Termination of Rewriting with – and Automated Synthesis of – Forbidden Patterns

Bernhard Gramlich
TU Vienna, gramlich@logic.at

Felix Schernhammer*
TU Vienna, felixs@logic.at

1 Introduction and Overview

Rewriting with *forbidden patterns* [5] is a proper restriction of term rewriting where subterms of terms may be forbidden for reduction whenever they appear in a certain context and have a certain shape. The main goal of rewriting with restrictions is to allow reductions that are essential for computing results (i.e., normal forms) and to disallow reductions that are not needed and may lead to infinite computations.

In [5] first criteria for completeness and termination of rewriting with forbidden patterns were introduced. Here, by *completeness* we mean the power of restricted rewriting to compute useful results, which in [5] were head-normal forms. The termination criterion of [5] is based on a transformation from rewrite systems with forbidden patterns into ordinary TRSs such that termination of both coincides.

In this work we provide another direct termination proof approach based on a *contextual* extension of the well-known dependency pairs (DP) approach of [3], cf. also [2, 1]. Our new approach is applicable to a wider class of rewrite systems with forbidden patterns than the transformational approach of [5]. We also address the important question of how to synthesize suitable forbidden patterns for a given TRS in an automated way.

2 Termination of Rewriting with Forbidden Patterns

A *forbidden pattern* is a triple $\langle t,p,\lambda\rangle$, consisting of a term t, a position $p \in Pos(t)$ and a flag $\lambda \in \{h,b,a\}$. Given a term s and a forbidden pattern $\pi = \langle t,p,\lambda\rangle$, t and p determine a set of positions $P_{t,p}(s) \subseteq Pos(s)$ by $q \in P_{t,p}(s) \Leftrightarrow s|_o = t\sigma \wedge q = o.p$ for some substitution σ and some position o. Moreover, for $\pi = \langle t,p,\lambda\rangle$, $P_\pi(s) = \{o \in Pos(s) \mid \exists q \in P_{t,p}(s)\colon\, o < q\}$ if $\lambda = a$, $P_\pi(s) = \{o \in Pos(s) \mid \exists q \in P_{t,p}(s)\colon\, o > q\}$ if $\lambda = b$ and $P_\pi(s) = P_{t,p}(s)$ if $\lambda = h$. Given a set of forbidden patterns Π, the set of *forbidden* postions $P_\Pi(s)$ w.r.t. Π of a term s is $\bigcup_{\pi\in\Pi} P_\pi(s)$. The *allowed* positions of s (w.r.t. Π) are $Pos(s) \setminus P_\Pi(s)$. Rewriting with forbidden patterns (we write $\to_{\mathscr{R},\Pi}$, or just $\to_\Pi$ – or even only Π as in Π-termination – if $\mathscr{R}$ is clear from the context) is rewriting at positions that are allowed (w.r.t. Π).

Example 1. *Consider the following rewrite system, cf. e.g. [6]:*

$$\mathsf{inf}(x) \to x : \mathsf{inf}(s(x)) \qquad \mathsf{2nd}(x : (y : zs)) \to y$$

We use one forbidden pattern $\Pi = \{\langle x : (y : z), 2.2, h\rangle\}$. Then the term $s = 0 : s(0) : inf(s(s(0)))$ is a normal form w.r.t. rewriting with forbidden patterns, we also say it is a Π-normal form. Here, $x : (y : z)$ matches s and the only potential redex $inf(s(s(0)))$ cannot be reduced, as it occurs at the forbidden position 2.2 in s.

Throughout this abstract we only consider forbidden patterns with h- and b-flags and call them h- and b-*patterns*, respectively.

We base our approach to prove termination of rewriting with forbidden patterns on the well-known dependency pair (DP) framework of [4], which is in turn based on dependency pairs of [3]. The central

*The author has been supported by the Austrian Academy of Sciences under grant 22.361.

observation of the (ordinary) dependency pair approach is that given a non-terminating rewrite system $\mathcal{R}$, there exists an infinite reduction sequence (starting w.l.o.g. with a root reduction step), such that no redex contracted in this sequence contains a non-terminating proper subterm. Such reduction sequences roughly correspond to minimal dependency pair chains whose existence or non-existence is analyzed in the DP framework. For rewriting with forbidden patterns the above observation does not hold.

Example 2. *Consider the following TRS $\mathcal{R}$*

$$a \to f(a) \qquad\qquad f(x) \to g(x)$$

and an associated set of forbidden patterns $\Pi = \{\langle f(x), 1, h\rangle\}$. $\mathcal{R}$ is not Π-terminating: $a \to_\Pi f(a) \to_\Pi g(a) \to_\Pi g(f(a)) \to_\Pi g(g(a)) \to_\Pi \dots$ Note that since position 1 is forbidden in $f(a)$, we do not have $f(a) \to_\Pi f(f(a))$. Obviously, every non-Π-terminating term s must contain exactly one a. After this a is reduced, the single a-symbol in the contracted term is forbidden (as it occurs in the first argument of f). Hence, the redex of the following reduction must properly contain a.

In Example 2 reductions whose redexes properly contain non-Π-terminating terms are crucial for the existence of infinite Π-derivations. Hence, instead of ordinary non-termination we focus on a restricted form of non-Π-termination, namely non-Π-termination in a context.

Definition 1 (*termination in a context*)**.** *Let $\mathcal{R}$ be a TRS and Π be a set of forbidden patterns. A term s is* Π-terminating in context $C[\Box]_p$ *if $C[s]_p$ does not admit an infinite Π-reduction sequence where each redex contracted occurs at, below or parallel to p and where infinitely many steps are at or below p.*

We say a term s is *minimal* non-Π-terminating in a context $C[\Box]_q$ (w.r.t. a rewrite system $\mathcal{R}$ and a set of forbidden patterns Π) if s is non-Π-terminating in $C[\Box]_q$ and every proper subterm $s|_p$ of s is Π-terminating in $C[s[\Box]_p]_q$. It is easy to see that every minimal non-Π-terminating (in a context $C[\Box]_q$) term s admits a reduction sequence $C[s]_q \xrightarrow{\not\leq q *}_\Pi C[s']_q = C[l\sigma]_q \xrightarrow{q}_\Pi C[r\sigma]_q = C[t]_q$ such that t contains a minimal non-Π-terminating subterm $t|_p$ (in the context $C[t[\Box]_p]_q$). However, in contrast to ordinary rewriting and standard minimal non-terminating terms one can in general not assume that $p \in Pos_{\mathcal{F}}(r)$ (this effect similarly exists in context-sensitive rewriting, cf. [2, 1]).

Example 3. *Consider $\mathcal{R}$ and Π of Example 2 and the term $f(a)$ which is minimally non-Π-terminating (in the empty context), since position 1 is forbidden in $f(a)$ according to Π. Now consider the reduction $f(a) = f(x)\sigma \xrightarrow{\varepsilon}_\Pi g(a) = g(x)\sigma$ ($x\sigma = a$). The term $g(a)$ contains only one proper minimal non-Π-terminating subterm $g(a)|_1 = a$ despite the fact that $1 \notin Pos_{\mathcal{F}}(g(x))$.*

In our approach we pay tribute to this phenomenon by having additional dependency pairs to explicitely mimic the necessary subterm decreases in DP chains (cf. V_c, A_c and S_c in Definition 2 below). Moreover, to each dependency pair we associate a context which is the calling context of the recursive function calls the dependency pair corresponds to. Informally, this amounts to an extended *contextual* version of dependency pairs which incorporates the full information of the given rules (especially the complete right-hand sides) in the form of associated contexts, but which still enables the typical DP-based reasoning enriched by *structural* DP-rules that allow to descend into variable subterms of right-hand sides as well as to control where subsequent DP-reductions are allowed to take place.

Definition 2 (extended *contextual* dependency pairs)**.** *Let $(\mathcal{F}, R)$ be a TRS where the signature is partitioned into defined symbols $\mathcal{D}$ and constructors $\mathcal{C}$. The set of (extended)* contextual dependency pairs (CDPs) *$CDP(\mathcal{R})$ is given by $DP_c(\mathcal{R}) \uplus V_c(\mathcal{R}) \uplus A_c(\mathcal{R}) \uplus S_c(\mathcal{R})$, where*

$$\begin{aligned}
DP_c(\mathcal{R}) &= \{l^\# \to r|_p^\# \,[c] \mid l \to r \in R, p \in Pos_{\mathcal{D}}(r), c = r[\Box]_p\} \\
V_c(\mathcal{R}) &= \{l^\# \to T(r|_p) \,[c] \mid l \to r \in R, r|_p = x \in Var, c = r[\Box]_p\} \\
A_c(\mathcal{R}) &= \{T(f(x_1,\dots,x_{ar(f)})) \to f^\#(x_1,\dots,x_{ar(f)}) \,[\Box] \mid l \to r \in R, root(r|_p) = f \in \mathcal{D}\} \\
S_c(\mathcal{R}) &= \{T(f(\vec{x})) \to T(x_i)[f(\vec{x})[\Box]_i] \mid \vec{x} = x_1,\dots,x_{ar(f)}, l \to r \in R, root(r|_p) = f, i \in \{1,\dots,ar(f)\}\}.
\end{aligned}$$

Here, T is a new auxiliary function symbol (the token *symbol for "shifting attention"). We call $V_c(\mathscr{R})$* variable descent *CDPs, $S_c(\mathscr{R})$* shift *CDPs and $A_c(\mathscr{R})$* activation *CDPs.* Contextual *rules of the shape $l \to r\ [c]$ can be interpreted as $l \to c[r]$ (provided that $Var(c[r]) \subseteq Var(l)$) when used as rewrite rules. Slightly abusing notation, for a set $\mathscr{P}$ of such contextual rewrite rules (i.e. a* contextual TRS*) we denote by $\to_{\mathscr{P}}$ the corresponding induced ordinary rewrite relation.*

Example 4. *Consider the TRS $\mathscr{R}$ of Example 2. Here, $CDP(\mathscr{R})$ consist of:*

$$\begin{array}{cccc} a^\# \to a^\#[f(\Box)] & a^\# \to f^\#(a)[\Box] & f^\#(x) \to T(x)[g(\Box)] & T(a) \to a^\#[\Box] \\ T(f(x)) \to f^\#(x)[\Box] & T(g(x)) \to g^\#(x)[\Box] & T(f(x)) \to T(x)[f(\Box)] & T(g(x)) \to T(x)[g(\Box)] \end{array}$$

Definition 3 (forbidden pattern CDP problem)**.** *A forbidden pattern CDP problem (FP-CDP problem) is a quadruple $(\mathscr{P}, \mathscr{R}, \Pi, T)$ where $\mathscr{P}$ is a contextual TRS, $\mathscr{R} = (\mathscr{F}, R)$ is a TRS, Π is a set of forbidden patterns over $\mathscr{F}$ and T is a designated function symbol with $T \notin \mathscr{F}$ that occurrs only at the root position of left- and right-hand sides of rules in $\mathscr{P}$ (but not e.g. in contexts).*

Definition 4 (forbidden pattern CDP chain)**.** *Let $(\mathscr{P}, \mathscr{R}, \Pi, T)$ be a CDP problem where $\mathscr{R} = (\mathscr{F}, R)$. The sequence $S\colon s_1 \to t_1\ [c_1[\Box]_{p_1}], s_2 \to t_2\ [c_2[\Box]_{p_2}], \ldots$ is a $(\mathscr{P}, \mathscr{R}, \Pi, T)$-CDP chain if*

- *there exists a substitution $\sigma : Var \to \mathscr{T}(\mathscr{F}, V)$, such that*

$$\begin{array}{rcl} s_1\sigma & \to_{\mathscr{P}} & c_1[t_1\sigma]_{p_1} = c'_1[t_1\sigma]_{p'_1} \\ \xrightarrow{\not\leq p'_1}{}^*_{\mathscr{R}}\ c''_1[s_2\sigma]_{p'_1} & \to_{\mathscr{P}} & c''_1[c_2[t_2\sigma]_{p_2}]_{p'_1} = c'_2[t_2\sigma]_{p'_2} \\ \xrightarrow{\not\leq p'_2}{}^*_{\mathscr{R}}\ c''_2[s_3\sigma]_{p'_2} & \to_{\mathscr{P}} & c''_2[c_3[t_3\sigma]_{p_3}]_{p'_2} = c'_3[t_3\sigma]_{p'_3}\ \ldots \end{array}$$

  where $c'_i = c''_{i-1}[c_i]$ and $p'_i = p_i p'_{i-1}$ for all $1 \leq i$,

- *the $\mathscr{R}$-reduction $c'_i[t_i\sigma]_{p'_i} \xrightarrow{\not\leq p'_i}{}^*_{\mathscr{R}} c''_i[s_{i+1}\sigma]_{p'_i}$ is empty (i.e., $c'_i[t_i\sigma]_{p'_i} = c''_i[s_{i+1}\sigma]_{p'_i}$) whenever $root(t_i) = T$ (i.e., the token symbol), and*

- *for each single reduction $s \xrightarrow{q}_{\mathscr{P}} t$ or $s \xrightarrow{q}_{\mathscr{R}} t$ in this reduction sequence position q is allowed in $erase(s)$ according to Π. Here $erase(s)$ is obtained from s by replacing all marked dependency pair symbols $f^\#$ by their unmarked versions f and by replacing terms $T(s')$ by s'.*[1]

Moreover, S is minimal if for every $i \geq 0$ every subterm of $c'_i[t_i\sigma]_{p'_i}$ at position $q > p'_i$ is Π-terminating in its context (here: $s_1\sigma = c''_0[s_1\sigma]_{p'_0}$ with $p'_0 = \varepsilon$, $c''_0[\Box] = \Box$).

Example 5. *Consider the TRS $\mathscr{R}$ and Π from Example 2 ($CDP(\mathscr{R})$ is given in Example 4) and the corresponding FP-CDP $P = (CDP(\mathscr{R}), \mathscr{R}, \Pi, T)$. P admits an infinite CDP chain:*

$$a^\# \to f^\#(a)\ [\Box], f^\#(x) \to T(x)\ [g(\Box)], T(a) \to a^\#\ [\Box], \ldots$$

We say a CDP problem is *finite* if it does not admit an infinite minimal CDP chain. Note that the restrictions imposed by forbidden patterns are not used in the construction of the CDPs but only in the definition of CDP chains.

Theorem 1. *Let $\mathscr{R}$ be a TRS with an associated set of forbidden patterns Π. $\mathscr{R}$ is Π-terminating if and only if the FP-CDP problem $(CDP(\mathscr{R}), \mathscr{R}, \Pi, T)$ is finite.*

[1]Note that this definition makes sense since whenever a T occurs in s, then q is not below the occurrence of T. Moreover, this definition of *erase* is formally not compatible with the strict modularity of the DP framework. However, to restore full modularity the *erase* function could be made part of the notion of CDP problem. We refrain from doing so for notational simplicity.

Now, following the dependency pair framework of [4] we define CDP processors as functions mapping CDP problems to sets of CDP problems.

It is easy to observe that each FP-CDP chain w.r.t. a FP-CDP problem $(\mathscr{P},\mathscr{R},\Pi,T)$ is also an ordinary (though not minimal) DP chain w.r.t. $(\mathscr{P},\mathscr{R})$ (when disregarding the contexts of DPs). Hence, in some cases processors that are sound in the ordinary DP framework of [4] and do not rely on minimality can be adapted to work also in the forbidden pattern contextual extension of the DP framework.

Now we sketch the definition of a CDP processor concerning the contexts of dependency pairs. The basic idea is that in an FP-CDP chain where contexts of dependency pairs are nested, these nested contexts are reduced and thus modified only at positions parallel to the "hole" position.

In order to find patterns that match nested contexts such that this matching is oblivious to reductions parallel to the "hole" position we restrict the attention to patterns of a certain shape.

Definition 5 (Π_{orth}). *Let $\mathscr{R}$ be a TRS and Π be a set of corresponding forbidden patterns. The subset $\Pi_{orth} \subseteq \Pi$ (of forbidden patterns that are "orthogonal to $\mathscr{R}$") consists of those forbidden patterns $\langle t,p,\lambda\rangle$ where $\lambda \in \{h,b\}$, t is linear and not overlapped below the root by any rule of $\mathscr{R}$.*

Lemma 1. *Let $(\mathscr{P},\mathscr{R},\Pi,T)$ be an FP-CDP problem and let $s_1 \to t_1[c_1[\Box]_{p_1}],\ldots,s_n \to t_n[c_n[\Box]_{p_n}]$ be a FP-CDP chain. If position $q = p_1.\cdots.p_n$ is forbidden in the term $c_1[c_2[\ldots c_n[erase(t_n)]_{p_n}\ldots]_{p_2}]_{p_1}$ by a forbidden pattern from Π_{orth}, then q is forbidden in $c'_1[c'_2[\ldots c'_n[erase(t_n)]_{p_n}\ldots]_{p_2}]_{p_1}$ where $c_i \to^*_{\mathscr{R}} c'_i$ with reductions parallel to p_i for all $1 \le i \le n$.*

Based on Lemma 1 we define an FP-CDP processor, the *context-processor*, that erases a CDP $s_1 \to t_1[c_1]$ of a FP-CDP problem $(\mathscr{P},\mathscr{R},\Pi,T)$ if, in all potential sequences of CDPs $s_1 \to t_1[c_1[\Box]_{p_1}],\ldots,s_n \to t_n[c_n[\Box]_{p_n}]$ of length n, position $p_1.\cdots.p_n$ is Π_{orth}-forbidden in the term $c_1[c_2[\ldots c_n[\Box]_{p_n}\ldots]_{p_2}]_{p_1}$.

We evaluated our termination approach based on contextual dependency pairs by using it for proving outermost termination of the TRSs in the outermost category of the TPDB. This is feasible, since outermost rewriting is a special case of rewriting with forbidden patterns (cf. [5] for details).

As processors we used only a CDP graph processor, a polynomial ordering processor and the context processor considering CDP chains of length 3. Out of 133 (potentially terminating) examples our implementation was able to verify outermost termination of 33 which is worse than the best transformational approaches in the termination competition 2009.[2] However, our method is very fast using less than 1 second on average in successful proofs. Moreover, various conceivable and promising refinements of our direct approach are not yet implemented.

3 Synthesizing Forbidden Patterns

Now we are going to utilize the machinery of Section 2 in order to synthesize suitable forbidden patterns for a given rewrite system $\mathscr{R}$. Note that this is possible, since the forbidden patterns are not used to construct the CDPs, but only in the CDP processors. Hence, it makes sense to perform the synthesis of (appropriate) forbidden patterns on the fly in the processors as well. To this end we reconsider the context processor. Instead of checking whether a term constructed from nested contexts is matched by a forbidden pattern from Π_{orth}, now a pattern matching these nested contexts is constructed. In the simplest case this pattern could be the context itself.

Example 6. *Consider the TRS of Example 2. This TRS yields (among others) the CDP*

$$a^\# \to a^\# \,[f(\Box)].$$

[2]The best result was 72 successful proofs (by *Jambox*).

A context processor considering CDP sequences of length 2 *encounters e.g. the term* $f(f(erase(a^\#)) = f(f(a))$*. Thus, a forbidden pattern* $\pi = \langle f(f(a)), 1.1, h\rangle$ *could be used. Indeed, when this forbidden pattern is used, there is no infinite FP-CDP chain w.r.t. to this single CDP.*

Usually one wants to restrict the shape of the generated patterns for instance by demanding that all forbidden patterns contain allowed redexes and do not overlap (each other), in order to ensure that normal forms and Π-normal forms coincide (then weak termination of $\rightarrow_{\mathscr{R}}$ coincides with termination of $\rightarrow_{\Pi}$).

A second choice for restrictions on the shape of forbidden patterns might be *canonical* forbidden patterns as defined in [5]. If only canonical forbidden patterns are used, forbidden pattern normal forms are head-normal forms of unrestricted rewriting.

Synthesis of forbidden patterns adhering to these syntactical restrictions can be achieved by systematically considering more general versions of nested contexts obtained e.g. by replacing some subterm by a fresh variable and checking whether these more general patterns yield termination.

Example 7. *Consider the TRS* $\mathscr{R}$ *of Example 1. There is a contextual dependency pair*

$$inf^\#(x) \rightarrow inf^\#(s(x))[x : \Box]$$

Applying a context processor that considers CDP sequences of length 2 *we get a term* $x : x' : inf(s(x'))$*, which we cannot use as forbidden pattern since it is not linear and thus not in* Π_{orth} *(hence a termination proof with the context processor would not be possible). Instead we linearize the term obtaining* $x : y : inf(s(z))$ *which we can use as canonical forbidden pattern. Indeed,* $\mathscr{R}$ *is* Π*-terminating when choosing* $\Pi = \langle x : y : inf(s(z)), 2.2, h\rangle$.

4 Conclusion

We introduced a novel approach to prove termination of rewriting with forbidden *h*- and *b*-patterns based on an extended contextual version of the dependency pair framework. Moreover we have sketched the definition of a context processor that analyzes the contexts and the possible nestings of contexts obtained by sequences of contextual dependency pairs. One advantage of this approach is that suitable forbidden patterns for a TRS can be generated on-the-fly during the CDP analysis.

References

[1] B. Alarcón, F. Emmes, C. Fuhs, J. Giesl, R. Gutiérrez, S. Lucas, P. Schneider-Kamp, and R. Thiemann. Improving context-sensitive dependency pairs. In I. Cervesato, H. Veith and A. Voronkov, eds., *Proc. LPAR'08*, LNCS 5530, pp. 636–651, Springer, 2008.

[2] B. Alarcón, R. Gutiérrez, and S. Lucas. Context-sensitive dependency pairs. In S. Arun-Kumar and N. Garg, eds., *Proc. FST&TCS'06*, LNCS 4337, pp. 297–308, Kolkata, India, Springer, 2006.

[3] T. Arts and J. Giesl. Termination of term rewriting using dependency pairs. *Theor. Comput. Sci.*, 236(1-2):133–178, 2000.

[4] J. Giesl, R. Thiemann, P. Schneider-Kamp, and S. Falke. Mechanizing and improving dependency pairs. *J. Autom. Reason.*, 37(3):155–203, 2006.

[5] B. Gramlich and F. Schernhammer. Extending context-sensitivity in term rewriting. In M. Fernandez, ed., *Final Proc. WRS'09, Brasília, Brazil, June 28, 2009*, *EPTCS*, Vol. 15, pp. 56-68, January 2010.

[6] S. Lucas. Termination of on-demand rewriting and termination of OBJ programs. In H. Sondergaard, ed., *Proc. PPDP'01*, pp. 82–93, Firenze, Italy, ACM Press, New York, September 2001.