On Complexity of Confluence and Church-Rosser Proofs

Arnold Beckmann
Department of Computer Science, Swansea University, UK

Georg Moser
Department of Computer Science, University of Innsbruck, Austria

Abstract

In this paper, we investigate *confluence* and the *Church-Rosser property* – two well-studied properties of rewriting and the λ-calculus – from the viewpoint of proof complexity. With respect to confluence, and focusing on orthogonal term rewrite systems, our main contribution is that the size, measured in number of symbols, of the smallest rewrite proof is polynomial in the size of the peak. For the Church-Rosser property we obtain exponential lower bounds for the size of the join in the size of the equality proof. Finally, we study the complexity of proving confluence in the context of the λ-calculus. Here, we establish an exponential (worst-case) lower bound of the size of the join in the size of the peak.

2012 ACM Subject Classification Theory of computation → Proof complexity

Keywords and phrases logic, bounded arithmetic, consistency, rewriting

Digital Object Identifier 10.4230/LIPIcs.MFCS.2024.21

Funding *Arnold Beckmann*: Royal Society International Exchanges Grant, IES\R3\223051
Georg Moser: Royal Society International Exchanges Grant, IES\R3\223051

1 Introduction

Confluence and the Church-Rosser property are two (very) well-known properties of rewriting that have been studied for several decades. *Confluence* expresses that if we have terms s, t, t', where s can be successively rewritten to t, as well as to t', then t and t' have a common descendent in the rewriting relation, cf. Figure 1 i). In short, if there is a *peak*: $t \,{}^*\!\leftarrow s \rightarrow^* t'$, we conclude the existence of a *rewrite proof*: $t \rightarrow^* \cdot \,{}^*\!\leftarrow t'$. The *Church-Rosser property* – illustrated in Figure 1 ii) – expresses that from the equality between t and t' ($t \leftrightarrow^* t'$), we conclude the existence of a rewrite proof: $t \rightarrow^* \cdot \,{}^*\!\leftarrow t'$. It is a folklore result that both properties are equivalent. And, as indicative in the name, their intensive study goes back to work by Church and Rosser [8].

Despite the large body of work on confluence and the Church-Rosser property, it seems that the, to us, natural question about the inherent proof complexities has only received scarce attention. A noteworthy exception is work by Ketema and Grue Simonsen [11]. Focusing on orthogonal term rewrite systems and employing the number of reductions as measure of proof complexity, they obtain in the context of confluence optimal exponential upper bounds on the size of the rewrite proof in relation to the size of the peak. With respect to the Church-Rosser property only a non-elementary upper bound can be shown. Related results have been obtained for the λ-calculus, where again non-elementary bounds are obtained for both properties, cf. [10].

If, however, proof complexity is measured more in the tradition of computational complexity, that is, as the number of symbols occurring in a proof, then more tractable results are possible. For example for orthogonal term rewrite systems, we prove that for confluence the size of the least rewrite proof is always polynomially bounded in the size of the peak.

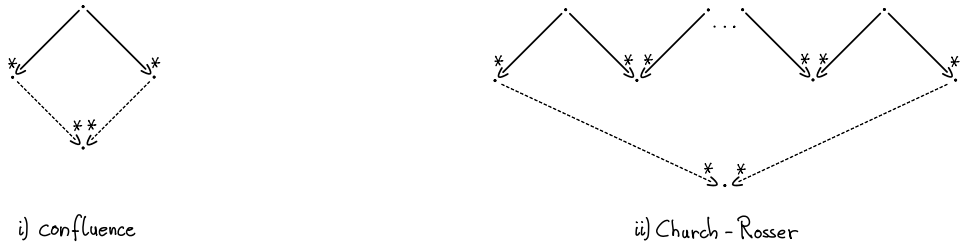

Figure 1 Confluence and Church-Rosser property.

Motivation. These results may open the way for the application of rewriting techniques in complexity theoretic studies, in particular in the context of Bounded Arithmetic [6]. A major open problem in Bounded Arithmetic is the separation of its fragments, which has deep connections to similar questions about the separation of computational complexity classes like the Polynomial Time Hierarchy, including the P vs. NP problem. Consider equational theories, restricted to term equations that define functions symbols exclusively by recursion. As established in [5] by the first author, consistency of such equational theories can be proved in the fragment of Bounded Arithmetic S^1_2. This is remarkable, as it disproves the general impression in Bounded Arithmetic, that consistency statements cannot be used for separation arguments - consistency of equational theories with a richer set of axioms are usually unprovable in Bounded Arithmetic [7].

In the proof in [5], the given equational proof is reconstructed in S^1_2 using a technically involved process of "approximation" and "calculation". An alternative, much more elegant, proof could employ the Church-Rosser property of the induced term rewrite system. To our best knowledge it is, however, unclear whether this property (or confluence) is formalisable in S^1_2. The results of this paper are conceivable as a first step towards this direction.

Contributions. In summary, we make the following contributions, where we are only concerned with *orthogonal* term rewrite systems.

1) Our main result, Theorem 17, shows that the size – measured in the number of symbols – of the smallest possible rewrite proofs is in the worst-case polynomially bounded in the size of the peak, cf. Figure 1. This shows that confluence properties are polynomial time computable, hence are formalisable in Bounded Arithmetic.
   The polynomial (in fact biquadratic) upper bound stems from a quadratic bound on the number of reductions in the rewrite proof in the size of the peak, and a quadratic bound on the size of each term in the rewrite proof.
2) For the Church-Rosser property we give an exponential worst-case lower bound to the size of the join in the size of the equality proof, cf. Theorem 19. This shows that it is not possible to formalise Church-Rosser properties directly in Bounded Arithmetic. The (worst-case) bound is precise.
3) We give matching (worst-case) upper and lower bounds based on different complexity measures. For confluence, we show that the size of the join is linear in the size of the product of the end terms in the peak, cf. Corollary 15 and Proposition 10. For the Church-Rosser property, we show that the size of the join is polynomial in the product of the sizes of the intermediary terms in the equational proof, cf. Theorem 22 and Proposition 21.
4) Finally, we study the complexity of proving confluence in the context of the λ-calculus. We obtain that the size of the join is at least exponential in the size of the peak. Hence, confluence is also not formalisable directly in Bounded Arithmetic.

Outline

The next section introduces basic notions and results. In Section 3 we establish the mentioned lower bound results for rewriting. Section 4 introduces technical notions that underly the methodology of our main results, to be presented in Section 5. In Section 6 we study lower and upper bounds on the complexity of Church-Rosser proofs. The lower bound of confluence proofs is established in Section 7. Section 8 discusses related works. Finally, in Section 9, we conclude and present future work.

2 Preliminaries

We assume (at least nodding) acquaintance with term rewriting [2, 12], however recall basic definitions and notations for ease of readability.

General. Let R be a binary relation. We write R^n for the n-fold iteration of R and R^* for the reflexive and transitive closure of R. Let $\mathcal{V}$ denote a countable infinite set of variables, and $\mathcal{F}$ a countable infinite set of function symbols (also called signature). The set of terms over $\mathcal{F}$ and $\mathcal{V}$ is denoted by $\mathcal{T}(\mathcal{F},\mathcal{V})$.

Let t be a term (over $\mathcal{F}$ and $\mathcal{V}$). A *position* p is a finite sequence of positive integers. Via positions, we uniquely identifying subterms of t, denoted as $t|_p$. We write $p\|q$ to indicate parallel positions, generalising the notions suitably to sets of positions. We write $\mathsf{Var}(t)$ to denote the set of variables occurring in t, ie. $\mathsf{Var}\,t = \{x \mid t|_p \text{ is a variable for some position } p\}$ and we write $\mathsf{rt}(t)$ to denote its root symbol. For example, for $\{x,y\} \subseteq \mathcal{V}$, $\mathsf{Var}(x+y) = \{x,y\}$ and $\mathsf{rt}(x+y) = +$. The *size* $|t|$ of term t is defined as the number of symbol occurrences in t, for example, $|x+y| = 3$. A term t is *linear* if every variable in t occurs only once.

Term Rewriting. A *rewrite rule* is a pair $l \to r$ of terms, such that (i) the left-hand side l is not a variable and (ii) $\mathsf{Var}(l) \supseteq \mathsf{Var}(r)$. A *term rewrite system* (TRS) over $\mathcal{F}$ is a finite set of rewrite rules $\mathcal{R}$; it will be denoted by the pair $(\mathcal{F},\mathcal{R})$. If the signature $\mathcal{F}$ is clear from context, we simply denote a TRS by its set of rules $\mathcal{R}$. If $l \to r$ is a rewrite rule and σ a renaming, then the rule $l\sigma \to r\sigma$ is called a *variant* of $l \to r$. A TRS is said to be *variant-free*, if it does not contain rewrite rules that are variants. In the following we assume that TRSs are variant-free.

The rewrite relation based on $\mathcal{R}$ is denoted as $\to_{\mathcal{R}}$ and its transitve and reflexive closure as $\to^*_{\mathcal{R}}$. If the TRS is clear from context, we will simply write $\to$ and $\to^*$ respectively. Let s be a redex in term t. Here a *redex* is an occurrence of a term s that is an instance of the left-hand side l of a rule $l \to r \in \mathcal{R}$. We write $t \xrightarrow{s}_{\mathcal{R}} t'$ to indicate that redex s is contracted in the rewrite step. A term t over $\mathcal{T}(\mathcal{F},\mathcal{V})$ is in *normal form* with respect to a TRS $\mathcal{R}$, if t does not contain any redex. We call a substitution σ *normalised* (with respect to $\mathcal{R}$), if all terms in the range of σ are in normal form. The *innermost rewrite relation* $\xrightarrow{\mathrm{i}}_{\mathcal{R}}$ of a TRS $\mathcal{R}$ is defined as follows: $s \xrightarrow{\mathrm{i}}_{\mathcal{R}} t$ if there exists a rewrite rule $l \to r \in \mathcal{R}$, a context C, and a substitution σ such that $s = C[l\sigma]$, $t = C[r\sigma]$, and all proper subterms of $l\sigma$ are normal forms of $\mathcal{R}$.

An *overlap* for $\mathcal{R}$ is a triple $\langle l \to r, p, l' \to r' \rangle$, such that (i) $l \to r$, $l' \to r'$ are rules in $\mathcal{R}$, whose variables are disjoint, (ii) p is not a variable position in l', (iii) l and $l'|_p$ are unifiable, (iv) if $p = \varepsilon$, then $l \to r$, $l' \to r'$ are not variants. A TRS is *left-linear* if the left-hand sides of all rules are linear. A TRS $\mathcal{R}$ without overlap is called *non-ambiguous*; a left-linear, non-ambiguous TRS is called *orthogonal*.

Let s and t be terms. Then an *(innermost) derivation* $D\colon s \to^*_{\mathcal{R}} t$ with respect to a TRS $\mathcal{R}$ is a finite sequence of (innermost) rewrite steps. Given an equational system $\mathcal{E}$, we can define, as usual, a TRS $\mathcal{R}$ such that

$$s =_{\mathcal{E}} t \quad \text{iff} \quad s \leftrightarrow^*_{\mathcal{R}} t\,.$$

(See [2, 12] for the straightforward construction.) A finite sequence of equational steps: $t_1 \leftrightarrow_{\mathcal{R}} t_2 \cdots \leftrightarrow_{\mathcal{R}} t_n$ is called an *equational proof*.

A term $s \in \mathcal{T}(\mathcal{F},\mathcal{V})$ is *confluent*, if for all $t, t' \in \mathcal{T}(\mathcal{F},\mathcal{V})$ with $t \,{}^*\!\leftarrow s \to^* t'$, there exists a common reduct v, that is, $t \to^* v \,{}^*\!\leftarrow t'$. A TRS $(\mathcal{F},\mathcal{R})$ is *confluent* if all terms in $\mathcal{T}(\mathcal{F},\mathcal{V})$ are confluent. We call the equational proof $t \,{}^*\!\leftarrow s \to^* t'$ a *peak*, the term v the *join* and the derivations $t \to^* v \,{}^*\!\leftarrow t'$ a *rewrite proof*. A peak is *local*, if it consists of one step each: $t \leftarrow s \to t'$. Confluence is equivalent to the *Church-Rosser property*, which states that for any equational proof $t \leftrightarrow^* t'$ there is a rewrite proof $t \to^* v \,{}^*\!\leftarrow t'$. A rewrite relation $\to$ has the *diamond property*, if any local peak over $\to$ can be joined immediately, that is, if $\leftarrow \cdot \to \subseteq \to \cdot \leftarrow$ holds.

Descendants and Residuals. Let $(\mathcal{F},\mathcal{R})$ be a TRS and let L be a set of labels. The *labelled TRS* $(\mathcal{F}^L,\mathcal{R}^L)$ is defined by setting (i) $\mathcal{F}^L := \mathcal{F} \cup \{f^\ell \mid f \in \mathcal{F} \text{ and } \ell \in L\}$, (ii) the projection $\langle t\rangle$ of a term $t \in \mathcal{T}(\mathcal{F}^L,\mathcal{V})$ removes all labels, and (iii) $\mathcal{R}^L := \{l \to r \mid \langle l\rangle \to r \in \mathcal{R}\}$. The next proposition is from Terese [12, Proposition 4.2.3].

▶ **Proposition 1.** *Consider a left-linear TRS $(\mathcal{F},\mathcal{R})$ and a set of labels L. Let $s \in \mathcal{T}(\mathcal{F},\mathcal{V})$ and let s' be a labelled term such that $\langle s'\rangle = s$. Then each reduction step $s \to t$ can be lifted to a reduction step $s' \to t'$ in the labelled TRS $(\mathcal{F}^L,\mathcal{R}^L)$ such that $\langle t'\rangle = t$.*

In the following, we write $\mathcal{R}^L$ in short for the labelled TRS $(\mathcal{F}^L,\mathcal{R}^L)$, if the (labelled) signature is clear from context.

▶ **Definition 2.** *Let t be a term in a TRS $\mathcal{R}$, let s be a redex and let f be a function symbol occurring at position p in t, ie. $\mathsf{f} = \mathsf{rt}(t|_p)$. Let t_{f} denote the term that results from t by labelling this occurrence of f with label $\ell \in L$. Then the reduction step $t \xrightarrow{s} t'$ (contracting redex s) is lifted to a reduction step $t_{\mathsf{f}} \to t''$ in $\mathcal{R}^L$.*

The occurrences of f in t' that have label ℓ in t'' are the descendants *of the original symbol occurence of f in t. Conversely, the original f is called the* ancestor *of its descendants.*

The descendant/ancestor relation is extended to subterm occurrences via their root symbols. The descendant of a redex is called a *residual*. For a set of redexes S, we call the set of residuals of redexes in S simply the set of residuals of S. The descendant/ancestor relation naturally generalises to sequence of rewrite steps, that is, derivations. Note that the ancestor relation is unique, that is, for any derivation $D\colon s \to^* t$ the ancestor of a subterm u in t is given as a unique occurrence of a subterm u' in s, if it exists, cf. [12, Chapter 4].

Orthogonality. It is well-known that every orthogonal TRS is confluent, which can for example be verified by repeated applications of the Parallel Moves Lemma, cf. [12, Lemma 4.3.3].

▶ **Lemma 3** (Parallel Moves Lemma)**.** *In an orthogonal TRS, let $t \to^* t_2$ be given. Let $t \xrightarrow{s} t_1$ be a one-step reduction by contraction of redex s. Then a common reduct t_3 of t_1 and t_2 can be found by contracting in t_2 of all residuals of redex s. Observe that all residuals will be pairwise disjoint.*

In order to prove the Parallel Moves Lemma, one makes use of the parallel rewriting relation, formalising the notion of contraction of pairwise disjoint redexes.

▶ **Definition 4.** *Let $\mathcal{R}$ be a TRS. We define the* parallel rewriting *relation $\Rightarrow_\mathcal{R}$ as follows*
1. *$x \Rightarrow_\mathcal{R} x$ for any variable x,*
2. *$f(\vec{s}) \Rightarrow_\mathcal{R} f(\vec{t})$ for any function symbol f, if for all i $s_i \Rightarrow_\mathcal{R} t_i$, and*
3. *$l\sigma \Rightarrow_\mathcal{R} r\sigma$, if $l \to r \in \mathcal{R}$ and σ a substitution.*

We often omit $\mathcal{R}$ and simply write $s \Rightarrow t$, if the TRS is clear from context.

Note that $\to_\mathcal{R} \subseteq \Rightarrow_\mathcal{R} \subseteq \to_\mathcal{R}^*$, in particular we have that $\to_\mathcal{R}^* = \Rightarrow_\mathcal{R}{}^*$. Making use of parallel rewriting, we can state the Parallel Moves Lemma succinctly as follows. A strengthening of the lemma has been stated and proven in [11].

▶ **Lemma 5.** *Parallel rewriting has the diamond property for every orthogonal TRS $\mathcal{R}$, that is, if $t \Leftarrow_\mathcal{R} s \Rightarrow_\mathcal{R} t'$, then there exists a join t'' such that $t' \Rightarrow_\mathcal{R} t'' \Leftarrow_\mathcal{R} t$.*

Let TRS $\mathcal{R}$ be fixed and let $s \Rightarrow t$ denote a paralel rewriting step with respect to $\mathcal{R}$. Suppose the (occurrences of) disjoint redexes contracted are collected in set S. Then we succinctly write $s \overset{S}{\Rightarrow} t$. Due to the Parallel Moves Lemma, we obtain the following proposition, cf. [12, Proposition 4.5.6].

▶ **Proposition 6.** *Let $\mathcal{R}$ be an orthogonal TRS, and let $t \in \mathcal{T}(\mathcal{F}, \mathcal{V})$. Let S, T be sets of pairwise disjoint redexes in t and let $t \overset{S}{\Rightarrow} t'$. Then the set of residuals of T in t' is unique, that is, independent of the order in which redexes in S are contracted.*

Proof. This is a direct consequence of the diamond property of $\Rightarrow$. Actually a stronger results holds. The single parallel rewriting step employed, is generalisable to a complete development step, without affecting the validity of the proposition, cf. [12, Proposition 4.5.6]. ◀

Based on Proposition 6 we denote with $T/\!\overset{S}{}$ the (unique) set of residuals of T in t' that are obtained by the parallel rewriting step $t \overset{S}{\Rightarrow} t'$. With Lemma 3 we observe that T/S consists of pairwise disjoint redexes in t'.

Following the definition of the functions $\mathsf{cvs}_\mathcal{R}$ and $\mathsf{vs}_\mathcal{R}$ in [11], we define functions that compute the worst case of joining derivations based on peaks, resp. equation proofs, of a given size in the most effective way. Let $\|D\|$ denote the number of symbol occurrences in D.

▶ **Definition 7.** *Let $\mathcal{R}$ be an orthogonal term rewrite system. With $\mathsf{j}_\mathcal{R}(t, t')$ we denote the minimal size of a joining derivation of terms t and t', if it exist:*

$$\mathsf{j}_\mathcal{R}(t,t') \;=\; \begin{cases} \min\{\|D'\| \colon D' : t \to_\mathcal{R}^* \cdot {}^*\!\leftarrow_\mathcal{R} t'\} & \text{if } t \text{ and } t' \text{ have a joining derivation} \\ \infty & \text{otherwise} \end{cases}$$

The worst case join complexities *for confluence* Conf *and Church-Rosser* CR *are defined as*

$$\begin{aligned} \mathsf{Conf}(n) &= \max\{\mathsf{j}_\mathcal{R}(t,t') \colon \exists D; \|D\| = n,\ D : t \leftarrow_\mathcal{R}^* \cdot \to_\mathcal{R}^* t',\ \mathcal{R} \text{ orthogonal TRS}\} \\ \mathsf{CR}(n) &= \max\{\mathsf{j}_\mathcal{R}(t,t') \colon \exists D; \|D\| = n,\ D : t \leftrightarrow_\mathcal{R}^* t',\ \mathcal{R} \text{ orthogonal TRS}\}\ . \end{aligned}$$

In the following we will give some (worst-case) upper and (worst-case) lower bounds to those functions. Our main result will be a polynomial upper bound to Conf in Corollary 18. We also provide an exponential lower bound to CR in Corollary 20.

For the remainder of the paper, we restrict to *orthogonal* TRSs.

3 Lower Bounds for Confluence

For our lower bound considerations we use the following big-O facts, which follow easily from definitions.

▶ **Lemma 8.**
1. *If* $e_1(n) = \mathsf{O}(e(n))$ *and* $e_2(n) = \Omega(e(n))$ *then* $e_2(n) = \Omega(e_1(n))$.
2. *If* $e_1(n) = e(n)^{\mathsf{O}(1)}$ *and* $e_2(n) = e(n)^{\Omega(1)}$, *then* $e_2(n) = e_1(n)^{\Omega(1)}$.

We first give a linear lower bound to the number of steps for joining a peak in the size of the splitting sequence. We will provide a corresponding upper bound in Corollary 16.

▶ **Proposition 9.** *There is an orthogonal TRS* $\mathcal{R}$ *satisfying the following: Let* $D_1 \colon a \to^* b$ *and* $D_2 \colon a \to^* c$ *be derivations over* $\mathcal{R}$, *such that* $b \to^k d$, *and* $c \to^l d$ *holds for numbers* k, l, *and term* d. *Then* $k + l = \Omega(\|D_1\| + \|D_2\|)$, *that is,* $k + l$ *is at least linear in the number of symbols in* D_1 *and* D_2 *together.*

Proof. Consider the TRS $\mathcal{R}_1$ given by

$$\mathsf{f}(x) \to \mathsf{g}(x,x) \qquad \mathsf{a}(x) \to \mathsf{b}(x,x)\ . \tag{1}$$

We define meta term symbols via $A(T) := \mathsf{a}(T)$, $B(T) := \mathsf{b}(T,T)$, $F(T) := \mathsf{f}(T)$, $G(T) := \mathsf{g}(T,T)$. For a meta term symbol T let $T^{(n)}$ denote its n-fold iteration.

We define

$$S_n = F^{(n)}(A^{(n)}(0)) \qquad U_n = F^{(n)}(B^{(n)}(0))$$
$$V_n = G^{(n)}(A^{(n)}(0)) \qquad W_n = G^{(n)}(B^{(n)}(0))\ ,$$

and compute

$$|S_n| = \mathsf{O}(n) \qquad |U_n| = \mathsf{O}(2^n) \qquad |V_n| = \mathsf{O}(n2^n)\ .$$

Consider the following peak in $\mathcal{R}_1$, rewriting innermost redexes first.

$$D_1 \colon\ S_n \xrightarrow{\mathsf{a}} F^{(n)}(A^{(n-1)}(B(0))) \xrightarrow{\mathsf{a}} F^{(n)}(A^{(n-2)}(B^{(2)}(0))) \xrightarrow{\mathsf{a}} \cdots \xrightarrow{\mathsf{a}} U_n$$
$$D_2 \colon\ S_n \xrightarrow{\mathsf{f}} F^{(n-1)}(G(A^{(n)}(0))) \xrightarrow{\mathsf{f}} F^{(n-2)}(G^{(2)}(A^{(n)}(0))) \xrightarrow{\mathsf{f}} \cdots \xrightarrow{\mathsf{f}} V_n\ .$$

To discern ambiguity, we have identified the root symbol of the redex above the rewrite relation.

The size of each term in the first derivation is $\mathsf{O}(2^n)$, hence the overall size of D_1 is $\mathsf{O}(n2^n)$. The size of the k-th term in the second derivation is $\mathsf{O}(n2^k)$, so adding them up for $k \leqslant n$ gives a bound of $\mathsf{O}(n2^n)$ for the overall derivation length of D_2 as well. Hence $(\|D_1\| + \|D_2\|) = \mathsf{O}(n2^n)$.

The 'fastest' join of U_n and V_n is given by rewriting innermost redexes first:

$$U_n \xrightarrow{\mathsf{f}}{}^{1} F^{(n-1)}(G(B^{(n)}(0))) \xrightarrow{\mathsf{f}}{}^{1} F^{(n-2)}(G^{(2)}(B^{(n)}(0))) \xrightarrow{\mathsf{f}}{}^{1} \cdots \xrightarrow{\mathsf{f}}{}^{1} W_n$$
$$V_n \xrightarrow{\mathsf{a}}{}^{2^n} G^{(n)}(A^{(n-1)}(B(0))) \xrightarrow{\mathsf{a}}{}^{2^n} G^{(n)}(A^{(n-2)}(B^{(2)}(0))) \xrightarrow{\mathsf{a}}{}^{2^n} \cdots \xrightarrow{\mathsf{a}}{}^{2^n} W_n\ .$$

The length of the first derivation is n, and of the second $n2^n$, respectively.

Thus, a lower bound to the number of steps S_{join} of any derivations that join U_n and V_n is $n2^n$: $S_{\text{join}} = \Omega(n2^n)$. Together with $(\|D_1\| + \|D_2\|) = \mathsf{O}(n2^n)$ and Lemma 8.(1), we obtain $S_{\text{join}} = \Omega(\|D_1\| + \|D_2\|)$. Hence, S_{join} must be at least linear in the size of the derivations D_1 and D_2 constituting the peak. ◀

We also give a linear lower bound to the size of the join of the diamond in the product of the sizes of meet-able terms in a peak. The corresponding upper bound will be given in Corollary 15.

▶ **Proposition 10.** *There is an orthogonal TRS $\mathcal{R}$ satisfying the following: Let $b \,{}^*\!\leftarrow a \to^* c$ be a peak over $\mathcal{R}$ with consequent join d such that $b \to^* d$ and $c \to^* d$. Then $|d| = \Omega(|b|\cdot|c|)$, that is, the size $|d|$ of d is at least linear in $|b|\cdot|c|$.*

Proof. Fix n. We will basically follow the example from the proof of Proposition 9, with a slight modification to obtain optimal bounds.

With the notation from the proof of Proposition 9, expand TRS $\mathcal{R}_1$, cf. (1), with the rule $\mathsf{h} \to A^{(n)}(\mathsf{0})$. Let the resulting TRS be denoted as $\mathcal{R}_2$. We define

$$S'_n = F^{(n)}(\mathsf{h}) \qquad U_n = F^{(n)}(B^{(n)}(\mathsf{0}))$$
$$V'_n = G^{(n)}(\mathsf{h}) \qquad W_n = G^{(n)}(B^{(n)}(\mathsf{0}))\ ,$$

and compute

$$|U_n| = \mathsf{O}(2^n) \qquad |V'_n| = \mathsf{O}(2^n) \qquad |W_n| = \Omega(2^{2n})\ .$$

Consider the following peak:

$$S'_n \xrightarrow{\mathsf{h}} F^{(n)}(A^{(n)}(\mathsf{0})) \xrightarrow{\mathsf{a}}{}^* U_n$$
$$S'_n \xrightarrow{\mathsf{f}} F^{(n-1)}(G(\mathsf{h})) \xrightarrow{\mathsf{f}} F^{(n-2)}(G^{(2)}(\mathsf{h})) \xrightarrow{\mathsf{f}}{}^* V'_n\ .$$

The 'smallest' join of U_n and V_n is given by rewriting only residuals:

$$U_n \xrightarrow{\mathsf{f}}{}^* W_n$$
$$V'_n \xrightarrow{\mathsf{h}}{}^* G^{(n)}(A^{(n)}(\mathsf{0})) \xrightarrow{\mathsf{a}}{}^* W_n\ .$$

We compute $|U_n|\cdot|V'_n| = \mathsf{O}(2^{2n})$. Together with $|W_n| = \Omega(2^{2n})$ and (1) we obtain $|W_n| = \Omega(|U_n|\cdot|V'_n|)$. Hence, the size of any join must be at least linear in the product of the sizes of U_n and V'_n. ◀

4 Injectivity

For the sequel, we fix an orthogonal TRS $\mathcal{R}$. Let $t' \,{}^*\!\leftarrow s \to^* t$ denote a peak over $\mathcal{R}$.

Consider the tiling diagramme in Figure 2 obtained by repeated applications of Lemma 5. We assume that $H_{0,\nu}$ denotes a singleton set of one redex in $s_{0,\nu}$, for $\nu = 0\ldots,i{-}1$, and that $V_{\mu,0}$ denotes a singleton set of one redex in $s_{\mu,0}$, for $\mu = 0\ldots,j{-}1$. Note that this implies $|H_{0,\nu}| = 1$ and $|V_{\mu,0}| = 1$. Further, we obtain

$$V_{\mu,\nu+1} = V_{\mu,\nu}/H_{\mu,\nu} \qquad H_{\mu+1,\nu} = H_{\mu,\nu}/V_{\mu,\nu}\ ,$$

as sets of residuals using Proposition 6. Moreover, using Proposition 6, we have that $H_{\mu,\nu}$ and $V_{\mu,\nu}$ are sets of pairwise disjoint redexes in $s_{\mu,\nu}$, for all $\mu = 0\ldots,j{-}1$, $\nu = 0\ldots,i{-}1$. Recall that a *redex* is an occurrence of a term t that is an instance of the left-hand side l of a rule $l \to r \in \mathcal{R}$.

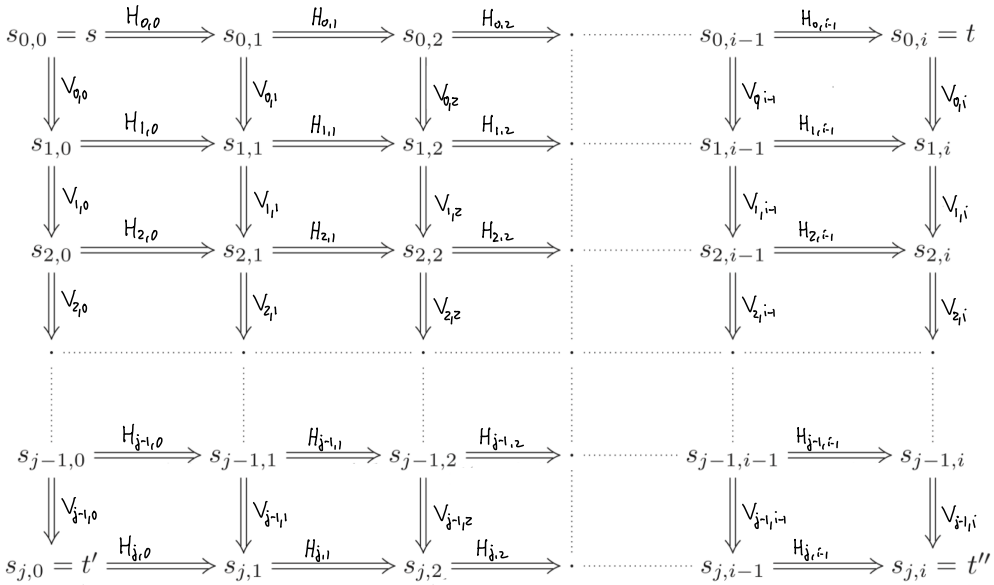

Figure 2 The tiling situation.

Generalised Ancestors

Given a sequence of rewrite steps

$$t \to_{s'} t' \to_{s''} t'' \to \ldots \to_{s^{(n-1)}} t^{(n-1)} \to_{s^{(n)}} t^{(n)}$$

we generalise the notion of ancestor to trace any subterm in the sequence back to t – we denote this *generalised ancestor*, or short *g.-ancestor*.

Ancestors are also g.-ancestors. Consider a subterm u_j in $t^{(j)}$, and its ancestors u_{j-1} in $t^{(j-1)}$, etc., until u_i in $t^{(i)}$ cannot be extended any further. Let f denote the root symbol of u_i in $t^{(i)}$. As f does not have an ancestor in $t^{(i-1)}$, we must be in the following situation: There exist a context $C[*]$, substitution σ, and rule $l \to r$ in $\mathcal{R}$, such that $t^{(i-1)} = C[l\sigma]$, $t^{(i)} \equiv C[r\sigma]$, and f occurs in r. We now define the *generalised ancestor* of f in $t^{(i)}$ as the root symbol of l in $C[l\sigma] = t^{(i-1)}$. Continue until t is reached.

▶ **Proposition 11.** *In the tiling diagramme in Figure 2, the generalised ancestors of any symbol occurrence are unique, that is, independent of the path chosen to compute them.*

Proof. Arguing inductively, it suffices to prove the statement for a single square:

$$\begin{array}{ccc} s_{\mu,\nu} & \overset{H_{\mu,\nu}}{\Longrightarrow} & s_{\mu,\nu+1} \\ \Downarrow{\scriptstyle V_{\mu,\nu}} & & \Downarrow{\scriptstyle V_{\mu,\nu+1}} \\ s_{\mu+1,\nu} & \overset{H_{\mu+1,\nu}}{\Longrightarrow} & s_{\mu+1,\nu+1}\,. \end{array}$$

Recall that using Proposition 6, we have that $H_{\mu,\nu}$ and $V_{\mu,\nu}$ are sets of disjoint redexes in $s_{\mu,\nu}$, for all $\mu = 0\ldots,j{-}1$, $\nu = 0\ldots,i{-}1$. Thus, in proof of the claim, we can assume without loss of generality that $|H_{\mu,\nu}| = |V_{\mu,\nu}| = 1$.

Let u be a subterm of $s_{\mu+1,\nu+1}$. First, suppose u has an *ancestor* in $s_{\mu,\nu}$. Then, this ancestor is unique, as mentioned above.

Second, suppose u has only *generalised ancestors* in $s_{\mu,\nu}$. Then, we distinguish cases on the relative positioning of redexes in $H_{\mu,\nu}$ and $V_{\mu,\nu}$, respectively. Recall, that by assumption the redexes in $H_{\mu,\nu}$ and $V_{\mu,\nu}$ are pairwise disjoint.

Case. Suppose $H_{\mu,\nu} \| V_{\mu,\nu}$, that is, the redexes in $H_{\mu,\nu} \cup V_{\mu,\nu}$ are all pairwise disjoint. Then the claim is obvious.

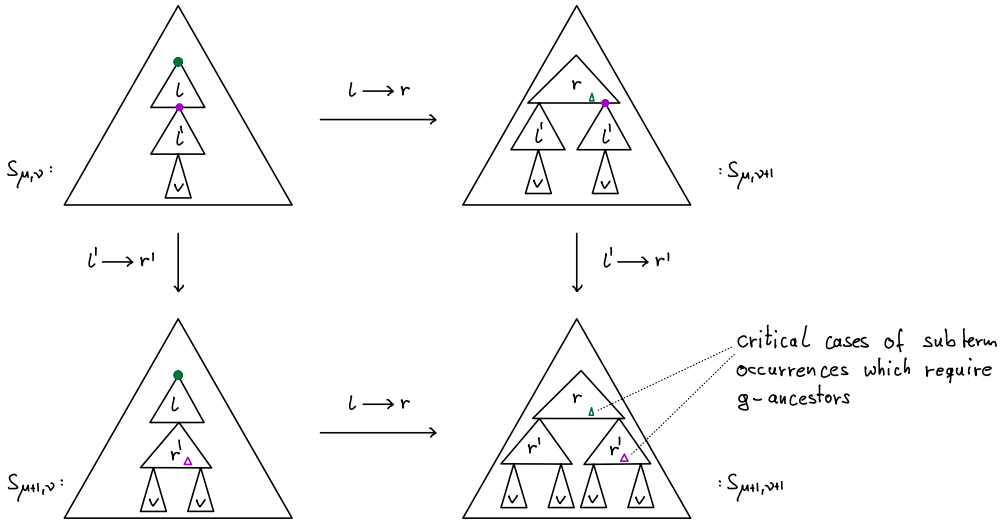

Figure 3 Critical cases where generalised ancestors occur.

Case. Suppose there exists rules $l \to r, l' \to r' \in \mathcal{R}$, and substitutions σ, σ' such that $l\sigma \in H_{\mu,\nu}$ and $l'\sigma' \in V_{\mu,\nu}$. Further $l'\sigma' \lhd l\sigma$. (The case $l\sigma = l'\sigma$ is trivial, because we must have $(l \to r) = (l' \to r')$ due to orthogonality of $\mathcal{R}$.) As u does not have an ancestor in $s_{\mu,\nu}$, $\mathsf{rt}(u)$ either occurs in r or in r'. The situation of this case is depicted in Figure 3.

Wlog. $\mathsf{rt}(u)$ occurs in r' and thus u occurs in any of the occurrences of $r'\sigma'$ in $s_{\mu+1,\nu+1}$. By assumption on $l\sigma$ and $l'\sigma'$, u has an ancestor in $s_{\mu+1,\nu}$ and a generalised ancestor in $s_{\mu,\nu+1}$, which are both unique and consequently their join in $s_{\mu,\nu}$ is unique, too. ◀

▶ **Definition 12.** *Let the tiling diagramme in Figure 2 be given, and let $\mu < j$, $\nu < i$. Let f be a function symbol occurrence in $s_{\mu,\nu}$, and let $\mu' \leqslant \mu$, $\nu' \leqslant \nu$. We define $\mathrm{ga}^{\mu,\nu}_{\mu',\nu'}(\mathsf{f})$ as the g.-ancestor of f in $s_{\mu',\nu'}$.*

We now formulate the main result of this section.

▶ **Lemma 13.** *Let the tiling diagramme in Figure 2 be given, and let $\mu < j$, $\nu < i$, and $\mu' \leqslant \mu$, $\nu' \leqslant \nu$. The mapping of function symbol occurrences f in $s_{\mu,\nu}$ to the pair $(\mathrm{ga}^{\mu,\nu}_{\mu,\nu'}(\mathsf{f}), \mathrm{ga}^{\mu,\nu}_{\mu',\nu}(\mathsf{f}))$ is an injection.*

Proof. This claim can be proven by induction on $\nu - \nu'$. The case for $\nu = \nu'$ is obvious, because $\mathrm{ga}^{\mu,\nu}_{\mu,\nu}$ is the identity, which is injective.

For the induction step from $\nu' + 1$ to ν' we can assume by induction hypothesis that the claim is true for $(\mu', \nu' + 1)$. We then show the claim for (μ', ν'), depicted as follows.

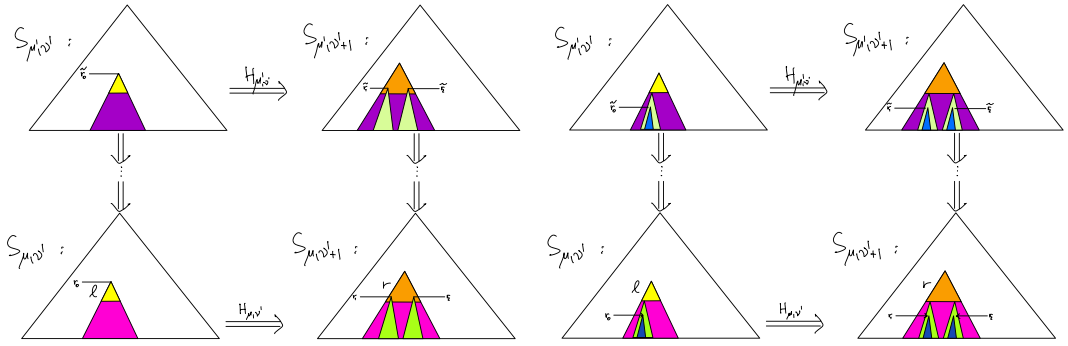

(a) Case A. **(b)** Case B.

Figure 4 Cases A and B in proof of Lemma 13.

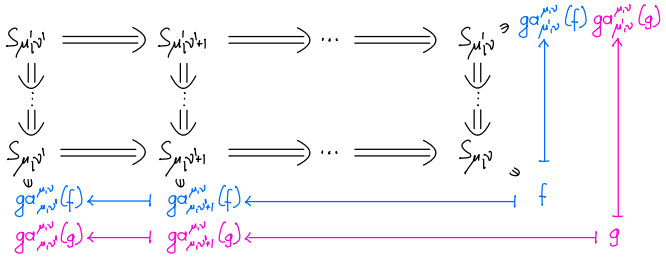

For sake of contradiction assume the claim is wrong for (μ', ν'). That is, there are f, g occurring in $s_{\mu,\nu}$ with f, g different symbol occurrences, such that $\mathrm{ga}^{\mu,\nu}_{\mu',\nu}(\mathsf{f}) = \mathrm{ga}^{\mu,\nu}_{\mu',\nu}(\mathsf{g})$ and $\mathrm{ga}^{\mu,\nu}_{\mu,\nu'}(\mathsf{f}) = \mathrm{ga}^{\mu,\nu}_{\mu,\nu'}(\mathsf{g})$. By i.h. we must have $\mathrm{ga}^{\mu,\nu}_{\mu,\nu'+1}(\mathsf{f}) \neq \mathrm{ga}^{\mu,\nu}_{\mu,\nu'+1}(\mathsf{g})$. Let $r_1 = \mathrm{ga}^{\mu,\nu}_{\mu,\nu'+1}(\mathsf{f})$, $r_2 = \mathrm{ga}^{\mu,\nu}_{\mu,\nu'+1}(\mathsf{g})$, and $r_0 = \mathrm{ga}^{\mu,\nu}_{\mu,\nu'}(\mathsf{f}) = \mathrm{ga}^{\mu,\nu}_{\mu,\nu'}(\mathsf{g})$. This situation is depicted below.

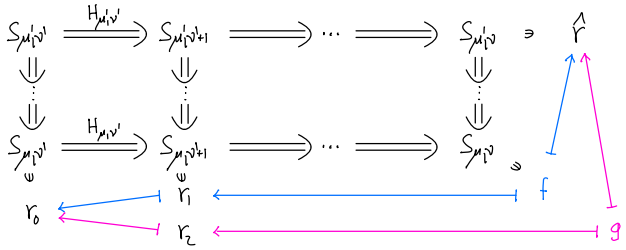

We must be in the following situation: There are rule $l \to r$ in $\mathcal{R}$, substitution ρ, terms $u_1, \ldots, u_k$, context $C[*_1, \ldots, *_k]$, such that $H_{\mu,\nu'} = \{u_1, \ldots, u_k\}$, $u_1 = l\rho$, $s_{\mu,\nu'} = C[u_1, \ldots, u_k]$, and r_1 and r_2 occur in $r\rho$ in $s_{\mu,\nu'+1} = C[r\rho, \ldots]$, and either

- the roots of r_1 and r_2 occur already in r in $C[r\rho, \ldots]$, hence their joint g.-ancestor r_0 is the root of l in $C[l\rho, u_2, \ldots, u_k]$, see Figure 4a;
- or we have a variable x occuring in l which occurs multiple times in r, e.g. as $C_r[*_1, *_2]$ with $r = C[x, x]$ – hence $r\rho = C_r\rho[x\rho, x\rho]$ – and r_1 occurs in the first $x\rho$, r_2 occurs in the second $x\rho$, and their joint ancestor r_0 occurs in $x\rho$ in $l\rho$ in $s_{\mu,\nu'}$, see Figure 4b.

Let $\hat{r} = \mathrm{ga}^{\mu,\nu}_{\mu',\nu}(\mathsf{f}) = \mathrm{ga}^{\mu,\nu}_{\mu',\nu}(\mathsf{g})$ be the g.-ancestor of f and g in $s_{\mu',\nu}$. $H_{\mu,\nu'}$ are residuals of $H_{\mu',\nu'}$, hence the ancestors $\tilde{r}_0$ of r_0 in $s_{\mu',\nu'}$ and $\tilde{r}_1, \tilde{r}_2$ of r_1, r_2 in $s_{\mu',\nu'+1}$ will occur in $l\rho'$ and $r\rho'$ for some ρ'. In particular in A), the roots of $\tilde{r}_1$ and $\tilde{r}_2$ are in r, and $\tilde{r}_0$ is at the root of l. In case B) we have that $r\rho' = C_r\rho'[x\rho', x\rho']$ with $\tilde{r}_1$ occuring in 1st and $\tilde{r}_2$ in 2nd of $x\rho'$.

In both cases we have that $\tilde{r}_1$ and $\tilde{r}_2$ are two distinct g.-ancestors of f and g in $s_{\mu',\nu'+1}$, resp., by following from $s_{\mu,\nu}$ the derivation first to $s_{\mu,\nu'+1}$ and then to $s_{\mu',\nu'+1}$. However, by following from $s_{\mu,\nu}$ the derivation to $s_{\mu',\nu}$, f and g have a joint ancestor $\hat{r}$, hence can only have one joint ancestor in $s_{\mu',\nu'+1}$ when following the derivation from $s_{\mu',\nu}$ to $s_{\mu',\nu'+1}$ to the left. This contradicts Proposition 11 that g.-ancestors are unique. ◀

▶ **Lemma 14.** *Let the tiling diagramme in Figure 2 be given, and let $\mu < j$, $\nu < i$.*

Assuming $|H_{0,\nu}| = 1$, the mapping of each redex in $H_{\mu,\nu}$ to their generalised ancestors in $s_{\mu,\nu'}$ for $\nu' < \nu$ is an injection.

Similar for $V_{\mu,\nu}$: Assuming $|V_{\mu,0}| = 1$, the mapping of each redex in $V_{\mu,\nu}$ to their generalised ancestors in $s_{\mu',\nu}$ for $\mu' < \mu$ is an injection.

Proof. We only consider the first assertion, the second is dual. Ie., we are in the following situation.

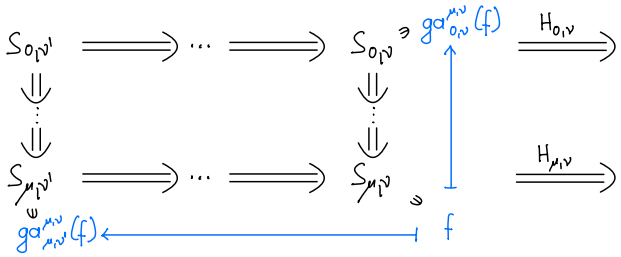

Let s be a term, H a set of redexes in s, and f a function symbol occurrence in s. We succinctly write $\mathsf{f} \in H$ to indicate that f is the occurrence of the root symbol of some redex in H.

By Lemma 13 we have that the mapping

$$\mathsf{f} \in H_{\mu,\nu} \mapsto (\mathrm{ga}^{\mu,\nu}_{\mu,\nu'}(\mathsf{f}), \mathrm{ga}^{\mu,\nu}_{0,\nu}(\mathsf{f}))$$

is an injection. By assumption we have that $|H_{0,\nu}| = 1$, hence $H_{0,\nu} = \{\hat{r}\}$ for some $\hat{r}$. This implies that $\mathrm{ga}^{\mu,\nu}_{0,\nu}(\mathsf{f}) = \hat{r}$ for all $\mathsf{f} \in H_{\mu,\nu}$. Hence

$$\mathsf{f} \in H_{\mu,\nu} \mapsto \mathrm{ga}^{\mu,\nu}_{\mu,\nu'}(\mathsf{f})$$

must be injective. ◀

5 Upper Bounds on Confluence

In this short section, we state and prove our main result that the size, that is, the number of symbols, of a rewrite proof is polynomial in the size of the peak, cf. Figure 1. First, we draw two easy corollaries from Lemma 13 and Lemma 14, respectively.

▶ **Corollary 15.** *Consider the tiling diagramme in Figure 2. The size of the join t'' is bounded by the product of the sizes of t and t':*

$$|t''| \quad \leqslant \quad |t| \cdot |t'| \,.$$

Proof. This is a direct consequence of Lemma 13. ◀

▶ **Corollary 16.** *Consider the tiling diagramme in Figure 2, assuming* $|H_{0,\nu}| = 1$ *and* $|V_{\mu,0}| = 1$. *In this situation, the number of (sequential) reduction steps needed to join* t *and* t' *via* t'', *is bounded by the square of the size of the initial sequence. More precisely:*

$$\sum_{\nu=0}^{i-1}|H_{j,\nu}| \;+\; \sum_{\mu=0}^{j-1}|V_{\mu,i}| \;\leqslant\; i\cdot|t'|+j\cdot|t| \;\leqslant\; \big(\sum_{\mu=0}^{j}|s_{\mu,0}|+\sum_{\nu=1}^{i}|s_{0,\nu}|\big)^2 .$$

Proof. For the first inequality, observe that by Lemma 14, we have that $|H_{j,\nu}| \leq |s_{j,0}|$ for $\nu < i$ and $|V_{\mu,i}| \leq |s_{0,i}|$ for $\mu < j$. Thus, $|H_{j,\nu}| \leq |t'|$ and $|V_{\mu,i}| \leq |t|$ by definition. Then, the second inequality follows by elementary calculations. Finally, observe that if the set of redex S is disjoint then $\overset{S}{\Longrightarrow} \subseteq \to^S$, from which the claim follows. ◀

Now, our main result follows with ease.

▶ **Theorem 17.** *Let* $\mathcal{R}$ *be an orthogonal TRS and assume the existence of a peak* $D\colon t' \,{}^*\!\leftarrow s \to^* t$. *Then there exists a rewriting proof* $D'\colon t' \to^* t'' \,{}^*\!\leftarrow t$ *whose size is polynomially bounded in the size of* D. *In fact, the size of* D' *is biquadratic in the size of* D.

Proof. This is a consequence of Corollaries 15 and 16. Let D' be the joining derivation given by the tiling diagram in Figure 2, where $s_{0,0} = s$, $s_{0,\nu}$ is the ν-th term in $s \to^i t$, and $s_{\mu,0}$ the μ-th term in $s \to^j t'$. Employing the notation of that figure, we obtain

$$\|D\| \;=\; \sum_{\mu=0}^{j}|s_{\mu,0}|+\sum_{\nu=1}^{i}|s_{0,\nu}| .$$

Recall that $\|D\|$ denotes the number of symbol occurrences in D. Due to Corollary 15, we have, for each μ, ν $(0 \leqslant \mu \leqslant j,\ 0 \leqslant \nu \leqslant i)$, that

$$|s_{\mu,\nu}| \;\leqslant\; |s_{\mu,0}|\cdot|s_{0,\nu}| \;\leqslant\; \|D\|^2 . \tag{2}$$

Moreover, due to Corollary 16, the number of joining steps in D' is bounded by $\|D\|^2$:

$$\text{number of joining steps} \;\leqslant\; \sum_{\nu=0}^{i-1}|H_{j,\nu}| + \sum_{\mu=0}^{j-1}|V_{\mu,i}| \;\leq\; \|D\|^2 . \tag{3}$$

Combining (2) and (3), we conclude that $\|D'\| \leqslant \|D\|^4$. ◀

▶ **Corollary 18.** Conf *is biquadratically bounded, i.e.* $\mathsf{Conf}(n) = \mathsf{O}(n^4)$.

A closer inspection of the example in the proof of Proposition 10 establishes a cubic lower bound, i.e. $\mathsf{Conf}(n) = \Omega(n^3)$.

6 Lower and Upper Bounds for the Church-Rosser Property

In the case of the Church-Rosser property, we first give an exponential lower bound to the size of the join, which in particular gives an exponential lower bound to the join complexity CR.

▶ **Theorem 19.** *There is an orthogonal TRS* $\mathcal{R}$ *satisfying the following: Let* D *be a derivation of* $a \leftrightarrow^* b$ *over* $\mathcal{R}$, *such that* $a \to^* c$ *and* $b \to^* c$ *holds, then* $|c|$ *is exponential in* $\|D\|$ *in general, i.e.* $|c| = 2^{\|D\|^{\Omega(1)}}$.

Proof. Consider the TRS $\mathcal{R}_3$ given by

$$\mathsf{f}_i(x) \to \mathsf{a}_i(x,x) \quad \mathsf{g}_i(x) \to \mathsf{a}_i(x,x) \qquad (i = 1,\dots,k)\,. \tag{4}$$

We define meta term symbols via $A_i(T) := \mathsf{a}_i(T,T)$, define

$$S_i^k = \mathsf{g}_1(\dots \mathsf{g}_{i-1}(\mathsf{g}_i(\mathsf{f}_{i+1}(\dots \mathsf{f}_k(0)\dots)))\dots) \qquad U^k = A_1(\dots A_k(0)\dots)$$
$$T_i^k = \mathsf{g}_1(\dots \mathsf{g}_{i-1}(A_i(\mathsf{f}_{i+1}(\dots \mathsf{f}_k(0)\dots)))\dots)\,,$$

and compute

$$|S_i^k| = \mathsf{O}(k) \qquad |T_i^k| = \mathsf{O}(k) \qquad S_i^k \xrightarrow{g_i} T_i^k \qquad S_i^k \xrightarrow{f_{i+1}} T_{i+1}^k\,.$$

Consider the following derivation:

$$D \quad := \quad T_1^k \leftarrow S_1^k \to T_2^k \leftarrow S_2^k \to T_3^k \;\dots\; T_{k-1}^k \leftarrow S_{k-1}^k \to T_k^k$$

The unique Church-Rosser join is given by $T_i^k \to^* U$ for all $i = 1,\dots,k$. From now on we drop the superscript k.

Let $S_D = \|D\|$ and $S_U = |U|$. We compute $S_D = \mathsf{O}(n^2)$ and $S_U = \Omega(2^n)$. Thus $S_D \leqslant ck^2$ for some $c > 0$, hence $k \geqslant \sqrt{\frac{1}{c}S_D} \geqslant S_D{}^{\epsilon}$ for small $\epsilon > 0$. Thus $S_U \geqslant 2^k \geqslant 2^{S_D{}^{\epsilon}}$. ◀

▶ **Corollary 20.** $\mathsf{CR}(n)$ *is exponential in* n*, i.e.* $\mathsf{CR}(n) = 2^{n^{\Omega(1)}}$.

Inspecting our upper bounds, Corollaries 15 and 16, establishes that this bound is optimal up to the degree, i.e. $\mathsf{CR}(n) = 2^{n^{\mathsf{O}(1)}}$.

We now show that the size of the join in the case of Church-Rosser is polynomially related to the product of the sizes of the terms in the starting derivation. We first state the lower bound.

▶ **Proposition 21.** *There is an orthogonal TRS* $\mathcal{R}$ *satisfying the following: Let* $a_1 \leftrightarrow a_2 \leftrightarrow \cdots \leftrightarrow a_k$ *be a derivation over* $\mathcal{R}$ *such that* $a_1 \to^* b$ *and* $a_k \to^* b$ *for some* b*. Then* $|b|$ *is polynomial in* $|a_1|\cdot|a_2|\cdot\dots\cdot|a_k|$ *in general, i.e.* $|b| = (|a_1|\cdot|a_2|\cdot\dots\cdot|a_k|)^{\Omega(1)}$.

Proof. We modify the TRS from the previous proof so that the starting terms are of constant size: Expand the TRS from the proof of Theorem 19 by

$$\bar{\mathsf{f}}_i^k \to \mathsf{f}_i(\bar{\mathsf{f}}_{i+1}^k) \quad \bar{\mathsf{g}}_i^k(x) \to \bar{\mathsf{g}}_{i-1}^k(\mathsf{g}_i(x)) \qquad (i = 1,\dots,k) \tag{5}$$

where $\bar{\mathsf{f}}_{k+1}^k$ represents 0. We define

$$\bar{S}_i^k = \bar{\mathsf{g}}_i^k(\bar{\mathsf{f}}_{i+1}^k) \qquad \bar{T}_i^k = \bar{\mathsf{g}}_{i-1}^k(A_i(\bar{\mathsf{f}}_{i+1}^k))\,,$$

and compute

$$|\bar{S}_i^k| = \mathsf{O}(1) \qquad \bar{S}_i^k = \bar{\mathsf{g}}_i^k(\bar{\mathsf{f}}_{i+1}^k) \xrightarrow{\bar{\mathsf{g}}_i^k} \bar{\mathsf{g}}_{i-1}^k(\mathsf{g}_i(\bar{\mathsf{f}}_{i+1}^k)) \xrightarrow{\mathsf{g}_i} \bar{\mathsf{g}}_{i-1}^k(A_i(\bar{\mathsf{f}}_{i+1}^k)) = \bar{T}_i^k$$
$$|\bar{T}_i^k| = \mathsf{O}(1) \qquad \bar{S}_i^k = \bar{\mathsf{g}}_i^k(\bar{\mathsf{f}}_{i+1}^k) \xrightarrow{\bar{\mathsf{f}}_{i+1}^k} \bar{\mathsf{g}}_i^k(\mathsf{f}_{i+1}(\bar{\mathsf{f}}_{i+2}^k)) \xrightarrow{\mathsf{f}_{i+1}} \bar{\mathsf{g}}_i^k(A_{i+1}(\bar{\mathsf{f}}_{i+2}^k)) = \bar{T}_{i+1}^k\,.$$

From now on we will drop the superscript k. Consider the following derivation:

$$\bar{D} \quad := \quad \bar{T}_1 \leftarrow^2 \bar{S}_1 \to^2 \bar{T}_2 \leftarrow^2 \bar{S}_2 \to^2 \bar{T}_3 \;\dots\; \bar{T}_{k-1} \leftarrow^2 \bar{S}_{k-1} \to^2 \bar{T}_k\,.$$

The unique Church-Rosser join is again given by $\bar{T}_i \to^* r$ for all $i = 1,\dots,k$.

Let $\bar{S} = \Pi_{t\in\bar{D}}|t|$ and $S_r = |r|$. We compute $\bar{S} = c^{2k}$ for some $c = \mathsf{O}(1)$ which is an upper bound on the size of terms occurring in $\bar{D}$. Hence $\bar{S} = (2^k)^{\mathsf{O}(1)}$. We also have $S_r = (2^k)^{\Omega(1)}$. Hence $S_r = \bar{S}^{\Omega(1)}$ using Lemma 8(2). ◀

We also have a corresponding upper bound.

▶ **Theorem 22.** *Let $\mathcal{R}$ be an orthogonal TRS. Given a derivation $a_1 \leftrightarrow a_2 \leftrightarrow \cdots \leftrightarrow a_k$ over $\mathcal{R}$, there is a join $a_1 \to^* b \,{}^*\!\leftarrow a_k$ for some b, such that $|b|$ is bounded by $|a_1|\cdot|a_2|\cdot\cdots\cdot|a_k|$.*

Proof. The upper bound is obtained by induction on k using the related upper bound for confluence, Corollary 15: Assume $a_1 \leftrightarrow \cdots \leftrightarrow a_k \leftrightarrow a_{k+1}$. By induction hypothesis there are some b, $a_1 \to^* b$ and $a_k \to^* b$ such that $|b|$ is bounded by $|a_1|\cdot|a_2|\cdot\cdots\cdot|a_k|$. If $a_{k+1} \to a_k$ then b is also the join for a_1 and a_{k+1} and we are already done. Otherwise, $a_k \to a_{k+1}$. Using that $a_k \to^* b$, we can join this peak with some c of size $\leqslant |b|\cdot|a_{k+1}|$ using Corollary 15. Thus $|c| \leqslant |b|\cdot|a_{k+1}| \leqslant |a_1|\cdot|a_2|\cdot\cdots\cdot|a_k|\cdot|a_{k+1}|$. ◀

7 A Lower Bound for the Lambda Calculus

For this section, we assume(at least nodding) acquaintance with the (untyped) λ-calculus [3, 4]. While we refrain from re-stating (hopefully) well-known notions, the result should be easy to understand.

We show that for confluence in λ-calculus, the size of the join is exponential in the product of the sizes of the starting terms in general.

▶ **Proposition 23.** *Given a peak $D\colon b \leftarrow^*_\lambda a \to^*_\lambda c$, and a joining derivation $b \to^*_\lambda d \leftarrow^*_\lambda c$. Then $|d|$ is exponential in $\|D\|$ as well as in $|b|\cdot|c|$ in general: $|d| = 2^{\|D\|^{\Omega(1)}}$ and $|d| = 2^{(|b|\cdot|c|)^{\Omega(1)}}$.*

Proof. Let f, g, h, x, y be variables. Let $A := \lambda x.((\lambda y.hyy)(gx))$ and $B := \lambda x.(h(gx)(gx))$. We have $A \xrightarrow{\lambda y}_\lambda B$, $|A| = \Theta(1)$, $|B| = \Theta(1)$.

Define terms T^k, U^k, V^k, W^k as follows: Let $T^0 = U^0 = V^0 = W^0 = f$, and inductively

$$T^{k+1} = (A\,T^k), \quad U^{k+1} = (B\,U^k), \quad V^{k+1} = (\lambda y.hyy)(gV^k), \quad W^{k+1} = h(gW^k)(gW^k)\ .$$

Then $|T^k| = \mathsf{O}(k)$, $|U^k| = \mathsf{O}(k)$, $|V^k| = \mathsf{O}(k)$, and $|W^k| = \Omega(2^k)$. We have

$$T^k \xrightarrow{\lambda y}{}^k_\lambda U^k \qquad T^k \xrightarrow{\lambda x}{}^k_\lambda V^k \qquad U^k \xrightarrow{\lambda x}{}^k_\lambda W^k \qquad V^k \xrightarrow{\lambda y}{}^k_\lambda W^k$$

by induction on k. Let D be $U^k \leftarrow^*_\lambda T^k \to^*_\lambda V^k$. Then $\|D\| = \mathsf{O}(k^2)$, hence $k \geqslant (\|D\|)^\epsilon$ for some $\epsilon > 0$, hence $|W^k| = \Omega(2^k) = \Omega(2^{(\|D\|)^\epsilon})$. As $|b|\cdot|d| = \mathsf{O}(k^2)$ as well, the same calculation applies in this case as well. ◀

8 Related Works

Ketema and Grue Simonsen have studied similar properties in [11]. For a given TRS $\mathcal{R}$, they define functions $\mathsf{cvs}_\mathcal{R}$ and $\mathsf{vs}_\mathcal{R}$, estimating the least number of reduction steps necessary in a rewrite proof, assuming an equational proof or a peak, respectively. More precisely, $\mathsf{cvs}_\mathcal{R}(m, n)$ denotes the least number of reduction steps required to complete a rewrite proof, given an equational proof involving at most n steps between two terms t, t' of size at most m. Likewise, $\mathsf{vs}_\mathcal{R}(m, n)$ denotes the least number of reduction steps in a rewrite proof, given a peak $t \,{}^*\!\leftarrow s \to^* t'$, where the size of s is at most m and the reduction lengths are at most of size n. For orthogonal TRSs $\mathcal{R}$ they obtain optimal exponential upper bound on $\mathsf{vs}_\mathcal{R}$ and an upper bound on $\mathsf{cvs}_\mathcal{R}$ that belongs to the 4^{th}-level of the Grzegorczyk hierarchy. I.e. the upper bound on $\mathsf{cvs}_\mathcal{R}$ is at least non-elementary. Wrt. the λ-calculus, confluence already requires an non-elementary upper bound. In subsequent work, Fujita proved that for the λ-calculus $\mathsf{cvs}_\mathcal{R}$ is upper bounded in the 4^{th}-level of the Grzegorczyk hierarchy, cf. [10]. Only optimality of the bound on $\mathsf{vs}_\mathcal{R}$ for orthogonal rewrite systems has been established.

We emphasise that these results are orthogonal to our contributions, as we make use of a different notion of proof complexity: the number of symbols, rather than the number of reduction steps. While this measure is natural in the context of rewriting (or even the λ-calculus), it is less so in the context of computational complexity, from our point of view. In short, for orthogonal TRSs, this change allows us to provide (optimal) polynomial upper bounds on confluence proofs and (optimal) exponential upper bounds on Church-Rosser proofs, while we establish an exponential lower bound on confluence proofs for the λ-calculus. Note that our changed notion of size not only allows tractable upper bounds, but also differentiates precisely between the expressivity of (first-order) term rewrite systems and (higher-order) λ-calculus, a difference that got somewhat blurred in related works.

To the best of our knowledge, confluence or Church-Rosser properties in term-rewriting have not been studied in general in Bounded Arithmetic (though they have been used as tools in the analysis of related artefacts, as in work by Das [9]). The closest we are aware of are the results by the first author [5] that formalises a restricted and very involved property the resembles elements of Church-Rosser, and which are used to prove the consistency of any equational theory that exclusively is based on recursive defining equations, in a weak theory of bounded arithmetic. These results were improved by Yamagata [13] by also allowing rules for substituting terms into equations in the equational reasoning while proving consistency in a weak theory of bounded arithmetic. However, Yamagata formalised ideas from programming semantics with no connection to rewriting.

9 Conclusion

In this paper, we have investigated two well-studied properties of rewriting and the λ-calculus, namely confluence and the Church-Rosser property, through the lens of proof complexity. In particular, for orthogonal TRSs, we have shown that the shortest rewrite proof obtained in a confluence argument is polynomially related to the size of the peak.

This is in contrast to earlier results on upper bounds on the size of confluence and Church-Rosser proofs that used the number of steps as size measure. While this measure is natural in the context of rewriting (or even the λ-calculus), it is less so in the context of computational complexity, from our point of view. We emphasise that our changed notion of size not only allows tractable upper bounds, but also differentiates precisely between the expressivity of (first-order) term rewrite systems and (higher-order) λ-calculus, a difference, that got somewhat blurred in related works.

We have established preliminary steps towards our motivation to study consistency proofs in weak theories of arithmetic through the lens of rewriting technologies. In future work we want to expand this direction. It seems natural to us to employ techniques from graph rewriting [12, Chapter 13] (see also [1]) to overcome the exponential lower bound on the size of the join that we have established for the Church-Rosser property. Due to the succinct encoding of multiple occurrences in graph rewriting it could be possible to allow an alternative encoding of the join and of the rewrite proof, altogether. The latter could potentially give rise to a polynomial encoding. These investigations are left to future work.

References

1 Martin Avanzini and Georg Moser. Complexity of Acyclic Term Graph Rewriting. In *Prof. 1st FSCD*, volume 52 of *LIPIcs*, pages 10:1–10:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICS.FSCD.2016.10.

2 Franz Baader and Tobias Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.

3 Hendrik Pieter Barendregt. *The lambda calculus - its syntax and semantics*, volume 103 of *Studies in logic and the foundations of mathematics*. North-Holland, 1985.

4 Henk Barendregt and Giulio Manzonetto. *A Lambda Calculus Satellite*. College Publications, 2022.

5 Arnold Beckmann. Proving Consistency of Equational Theories in Bounded Arithmetic. *J. Symb. Log.*, 67(1):279–296, 2002. doi:10.2178/JSL/1190150044.

6 Samuel R. Buss. *Bounded Arithmetic*. Bibliopolis, Naples, Italy, 1986.

7 Samuel R. Buss and Aleksandar Ignjatović. Unprovability of consistency statements in fragments of bounded arithmetic. *Ann. Pure Appl. Logic*, 74(3):221–244, 1995.

8 Alonzo Church and J. Barkley Rosser. Some properties of conversion. *Transaction of the American Mathematical Society*, 39:472–482, 1936.

9 Anupam Das. From positive and intuitionistic bounded arithmetic to monotone proof complexity. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '16, pages 126–135, New York, NY, USA, 2016. Association for Computing Machinery.

10 Ken-etsu Fujita. The Church-Rosser theorem and quantitative analysis of witnesses. *Inf. Comput.*, 263:52–56, 2018. doi:10.1016/J.IC.2018.09.002.

11 Jeroen Ketema and Jakob Grue Simonsen. Least upper bounds on the size of confluence and church-rosser diagrams in term rewriting and λ-calculus. *ACM Trans. Comput. Log.*, 14(4):31:1–31:28, 2013. doi:10.1145/2528934.

12 Terese. *Term Rewriting Systems*. Cambridge University Press, 2003.

13 Yoyuki Yamagata. Consistency proof of a fraement of pv with substitution in bounded arithmetic. *The Journal of Symbolic Logic*, 83(3):1063–1090, 2018. doi:10.1017/jsl.2018.14.