

else replace w^* in w by $a_{i+1}^{c_{i,i,i+1}} \cdots a_n^{c_{i,i,n}}$ **fi**;

od;

Figure 6: Collection from the left, of w in G .

would always evaluate to 0 with this $c_{i,j,k}$.

Another obvious point is that we should substitute the actual value $c_{i,j,k}$ for each indeterminate $C_{i,j,k}$ in a polynomial produced by Deep Thought, and try to simplify such polynomials to make them easier to evaluate. For example, we should make use of the fact that if α and β are non-atoms with $\text{left}(\alpha) = \text{left}(\beta)$ and $\text{right}(\alpha) = \text{right}(\beta)$ then the polynomials g_α and g_β are nearly the same.

One small trick which has proved fruitful is to precompute the values of $\binom{m}{k}$ for small $|m|$ and k (say $0 \leq |m|, k \leq 20$), and use these precomputed values when evaluating Deep Thought polynomials.

If a generator a_i of G has finite order, we should work modulo this order when calculating with the exponent of a_i , to reduce the work involved in integer arithmetic. The calculation of the order of an element of G is described in Section 7.3, above.

9. Hybrid collection

We have discussed how to use Deep Thought polynomials to multiply and invert elements in the finitely generated nilpotent group G defined by the consistent power-commutator presentation (3). However, it may sometimes be more efficient (in terms of space or time) to adopt another strategy, such as collection from the left [4], detailed in Figure 6. (For ease of exposition, we shall assume that all exponents are non-negative. This is certainly the case for our implementation of collection from the left in (finite) p -groups.) Deep Thought tends to be best for groups of low class and high exponent, and we now describe how to combine Deep Thought with collection from the left to be able to multiply and invert more efficiently than one or the other approach could on its own.

The trick is to determine an integer $d \leq n$, so that Deep Thought deals efficiently with $G_d = \langle a_d, \dots, a_n \rangle$. This may require some experimentation. We then calculate the Deep Thought polynomials f_{rs} for $d \leq s \leq r \leq n$, to be able to multiply $a_1^{x_1} \cdots a_n^{x_n} a_s^{y_s}$ efficiently using Deep Thought.

Now, in the process of collection, our hybrid collector is either in Deep Thought mode or in from-the-left mode, and is multiplying a word $x = a_1^{x_1} \cdots a_n^{x_n}$ in normal form (stored as an exponent vector in practice) times some other word v (stored on a stack in practice). (We have finished exactly when $v = 1$.)

Suppose we are in Deep Thought mode and v starts with a word of the form $v^* = a_s^{y_s}$. If $s \geq d$, then we remove v^* from v , calculate $x := x v^*$ using Deep Thought polynomials

(but not normalizing the result), and continue the process. If $s < d$ then we first normalize x using Deep Thought polynomials, switch to from-the-left mode, and continue the process.

Suppose we are in from-the-left mode and v starts with a word of the form $a_s^{y_s}$. If $s \geq d$, then we switch to Deep Thought mode, and continue the process. If $s < d$ then we remove the leading a_s from v , and collect this a_s into x using collection from the left in the ordinary way, and then continue the process.

10. *Implementations of Deep Thought*

We have implemented the Deep Thought algorithm and the hybrid collector described here (the hybrid for p -groups only), in the C programming language. Our Deep Thought implementation seems to work well for arbitrary finitely generated nilpotent groups up to about class 8 or more, and the hybrid collector can speed up multiplication in p -groups of much higher class. The Deep Thought C implementation has also been used by Paul Igodt and his colleagues to study various aspects of finitely generated torsion-free nilpotent groups (see, for example, [1]).

More recently, Deep Thought has been implemented by Wolfgang Merkwitz in the GAP system, and will thus be available for public use. This implementation is described in [5].

Merkwitz [5] comes to the conclusion that Deep Thought might be used to calculate in a p -group if $p \geq 11$. He gives a detailed account of the experiments leading to this conclusion. His evidence is very much in line with ours. He finds that, working on a Pentium PC operating at 166 MHz, the time taken to construct the Deep Thought polynomials in some p -groups of composition length 35 and nilpotency class 9 is about three seconds, regardless of the prime p . He finds that multiplication using Deep Thought out-performs collection (from the left) in evaluating the product of two random words by a factor of about 7 if $p = 7$, and a factor of about 10,000 if $p = 47$.

Merkwitz also finds that, for example, computing the derived subgroup of the groups he considers using collection takes 0.2 seconds, regardless of the prime. This presumably means that the collections performed are almost entirely trivial. The time required using Deep Thought is also, of course, effectively independent of the prime, but is almost five times as long. This emphasises the fact that, as currently implemented, collection from the left is out-performing Deep Thought multiplication in trivial calculations.

Acknowledgements

We thank Wolfgang Merkwitz, Joachim Neubüser, Werner Nickel, Charles Sims and Michael Vaughan-Lee for interesting and useful discussions. This research was partly supported by a European Union HCM grant in Computational Group Theory.

References

1. K. DEKIMPE and P. IGDOT, 'Computational aspects of affine representations for torsion free nilpotent groups via the Seifert construction', *J. Pure Appl. Algebra* 84 (1993) 165–190. 23
2. P. HALL, 'A contribution to the theory of groups of prime-power order', *Proc. London Math. Soc. (2)* 36 (1934) 29–95. 10, 10

3. P. HALL, 'Nilpotent groups', Notes of lectures given at the Canadian Mathematical Congress 1957 Summer Seminar, in *The collected works of Philip Hall* (Clarendon Press, Oxford, 1988) pp. 415–462. 9, 9
4. C.R. LEEDHAM-GREEN and L.H. SOICHER, 'Collection from the left and other strategies', *J. Symbolic Comp.* 9 (1990) 665–675. 10, 22
5. W.W. MERKWITZ, 'Symbolische multiplikation in nilpotenten Gruppen mit Deep Thought', Diplomarbeit, RWTH Aachen, 1997. 23, 23
6. D.J.S. ROBINSON, *A course in the theory of groups* (Second Edition) (Springer, New York and Berlin, 1996). 19
7. M. SCHÖNERT *et al.*, 'GAP: groups, algorithms and programming', version 3, release 4, Lehrstuhl D für Mathematik, RWTH Aachen, 1994. 14
8. C.C. SIMS, *Computation with finitely presented groups* (Cambridge University Press, Cambridge, 1994). 9

C. R. Leedham-Green C.R.Leedham-Green@qmw.ac.uk
Leonard H. Soicher L.H.Soicher@qmw.ac.uk

School of Mathematical Sciences
Queen Mary and Westfield College
Mile End Road, London E1 4NS, U.K.