

6 Termination Revisited

So far: Termination as a subordinate task for entailment checking.

TRS is generated by some saturation process; ordering must be chosen before the saturation starts.

Now: Termination as a main task (e. g., for program analysis).

TRS is fixed and known in advance.

Literature:

Nao Hirokawa and Aart Middeldorp: Dependency Pairs Revisited, RTA 2004, pp. 249-268 (in particular Sect. 1–4).

Thomas Arts and Jürgen Giesl: Termination of Term Rewriting Using Dependency Pairs, Theoretical Computer Science, 236:133-178, 2000.

6.1 Dependency Pairs

Invented by T. Arts and J. Giesl in 1996, many refinements since then.

Given: finite TRS R over $\Sigma = (\Omega, \emptyset)$.

$T_0 := \{ t \in T_\Sigma(X) \mid \text{there is an infinite derivation } t \rightarrow_R t_1 \rightarrow_R t_2 \rightarrow_R \dots \}$.

$T_\infty := \{ t \in T_0 \mid \forall p > \varepsilon : t/p \notin T_0 \}$ = minimal elements of T_0 w. r. t. \triangleright .

$t \in T_0 \Rightarrow$ there exists a $t' \in T_\infty$ such that $t \triangleright t'$.

R is non-terminating iff $T_0 \neq \emptyset$ iff $T_\infty \neq \emptyset$.

Assume that $T_\infty \neq \emptyset$ and consider some non-terminating derivation starting from $t \in T_\infty$. Since all subterms of t allow only finite derivations, at some point a rule $l \rightarrow r \in R$ must be applied at the root of t (possibly preceded by rewrite steps below the root):

$$t = f(t_1, \dots, t_n) \xrightarrow{>\varepsilon}_R^* f(s_1, \dots, s_n) = l\sigma \xrightarrow{\varepsilon}_R r\sigma.$$

In particular, $root(t) = root(l)$, so we see that the root symbol of any term in T_∞ must be contained in $D := \{ root(l) \mid l \rightarrow r \in R \}$. D is called the set of *defined symbols* of R ; $C := \Omega \setminus D$ is called the set of *constructor symbols* of R .

The term $r\sigma$ is contained in T_0 , so there exists a $v \in T_\infty$ such that $r\sigma \triangleright v$.

If v occurred in $r\sigma$ at or below a variable position of r , then $x\sigma/p = v$ for some $x \in var(r) \subseteq var(l)$, hence $s_i \triangleright x\sigma$ and there would be an infinite derivation starting from some t_i . This contradicts $t \in T_\infty$, though.

Therefore, $v = u\sigma$ for some non-variable subterm u of r . As $v \in T_\infty$, we see that $\text{root}(u) = \text{root}(v) \in D$. Moreover, u cannot be a proper subterm of l , since otherwise again there would be an infinite derivation starting from some t_i .

Putting everything together, we obtain

$$t = f(t_1, \dots, t_n) \xrightarrow{>\varepsilon}_R^* f(s_1, \dots, s_n) = l\sigma \xrightarrow{\varepsilon}_R r\sigma \supseteq u\sigma$$

where $r \supseteq u$, $\text{root}(u) \in D$, $l \not\supseteq u$, u is not a variable.

Since $u\sigma \in T_\infty$, we can continue this process and obtain an infinite sequence.

If we define $S := \{l \rightarrow u \mid l \rightarrow r \in R, r \supseteq u, \text{root}(u) \in D, l \not\supseteq u, u \notin X\}$, we can combine the rewrite step at the root and the subterm step and obtain

$$t \xrightarrow{>\varepsilon}_R^* l\sigma \xrightarrow{\varepsilon}_S u\sigma.$$

To get rid of the superscripts ε and $>\varepsilon$, it turns out to be useful to introduce a new set of function symbols f^\sharp that are only used for the root symbols of this derivation:

$$\Omega^\sharp := \{f^\sharp/n \mid f/n \in \Omega\}.$$

For a term $t = f(t_1, \dots, t_n)$ we define $t^\sharp := f^\sharp(t_1, \dots, t_n)$; for a set of terms T we define $T^\sharp := \{t^\sharp \mid t \in T\}$.

The set of *dependency pairs* of a TRS R is then defined by

$$DP(R) := \{l^\sharp \rightarrow u^\sharp \mid l \rightarrow r \in R, r \supseteq u, \text{root}(u) \in D, l \not\supseteq u, u \notin X\}.$$

For $t \in T_\infty$, the sequence using the S -rule corresponds now to

$$t^\sharp \xrightarrow{*}_R l^\sharp\sigma \rightarrow_{DP(R)} u^\sharp\sigma$$

where $t^\sharp \in T_\infty^\sharp$ and $u^\sharp\sigma \in T_\infty^\sharp$.

(Note that rules in R do not contain symbols from Ω^\sharp , whereas all roots of terms in $DP(R)$ come from Ω^\sharp , so rules from R can only be applied below the root and rules from $DP(R)$ can only be applied at the root.)

Since $u^\sharp\sigma$ is again in T_∞^\sharp , we can continue the process in the same way. We obtain: R is non-terminating iff there is an infinite sequence

$$t_1 \xrightarrow{*}_R t_2 \rightarrow_{DP(R)} t_3 \xrightarrow{*}_R t_4 \rightarrow_{DP(R)} \dots$$

with $t_i \in T_\infty^\sharp$ for all i .

Moreover, if there exists such an infinite sequence, then there exists an infinite sequence in which all DPs that are used are used infinitely often. (If some DP is used only finitely often, we can cut off the initial part of the sequence up to the last occurrence of that DP; the remainder is still an infinite sequence.)

Such infinite sequences correspond to “cycles” in the “dependency graph”:

Dependency graph $DG(R)$ of a TRS R :

directed graph

nodes: dependency pairs $s \rightarrow t \in DP(R)$

edges: from $s \rightarrow t$ to $u \rightarrow v$ if there are σ, τ such that $t\sigma \rightarrow_R^* u\tau$.

Intuitively, we draw an edge between two dependency pairs, if these two dependency pairs can be used after another in an infinite sequence (with some R -steps in between). While this relation is undecidable in general, there are reasonable overapproximations:

The functions cap and ren are defined by:

$$cap(x) = x$$

$$cap(f(t_1, \dots, t_n)) = \begin{cases} y & \text{if } f \in D \\ f(cap(t_1), \dots, cap(t_n)) & \text{if } f \in C \cup D^\# \end{cases}$$

$$ren(x) = y, \quad y \text{ fresh}$$

$$ren(f(t_1, \dots, t_n)) = f(ren(t_1), \dots, ren(t_n))$$

The overapproximated dependency graph contains an edge from $s \rightarrow t$ to $u \rightarrow v$ if $ren(cap(t))$ and u are unifiable.

A cycle in the dependency graph is a non-empty subset $K \subseteq DP(R)$ such that there is a non-empty path from every DP in K to every DP in K (the two DPs may be identical).

Let $K \subseteq DP(R)$. An infinite rewrite sequence in $R \cup K$ of the form

$$t_1 \rightarrow_R^* t_2 \rightarrow_K t_3 \rightarrow_R^* t_4 \rightarrow_K \dots$$

with $t_i \in T_\infty^\#$ is called K -minimal, if all rules in K are used infinitely often.

R is non-terminating iff there is a cycle $K \subseteq DP(R)$ and a K -minimal infinite rewrite sequence.

6.2 Subterm Criterion

Our task is to show that there are no K -minimal infinite rewrite sequences.

Suppose that every dependency pair symbol $f^\#$ in K has positive arity (i. e., no constants). A *simple projection* π is a mapping $\pi : \Omega^\# \rightarrow \mathbb{N}$ such that $\pi(f^\#) = i \in \{1, \dots, \text{arity}(f^\#)\}$.

We define $\pi(f^\#(t_1, \dots, t_n)) = t_{\pi(f^\#)}$.

Theorem 6.1 (Hirokawa and Middeldorp) *Let K be a cycle in $DG(R)$. If there is a simple projection π for K such that $\pi(l) \supseteq \pi(r)$ for every $l \rightarrow r \in K$ and $\pi(l) \triangleright \pi(r)$ for some $l \rightarrow r \in K$, then there are no K -minimal sequences.*

Proof. Suppose that

$$t_1 \rightarrow_R^* u_1 \rightarrow_K t_2 \rightarrow_R^* u_2 \rightarrow_K \dots$$

is a K -minimal infinite rewrite sequence. Apply π to every t_i :

Case 1: $u_i \rightarrow_K t_{i+1}$. There is an $l \rightarrow r \in K$ such that $u_i = l\sigma$, $t_{i+1} = r\sigma$. Then $\pi(u_i) = \pi(l)\sigma$ and $\pi(t_{i+1}) = \pi(r)\sigma$. By assumption, $\pi(l) \supseteq \pi(r)$. If $\pi(l) = \pi(r)$, then $\pi(u_i) = \pi(t_{i+1})$. If $\pi(l) \triangleright \pi(r)$, then $\pi(u_i) = \pi(l)\sigma \triangleright \pi(r)\sigma = \pi(t_{i+1})$. In particular, $\pi(u_i) \triangleright \pi(t_{i+1})$ for infinitely many i (since every DP is used infinitely often).

Case 2: $t_i \rightarrow_R^* u_i$. Then $\pi(t_i) \rightarrow \pi(u_i)$.

By applying π to every term in the K -minimal infinite rewrite sequence, we obtain an infinite $(\rightarrow_R \cup \triangleright)$ -sequence containing infinitely many \triangleright -steps. Since \triangleright is well-founded, there must also exist infinitely many \rightarrow_R -steps (otherwise the infinite sequence would have an infinite tail consisting only of \triangleright -steps, contradicting well-foundedness.)

Now note that $\triangleright \circ \rightarrow_R \subseteq \rightarrow_R \circ \triangleright$. Therefore we can commute \triangleright -steps and \rightarrow_R -steps and move all \rightarrow_R -steps to the front. We obtain an infinite \rightarrow_R -sequence that starts with $\pi(t_1)$. However $t_1 \triangleright \pi(t_1)$ and $t_1 \in T_\infty$, so there cannot be an infinite \rightarrow_R -sequence starting from $\pi(t_1)$. \square

Problem: The number of cycles in $DG(R)$ can be exponential.

Better method: Analyze strongly connected components (SCCs).

SCC of a graph: maximal subgraph in which there is a non-empty path from every node to every node. (The two nodes can be identical.)³

Important property: Every cycle is contained in some SCC.

Idea: Search for a simple projection π such that $\pi(l) \supseteq \pi(r)$ for all DPs $l \rightarrow r$ in the SCC. Delete all DPs in the SCC for which $\pi(l) \triangleright \pi(r)$ (by the previous theorem, there cannot be any K -minimal infinite rewrite sequences using these DPs). Then re-compute SCCs for the remaining graph and re-start.

No SCCs left \Rightarrow no cycles left $\Rightarrow R$ is terminating.

Example: See Ex. 13 from Hirokawa and Middeldorp.

³There are several definitions of SCCs that differ in the treatment of edges from a node to itself.