3.4 Critical Pairs

Showing local confluence (Sketch):

Problem: If $t_1 \leftarrow_E t_0 \rightarrow_E t_2$, does there exist a term s such that $t_1 \rightarrow_E^* s \leftarrow_E^* t_2$?

If the two rewrite steps happen in different subtrees (disjoint redexes): yes.

If the two rewrite steps happen below each other (overlap at or below a variable position): yes.

If the left-hand sides of the two rules overlap at a non-variable position: needs further investigation.

Critical Pairs

Showing local confluence (Sketch):

Question:

Are there rewrite rules $l_1 \longrightarrow r_1$ and $l_2 \longrightarrow r_2$ such that some subterm l_1/p and l_2 have a common instance $(l_1/p)\sigma_1 = l_2\sigma_2$?

Observation:

If we assume w.o.l.o.g. that the two rewrite rules do not have common variables, then only a single substitution is necessary: $(l_1/p)\sigma = l_2\sigma$.

Further observation:

The mgu of l_1/p and l_2 subsumes all unifiers σ of l_1/p and l_2.

Critical Pairs

Let $l_i \longrightarrow r_i$ $(i = 1, 2)$ be two rewrite rules in a TRS R whose variables have been renamed such that $\mathrm{var}(\{l_1, r_1\}) \cap \mathrm{var}(\{l_2, r_2\}) = \emptyset$.

Let $p \in \mathrm{pos}(l_1)$ be a position such that l_1/p is not a variable and σ is an mgu of l_1/p and l_2.

Then $r_1\sigma \leftarrow l_1\sigma \longrightarrow (l_1\sigma)[r_2\sigma]_p$.

$\langle r_1\sigma, (l_1\sigma)[r_2\sigma]_p\rangle$ is called a critical pair of R.

The critical pair is joinable (or: converges), if $r_1\sigma \downarrow_R (l_1\sigma)[r_2\sigma]_p$.

Critical Pairs

Theorem 3.18 ("Critical Pair Theorem"):
A TRS R is locally confluent if and only if all its critical pairs are joinable.

Proof:
"only if": obvious, since joinability of a critical pair is a special case of local confluence.

Critical Pairs

Proof:

"if": Suppose s rewrites to t_1 and t_2 using rewrite rules $l_i \rightarrow r_i \in R$ at positions $p_i \in \text{pos}(s)$, where $i = 1, 2$.
Without loss of generality, we can assume that the two rules are variable disjoint, hence $s/p_i = l_i\theta$ and $t_i = s[r_i\theta]_{p_i}$.

We distinguish between two cases: Either p_1 and p_2 are in disjoint subtrees ($p_1 \; || \; p_2$), or one is a prefix of the other (w.o.l.o.g., $p_1 \leq p_2$).

Critical Pairs

Case 1: $p_1 \mid\mid p_2$.

Then $s = s[l_1\theta]_{p_1}[l_2\theta]_{p_2}$,
and therefore $t_1 = s[r_1\theta]_{p_1}[l_2\theta]_{p_2}$ and $t_2 = s[l_1\theta]_{p_1}[r_2\theta]_{p_2}$.

Let $t_0 = s[r_1\theta]_{p_1}[r_2\theta]_{p_2}$.
Then clearly $t_1 \longrightarrow_R t_0$ using $l_2 \rightarrow r_2$ and $t_2 \longrightarrow_R t_0$ using $l_1 \rightarrow r_1$.

Critical Pairs

Case 2: $p_1 \leq p_2$.

Case 2.1: $p_2 = p_1\,q_1\,q_2$, where l_1/q_1 is some variable x.

In other words, the second rewrite step takes place at or below a variable in the first rule. Suppose that x occurs m times in l_1 and n times in r_1 (where $m \geq 1$ and $n \geq 0$).

Then $t_1 \rightarrow_R^* t_0$ by applying $l_2 \rightarrow r_2$ at all positions $p_1\,q'\,q_2$, where q' is a position of x in r_1.

Conversely, $t_2 \rightarrow_R^* t_0$ by applying $l_2 \rightarrow r_2$ at all positions $p_1\,q\,q_2$, where q is a position of x in l_1 different from q_1, and by applying $l_1 \rightarrow r_1$ at p_1 with the substitution θ', where $\theta' = \theta[x \mapsto (x\theta)[r_2\theta]_{q_2}]$.

Critical Pairs

Case 2.2: $p_2 = p_1\,p$, where p is a non-variable position of l_1.

Then $s/p_2 = l_2\theta$ and $s/p_2 = (s/p_1)/p = (l_1\theta)/p = (l_1/p)\theta$, so θ is a unifier of l_2 and l_1/p.

Let σ be the mgu of l_2 and l_1/p, then $\theta = \tau \circ \sigma$ and $\langle r_1\sigma, (l_1\sigma)[r_2\sigma]_p\rangle$ is a critical pair.

By assumption, it is joinable, so $r_1\sigma \rightarrow_R^* v \leftarrow_R^* (l_1\sigma)[r_2\sigma]_p$.

Consequently, $t_1 = s[r_1\theta]_{p_1} = s[r_1\sigma\tau]_{p_1} \rightarrow_R^* s[v\tau]_{p_1}$ and $t_2 = s[r_2\theta]_{p_2} = s[(l_1\theta)[r_2\theta]_p]_{p_1} = s[(l_1\sigma\tau)[r_2\sigma\tau]_p]_{p_1} = s[((l_1\sigma)[r_2\sigma]_p)\tau]_{p_1} \rightarrow_R^* s[v\tau]_{p_1}$.

This completes the proof of the Critical Pair Theorem.

Critical Pairs

Note: Critical pairs between a rule and (a renamed variant of) itself must be considered – except if the overlap is at the root (i.e., $p = \varepsilon$).

Critical Pairs

Corollary 3.19:

A terminating TRS R is confluent if and only if all its critical pairs are joinable.

Proof:

By Newman's Lemma and the Critical Pair Theorem.

Critical Pairs

Corollary 3.20:

For a finite terminating TRS, confluence is decidable.

Proof:

For every pair of rules and every non-variable position in the first rule there is at most one critical pair $\langle u_1, u_2 \rangle$.

Reduce every u_i to some normal form u_i'. If $u_1' = u_2'$ for every critical pair, then R is confluent, otherwise there is some non-confluent situation $u_1' \leftarrow_R^* u_1 \leftarrow_R s \rightarrow_R u_2 \rightarrow_R^* u_2'$.

3.5 Termination

Termination problems:

Given a finite TRS R and a term t, are all R-reductions starting from t terminating?

Given a finite TRS R, are all R-reductions terminating?

Proposition 3.21:
Both termination problems for TRSs are undecidable in general.

Proof:
Encode Turing machines using rewrite rules and reduce the (uniform) halting problems for TMs to the termination problems for TRSs.

Termination

Consequence:

Decidable criteria for termination are not complete.

Reduction Orderings

Goal:

Given a finite TRS R, show termination of R by looking at finitely many rules $l \rightarrow r \in R$, rather than at infinitely many possible replacement steps $s \rightarrow_R s'$.

Reduction Orderings

A binary relation $\sqsupset$ over $\mathsf{T}_\Sigma(X)$ is called
compatible with Σ-operations,
if $s \sqsupset s'$ implies $f(t_1, \ldots, s, \ldots, t_n) \sqsupset f(t_1, \ldots, s', \ldots, t_n)$
for all $f/n \in \Omega$ and $s, s', t_i \in \mathsf{T}_\Sigma(X)$.

Lemma 3.22:
The relation $\sqsupset$ is compatible with Σ-operations, if and only if
$s \sqsupset s'$ implies $t[s]_p \sqsupset t[s']_p$
for all $s, s', t \in \mathsf{T}_\Sigma(X)$ and $p \in \mathrm{pos}(t)$.

(compatible with Σ-operations $=$ compatible with contexts)

Reduction Orderings

A binary relation $\sqsupset$ over $\mathsf{T}_\Sigma(X)$ is called stable under substitutions, if $s \sqsupset s'$ implies $s\sigma \sqsupset s'\sigma$ for all $s, s' \in \mathsf{T}_\Sigma(X)$ and substitutions σ.

Reduction Orderings

A binary relation $\sqsupset$ is called a rewrite relation, if it is compatible with Σ-operations and stable under substitutions.

Example: If R is a TRS, then $\rightarrow_R$ is a rewrite relation.

A strict partial ordering over $\mathrm{T}_\Sigma(X)$ that is a rewrite relation is called rewrite ordering.

A well-founded rewrite ordering is called reduction ordering.

Reduction Orderings

Theorem 3.23:
A TRS R terminates if and only if there exists a reduction ordering $\succ$ such that $l \succ r$ for every rule $l \rightarrow r \in R$.

Proof:
"if": $s \rightarrow_R s'$ if and only if $s = t[l\sigma]_p$, $s' = t[r\sigma]_p$.
If $l \succ r$, then $l\sigma \succ r\sigma$ and therefore $t[l\sigma]_p \succ t[r\sigma]_p$.
This implies $\rightarrow_R \subseteq \succ$.
Since $\succ$ is a well-founded ordering, $\rightarrow_R$ is terminating.

"only if": Define $\succ = \rightarrow_R^+$.
If $\rightarrow_R$ is terminating, then $\succ$ is a reduction ordering.

The Interpretation Method

Proving termination by interpretation:

Let $\mathcal{A}$ be a Σ-algebra;
let $\succ$ be a well-founded strict partial ordering on its universe.

Define the ordering $\succ_{\mathcal{A}}$ over $\mathrm{T}_{\Sigma}(X)$ by $s \succ_{\mathcal{A}} t$ iff
$\mathcal{A}(\beta)(s) \succ \mathcal{A}(\beta)(t)$ for all assignments $\beta : X \to U_{\mathcal{A}}$.

Is $\succ_{\mathcal{A}}$ a reduction ordering?

The Interpretation Method

Lemma 3.24:

$\succ_{\mathcal{A}}$ is stable under substitutions.

Proof:

Let $s \succ_{\mathcal{A}} s'$, that is,

$\mathcal{A}(\beta)(s) \succ \mathcal{A}(\beta)(s')$ for all assignments $\beta : X \to U_{\mathcal{A}}$.

Let σ be a substitution. We have to show that

$\mathcal{A}(\gamma)(s\sigma) \succ \mathcal{A}(\gamma)(s'\sigma)$ for all assignments $\gamma : X \to U_{\mathcal{A}}$.

Choose $\beta = \gamma \circ \sigma$, then by the substitution lemma,

$\mathcal{A}(\gamma)(s\sigma) = \mathcal{A}(\beta)(s) \succ \mathcal{A}(\beta)(s') = \mathcal{A}(\gamma)(s'\sigma)$.

Therefore $s\sigma \succ_{\mathcal{A}} s'\sigma$.

The Interpretation Method

A function $F : U_{\mathcal{A}}^n \longrightarrow U_{\mathcal{A}}$ is called monotone (w.r.t. $\succ$), if $a \succ a'$ implies

$F(b_1, \ldots, a, \ldots, b_n) \succ F(b_1, \ldots, a', \ldots, b_n)$

for all $a, a', b_i \in U_{\mathcal{A}}$.

The Interpretation Method

Lemma 3.25:

If the interpretation $f_{\mathcal{A}}$ of every function symbol f is monotone w.r.t. $\succ$, then $\succ_{\mathcal{A}}$ is compatible with Σ-operations.

Proof:

Let $s \succ s'$, that is, $\mathcal{A}(\beta)(s) \succ \mathcal{A}(\beta)(s')$ for all $\beta : X \to U_{\mathcal{A}}$.

Let $\beta : X \to U_{\mathcal{A}}$ be an arbitrary assignment.

Then $\mathcal{A}(\beta)(f(t_1, \ldots, s, \ldots, t_n))$

$= f_{\mathcal{A}}(\mathcal{A}(\beta)(t_1), \ldots, \mathcal{A}(\beta)(s), \ldots, \mathcal{A}(\beta)(t_n))$

$\succ f_{\mathcal{A}}(\mathcal{A}(\beta)(t_1), \ldots, \mathcal{A}(\beta)(s'), \ldots, \mathcal{A}(\beta)(t_n))$

$= \mathcal{A}(\beta)(f(t_1, \ldots, s', \ldots, t_n))$.

Therefore $f(t_1, \ldots, s, \ldots, t_n) \succ_{\mathcal{A}} f(t_1, \ldots, s', \ldots, t_n)$.

The Interpretation Method

Theorem 3.26:
If the interpretation $f_{\mathcal{A}}$ of every function symbol f is monotone w.r.t. $\succ$, then $\succ_{\mathcal{A}}$ is a reduction ordering.

Proof:
By the previous two lemmas, $\succ_{\mathcal{A}}$ is a rewrite relation.
If there were an infinite chain $s_1 \succ_{\mathcal{A}} s_2 \succ_{\mathcal{A}} \ldots$, then it would correspond to an infinite chain $\mathcal{A}(\beta)(s_1) \succ \mathcal{A}(\beta)(s_2) \succ \ldots$ (with β chosen arbitrarily).
Thus $\succ_{\mathcal{A}}$ is well-founded.
Irreflexivity and transitivity are proved similarly.

Polynomial Orderings

Polynomial orderings:

Instance of the interpretation method:

The carrier set $U_{\mathcal{A}}$ is some subset of the natural numbers.

To every n-ary function symbol f associate a polynomial $P_f(X_1, \ldots, X_n) \in \mathbb{N}[X_1, \ldots, X_n]$ with coefficients in $\mathbb{N}$ and indeterminates $X_1, \ldots, X_n$. Then define $f_{\mathcal{A}}(a_1, \ldots, a_n) = P_f(a_1, \ldots, a_n)$ for $a_i \in U_{\mathcal{A}}$.

Polynomial Orderings

Requirement 1:

If $a_1, \ldots, a_n \in U_{\mathcal{A}}$, then $f_A(a_1, \ldots, a_n) \in U_{\mathcal{A}}$.
(Otherwise, $\mathcal{A}$ would not be a Σ-algebra.)

Polynomial Orderings

Requirement 2:

f_A must be monotone (w.r.t. $\succ$).

From now on:

$U_{\mathcal{A}} = \{ n \in \mathbb{N} \mid n \geq 2 \}$.

If $f/0 \in \Omega$, then P_f is a constant ≥ 2.

If $f/n \in \Omega$ with $n \geq 1$, then P_f is a polynomial $P(X_1, \ldots, X_n)$, such that every X_i occurs in some monomial with exponent at least 1 and non-zero coefficient.

$\Rightarrow$ Requirements 1 and 2 are satisfied.

Polynomial Orderings

The mapping from function symbols to polynomials can be extended to terms:
A term t containing the variables $x_1, \dots, x_n$
yields a polynomial P_t with indeterminates $X_1, \dots, X_n$
(where X_i corresponds to $\beta(x_i)$).

Example:

$$\Omega = \{a/0,\ f/1,\ g/3\},$$
$$U_{\mathcal{A}} = \{\, n \in \mathbb{N} \mid n \geq 2 \,\},$$
$$P_a = 3, \quad P_f(X_1) = X_1^2, \quad P_g(X_1, X_2, X_3) = X_1 + X_2 X_3.$$

Let $t = g(f(a), f(x), y)$, then $P_t(X, Y) = 9 + X^2 Y$.

Polynomial Orderings

If P, Q are polynomials in $\mathbb{N}[X_1, \ldots, X_n]$, we write $P > Q$ if $P(a_1, \ldots, a_n) > Q(a_1, \ldots, a_n)$ for all $a_1, \ldots, a_n \in U_{\mathcal{A}}$.

Clearly, $l \succ_{\mathcal{A}} r$ iff $P_l > P_r$.

Question: Can we check $P_l > P_r$ automatically?

Polynomial Orderings

Hilbert's 10th Problem:

> Given a polynomial $P \in \mathbb{Z}[X_1, \ldots, X_n]$ with integer coefficients, is $P = 0$ for some n-tuple of natural numbers?

Theorem 3.27:
Hilbert's 10th Problem is undecidable.

Proposition 3.28:
Given a polynomial interpretation and two terms l, r, it is undecidable whether $P_l > P_r$.

Proof:
By reduction of Hilbert's 10th Problem.

Polynomial Orderings

One possible solution:

Test whether $P_l(a_1, \ldots, a_n) > P_r(a_1, \ldots, a_n)$
for all $a_1, \ldots, a_n \in \{\, x \in \mathbb{R} \mid x \geq 2 \,\}$.

This is decidable (but very slow).
Since $U_{\mathcal{A}} \subseteq \{\, x \in \mathbb{R} \mid x \geq 2 \,\}$, it implies $P_l > P_r$.

Polynomial Orderings

Another solution (Ben Cherifa and Lescanne):

Consider the difference $P_l(X_1, \ldots, X_n) - P_r(X_1, \ldots, X_n)$ as a polynomial with real coefficients and apply the following inference system to it to show that it is positive for all $a_1, \ldots, a_n \in U_{\mathcal{A}}$:

Polynomial Orderings

$P \; \Rightarrow_{BCL} \; \top,$

if P contains at least one monomial with a positive coefficient and no monomial with a negative coefficient.

$P + c\,X_1^{p_1} \cdots X_n^{p_n} - d\,X_1^{q_1} \cdots X_n^{q_n} \; \Rightarrow_{BCL} \; P + c'\,X_1^{p_1} \dots X_n^{p_n},$

if $c, d > 0$, $p_i \geq q_i$ for all i,
and $c' = c - d \cdot 2^{(q_1 - p_1) + \cdots + (q_n - p_n)} \geq 0$.

$P + c\,X_1^{p_1} \cdots X_n^{p_n} - d\,X_1^{q_1} \cdots X_n^{q_n} \; \Rightarrow_{BCL} \; P - d'\,X_1^{q_1} \dots X_n^{q_n},$

if $c, d > 0$, $p_i \geq q_i$ for all i,
and $d' = d - c \cdot 2^{(p_1 - q_1) + \cdots + (p_n - q_n)} > 0$.

Polynomial Orderings

Lemma 3.29:

If $P \Rightarrow_{BCL} P'$, then $P(a_1, \ldots, a_n) \geq P'(a_1, \ldots, a_n)$ for all $a_1, \ldots, a_n \in U_{\mathcal{A}}$.

Proof:

Follows from the fact that $a_i \in U_{\mathcal{A}}$ implies $a_i \geq 2$.

Proposition 3.30:

If $P \Rightarrow^{+}_{BCL} \top$, then $P(a_1, \ldots, a_n) > 0$ for all $a_1, \ldots, a_n \in U_{\mathcal{A}}$.