Theoretical Computer Science 913 (2022) 1–7

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Note

Loop-checking and the uniform word problem for join-semilattices with an inflationary endomorphism

Marc Bezem [a,*], Thierry Coquand [b]

[a] *Department of Informatics, University of Bergen, Norway*
[b] *Department of Computer Science and Engineering, Chalmers/University of Gothenburg, Sweden*

ARTICLE INFO

Article history:
Received 4 December 2020
Received in revised form 12 January 2022
Accepted 13 January 2022
Available online 19 January 2022
Communicated by A. Avron

Keywords:
Join-semilattices with one inflationary endomorphism
Uniform word problem
Loop-checking
Decidability

ABSTRACT

We solve in polynomial time two decision problems that occur in type checking when typings depend on universe level constraints.

1. Introduction

The *uniform word problem* for an equational theory T is to determine, given a finite set E of relations between generators, whether a given relation is provable from E in T. The theory of join-semilattices with a finite set of endomorphisms was shown to have an EXPTIME-complete uniform word problem by Baader et al. [1] and Hofmann [5].

Here we show that the special case of one inflationary endomorphism, denoted by $\_^+$, has a uniform word problem that can be solved in PTIME.

A *loop* is a term t such that $t^+ \leqslant t$. We show that *loop-checking*, i.e., testing whether or not a loop exists, is also decidable in PTIME.

In our special case, both the uniform word problem and loop-checking are relevant for dependent type theory with universes. Our decision procedure can be seen as forward reasoning with loop-detection on the fly.

We start from the equational definition of semilattices in which the join, denoted by $\vee$, is an associative, commutative, and idempotent binary operation. The endomorphism satisfies the following two equational axioms:

$$x \vee x^+ = x^+ \qquad (x \vee y)^+ = x^+ \vee y^+$$

The logic is ordinary equational logic. We denote the resulting theory by $\mathcal{L}$. For example, we can prove $(t^+)^+ \vee t = (t^+)^+$ in $\mathcal{L}$, for any term t. Also, we can infer $s^+ = (t^+)^+$ from $s = t^+$, but not conversely. As is customary, we let $s \geqslant t$ abbreviate $s \vee t = s$. Throughout this note we call a join-semilattice with an inflationary endomorphism simply a *semilattice*. We call t^+ the *successor* of t.

* Corresponding author.
E-mail addresses: Marc.Bezem@uib.no (M. Bezem), coquand@chalmers.se (T. Coquand).

https://doi.org/10.1016/j.tcs.2022.01.017
0304-3975/

2. Semilattice presentations and Horn clauses

A *semilattice presentation* consists of a set V of *generators* and a set C of *relations*. We will colloquially call the generators also variables, and the relations constraints. For any semilattice term t and natural number k, let $t+k$ denote the k-fold successor of t. Thus $t+0=t$, and we may use $t+1$ and t^+ interchangeably. A *term over* V is a term of the form $x_1+k_1 \vee \cdots \vee x_m+k_m$, with all $x_i \in V$ and $k_i \in \mathrm{N}$.

Since the endomorphism commutes with the join operation, every semilattice term t is equal to a term of the form $x_1+k_1 \vee \cdots \vee x_m+k_m$, with all variables x_i occurring in t and all $k_i \in \mathrm{N}$.

A *relation* is an equation $s=t$ with s,t terms over V. A constraint like $x=y^+$ with $x,y \in V$ expresses a relation between the generators x and y and should not be read as an implicitly universally quantified axiom in which the x and y can be instantiated.

The semilattice *presented by* (V,C) consists of terms over V modulo provable equality from C. The latter will be denoted by $C \vdash_{\mathcal{L}} s=t$. In the next sections we will prove that $C \vdash_{\mathcal{L}} s=t$ is decidable in polynomial time for finite semilattice presentations (V,C). The results in this section hold for arbitrary semilattice presentations.

We follow Lorenzen [7, Section 2] for an equivalent characterisation of $C \vdash_{\mathcal{L}} s=t$ using Horn clauses. Let (V,C) be a semilattice presentation. Let $s := x_1+k_1 \vee \cdots \vee x_m+k_m$ and $t := y_1+l_1 \vee \cdots \vee y_n+l_n$ be terms over V. From the constraint $s=t$ we can prove $m+n$ inequalities which we write as Horn clauses by replacing join by conjunction (written as ",") and $\geqslant$ by implication. In this note all clauses are propositional Horn clauses $A \to b$ with a non-empty body A and conclusion b. The atoms are of the form $x+k$ with $x \in V$ and $k \in \mathrm{N}$, and we will often call such Horn clauses simply *clauses*. We will express by $A \subseteq B$ that all atoms in A also occur in B.

Thus we get the set $S_{s=t}$ consisting of the following Horn clauses:

$$\begin{array}{c}x_1+k_1,\ldots,x_m+k_m \to y_1+l_1\\ \cdots\\ x_1+k_1,\ldots,x_m+k_m \to y_n+l_n\end{array}$$

and

$$\begin{array}{c}y_1+l_1,\ldots,y_n+l_n \to x_1+k_1\\ \cdots\\ y_1+l_1,\ldots,y_n+l_n \to x_m+k_m\end{array}$$

Let S_C be the union of all $S_{s=t}$ with $s=t$ a constraint in C.

We reflect for a moment on which other clauses we need. Consider the axiom $x \vee x^+ = x^+$. This would lead to three clauses: $x, x^+ \to x^+$, $x^+ \to x^+$, and $x^+ \to x$. Only the last is non-trivial, we call it a *predecessor clause*. The next question is: for which x do we need a predecessor clause? Since the axiom $x \vee x^+ = x^+$ is implicitly universally quantified, we would need all instances with x a term over V. For this it suffices to have all predecessor clauses $x+k+1 \to x+k$ with $x \in V$ and all $k \in \mathrm{N}$.

The axiom stating that endomorphism and join commute is built-in in the notion of term over V and does not require extra clauses. However, we should not forget the congruence axioms from equational logic. Congruence of the endomorphism means that $s=t$ implies $s+1=t+1$. This requires that we close the set of clauses under *shifting upwards*: if $A \to b$ is in the clause set, then so is $A+1 \to b+1$, where $A+1$ is the set of atoms of the form $a+1$ with $a \in A$. Congruence of join means that $s=t$ implies $r \vee s = r \vee t$. It is easy to see that this does not require extra clauses.

In summary: given a semilattice presentation (V,C), let $\overline{S_C}$ be the smallest set of clauses that is closed under shifting upwards and contains the set S_C coming from the constraints in C, as well as all predecessor clauses for each $v \in V$.

Given a set S of Horn clauses, let $S \vdash_{\mathcal{H}} A \to b$ denote provability from S. One convenient way to define this is by two inference rules:

$$\frac{}{S \vdash_{\mathcal{H}} A \to b}\ b \in A \qquad \frac{S \vdash_{\mathcal{H}} A, c \to b}{S \vdash_{\mathcal{H}} A \to b}\ \text{there exists } A' \to c \text{ in } S \text{ with } A' \subseteq A$$

This is the inductive way of defining forward reasoning, that is, using the Horn clauses in S to generate atoms from A. We can also use this definition if A is infinite. Then it is more customary to write $S, A \vdash_{\mathcal{H}} b$.

For X a set of atoms, define $X^+ := \{x^+ \mid x \in X\}$ and $X+k := \{x+k \mid x \in X\}$ for all $k \in \mathrm{N}$. For S a set of clauses, define $S^+ := \{X^+ \to y^+ \mid X \to y \text{ in } S\}$. The following lemma will be used later on.

Lemma 2.1. *Let V be a finite set of variables, and $A \to b$ a Horn clause. Let S and T be sets of Horn clauses, where all clauses in T have conclusion* in V. *Then the following three are equivalent: (1) $S \vdash_{\mathcal{H}} A \to b$; (2) $S^+ \vdash_{\mathcal{H}} A^+ \to b^+$; (3) $S^+ \cup T \vdash_{\mathcal{H}} V, A^+ \to b^+$.*

Proof. Immediate by induction on the definition of $\vdash_{\mathcal{H}}$. $\square$

We have the following theorem generalizing [7, Theorem 3].

Theorem 2.2. *For all terms $x_1+k_1,\dots,x_m+k_m,\ y+l$ over V we have:*

$$C \vdash_{\mathcal{L}} x_1+k_1 \vee \cdots \vee x_m+k_m \geqslant y+l \quad \textit{iff} \quad \overline{S_C} \vdash_{\mathcal{H}} x_1+k_1,\dots,x_m+k_m \to y+l.$$

Proof. The if-part is a straightforward structural induction on the definition of $\vdash_{\mathcal{H}}$: all steps can easily be mimicked in semilattice theory. The converse implication is more interesting. For any set of Horn clauses X, Y, let $X \vdash_{\mathcal{H}} Y$ mean that $X \vdash_{\mathcal{H}} A \to b$ for all clauses $A \to b$ in Y. For any two terms s, t over V, define $s \equiv t$ by $\overline{S_C} \vdash_{\mathcal{H}} S_{s=t}$. We have $s^+ \equiv t^+$ if $s \equiv t$, as $\overline{S_C}$ is closed under shifting upwards. We also have $s \vee r \equiv t \vee r$ and $r \vee s \equiv r \vee t$ if $s \equiv t$.

Now we can define $s \vee t$ on terms over V in the obvious way, and $(x_1+k_1 \vee \cdots \vee x_m+k_m)^+ := x_1+k_1+1 \vee \cdots \vee x_m+k_m+1$. Both are well-defined operations modulo the congruence $\equiv$.

Then one can show that all axioms and rules are satisfied modulo $\equiv$. For example, all semilattice axioms are satisfied, for example, the predecessor clauses prove $s \vee s^+ \equiv s^+$. Moreover, for each constraint $s = t$ in C we have $s \equiv t$, as $S_{s=t}$ is included in $\overline{S_C}$.

By soundness, if $C \vdash_{\mathcal{L}} s = t$, then $s \equiv t$. In particular we have the only-if part of the theorem. □

3. Decidability

In this section we first prove the decidability of $\overline{S_C} \vdash_{\mathcal{H}} A \to b$. In the next section we show that our decision procedure is in PTIME. By Theorem 2.2 this is sufficient for the decidability of $C \vdash_{\mathcal{L}} s = t$. We recall a basic fact about Horn clauses: the models (as satisfying sets of atoms) are closed under intersection. Moreover, every set X of atoms can be extended to a unique minimal model; this minimal model consists of all atoms that can be inferred from X using the Horn clauses as generating rules, as defined just before Lemma 2.1.

We proceed by defining an auxiliary notion that we call 'gain'. The *gain* of a clause $x_1+k_1,\dots,x_m+k_m \to y+l$ is l minus the minimum of $\{k_1,\dots,k_m\}$. For example, predecessor clauses have gain -1. The gain of a clause is invariant under shifting.

Let N^∞ be N extended with ∞, totally ordered by $n < \infty$ for all $n \in \mathsf{N}$. Given a finite semilattice presentation (V, C), we view a function $f : V \to \mathsf{N}^\infty$ as specifying the downward closed set of atoms $\{v+k \mid v \in V,\ k \in \mathsf{N},\ k \leqslant f(v)\}$. Note that this set contains all atoms $v+k$ if $f(v) = \infty$. We are interested in such sets as models of $\overline{S_C}$. A clause $A \to b$ with $A = x_1+k_1,\dots,x_m+k_m$ and $b = y+l$, all $x_i, y \in V$, is satisfied by f if $l \leqslant f(y)$ when all $k_i \leqslant f(x_i)$. Predecessor clauses are of course satisfied by downward closure.

Lemma 3.1. *Given $f : V \to \mathsf{N}^\infty$ and a clause $A \to b$, let P be the problem whether or not $A+k \to b+k$ is satisfied by f for all $k \in \mathsf{N}$. Then P is decidable.*

Proof. Assume $A = x_1+k_1,\dots,x_m+k_m$ and $b = y+l$ with $x_i, y \in V$. Let W consist of all variables x_i in A that satisfy $f(x_i) < \infty$. If W is empty, then P is equivalent to $f(y) = \infty$. Otherwise, let $k_0 = \min_{\{i \mid x_i \in W\}}(f(x_i) - k_i)$. If $k_0 < 0$, then P holds. If $k_0 \geqslant 0$, then P is equivalent to $l + k_0 \leqslant f(y)$. □

Given a finite semilattice presentation (V, C) and a subset W of V, we denote by $\overline{S_C}|W$ the set of clauses in $\overline{S_C}$ mentioning only variables in W, and by $\overline{S_C}{\downarrow}W$ the set of clauses in $\overline{S_C}$ with conclusion over W.

Theorem 3.2. *Let (V, C) be a finite semilattice presentation. For any $f : V \to \mathsf{N}^\infty$ we can compute the least $g \geqslant f$ that is a model of $\overline{S_C}$.*

We prepare the proof of this theorem with a lemma.

Lemma 3.3. *Let (V, C) be a finite semilattice presentation. Let W be a strict subset of V such that for any $f : W \to \mathsf{N}^\infty$ we can compute the least $g \geqslant f$ that is a model of $\overline{S_C}|W$. Then for any $f : V \to \mathsf{N}^\infty$ with $f(V - W) \subseteq \mathsf{N}$ we can compute the least $h \geqslant f$ that is a model of $\overline{S_C}{\downarrow}W$.*

Proof. Let conditions be as stated in the lemma. Since (V, C) is finite we can compute the smallest number $\mathsf{Maxgain} \geqslant 0$ such that each clause in S_C has gain at most $\mathsf{Maxgain}$. Let $f : V \to \mathsf{N}^\infty$ with $f(V - W) \subseteq \mathsf{N}$ be given and denote its restriction to W by f_W. By the definition of $\overline{S_C}{\downarrow}W$, any minimal $h \geqslant f$ that is a model of $\overline{S_C}{\downarrow}W$ coincides with f on $V - W$, so we focus on its values on W. By assumption we can compute the least $g_f \geqslant f_W$ that is a model of $\overline{S_C}|W$. Now we look at clauses in $\overline{S_C}{\downarrow}W - \overline{S_C}|W$. Such clauses are of the form $X, Y \to w+k$ with X a non-empty set of atoms over $V - W$, and possibly empty Y over W. If $X = \dots, x_i+k_i, \dots$ is satisfied by f, then by the definition of $\mathsf{Maxgain}$, using $f(V - W) \subseteq \mathsf{N}$, we have:

$$k \leq \min_i(f(x_i)) + \mathsf{Maxgain} \leqslant \max(f(V-W)) + \mathsf{Maxgain} \in \mathsf{N}. \tag{1}$$

Inequality (1) gives an upper bound on values that clauses in $\overline{S_C}{\downarrow}W - \overline{S_C}|W$ can generate. Define:

$$M(g_f) := \sum_{w \in W} \max(0, \max(f(V-W)) + \mathsf{Maxgain} - g_f(w)).$$

After these preparations we are ready to prove the lemma by induction on $M(g_f)$. More precisely, we prove for all $n \in \mathsf{N}$ and $f : V \to \mathsf{N}^\infty$ with $f(V-W) \subseteq \mathsf{N}$, if $M(g_f) = n$, then we can compute the least $h \geqslant f$ that is a model of $\overline{S_C}{\downarrow}W$.

In the base case $M(g_f) = 0$ we have $k \leqslant g_f(w)$ for all $w \in W$ and all clauses in $\overline{S_C}{\downarrow}W$ are satisfied by g_f. In this case we take $h = g_f$ on W and $h = f$ on $V - W$ and we are done.

For the induction step, let $M(g_f) > 0$ and assume the result has been proved for smaller values of $M(\_)$. We now make a case distinction that is decidable by Lemma 3.1 since S_C is finite (even though $\overline{S_C}$ is not). Thus we only have to check finitely many clauses. If all clauses in $\overline{S_C}{\downarrow}W - \overline{S_C}|W$ are satisfied by g_f we are again done, like in the base case. Otherwise, one such clause gives value $g_f(w) + k + 1$ for some $w \in W$ and $k \in \mathsf{N}$, using values of f on $V - W$ and values of g_f on W. Then we know by (1) that the term with $g_f(w)$ in the sum defining $M(g_f)$ is positive. Define $f' : V \to \mathsf{N}^\infty$ by

$$\begin{aligned} f'(x) &= f(x) && \text{for } x \text{ in } V - W, \\ f'(y) &= g_f(y) && \text{for } y \text{ in } W - \{w\}, \text{ and} \\ f'(w) &= g_f(w) + k + 1. \end{aligned}$$

We have $g_{f'}(w) \geqslant f'(w) > g_f(w)$, so $M(g_{f'}) < M(g_f)$ and we can apply the induction hypothesis to f'. The resulting h for f' also works for f, since every step in the sequence $h \geqslant f' > g_f \geqslant f$ is by adding atoms that can be inferred from f. □

We now return to the proof of Theorem 3.2.

Proof. Let (V, C) be a finite semilattice presentation and $\mathsf{Maxgain} \geqslant 0$ the smallest number such that each clause in S_C has gain at most $\mathsf{Maxgain}$. By induction on $|V|$ we compute, for any $f : V \to \mathsf{N}^\infty$, the least $g \geqslant f$ that is a model of $\overline{S_C}$.

In the base case $|V| = 0$ there is nothing to prove.

For the induction step, let $|V| > 0$ and assume the theorem has been proved for smaller values of $|V|$. Let $f : V \to \mathsf{N}^\infty$. If $f(v) = \infty$ for some $v \in V$, then we can eliminate v from $\overline{S_C}$. Recall that $f(v) = \infty$ means that all atoms of the form $v + k$ are true. This means that all clauses of the form $\ldots \to v + k$ can be left out from $\overline{S_C}$. Also, in each clause $A \to b$ in $\overline{S_C}$ we can leave out all atoms of the form $v + k$ from A. If we get a (forbidden) clause $\emptyset \to v' + l$, we can infer $f(v') = \infty$ and continue. We end up with a strict subset V' of V and a (restricted) $f : V' \to \mathsf{N}$.

Alternatively, we can do the simplification of the semilattice presentation and end up with (V', C'). Here C' is obtained from C by removing all $v + k$ from the joins, taking care to continue if a join becomes empty, and so on.[1] Both methods lead to the same set of clauses $\overline{S_{C'}}$, and this set satisfies all requirements, in particular each clause has gain at most $\mathsf{Maxgain}$.

In such a case we can directly apply the induction hypothesis to the simplified semilattice presentation, and extend the function with values ∞ for all $v \in V - V'$. Otherwise we have $f : V \to \mathsf{N}$. Since S_C is finite, using Lemma 3.1, we can decide whether f is a model of $\overline{S_C}$. If so, we are done. If not, we proceed as follows. For every $x \in V$ and $A \to x + l$ in S_C, consider the integer k_0 as in the proof of Lemma 3.1. If $k_0 < 0$, then A and all its upward shifts are false in f. If $k_0 \geqslant 0$, then we can infer the atom $x + l + k_0$. If moreover $l + k_0 > f(x)$, then this atom is new. Let W be the (non-empty) subset of variables $x \in V$ for which there is a clause in S_C that yields a new atom, and let $g(x)$ be the maximum of $f(x)$ and the possible values $l + k_0$ for x obtained in this way. Since S_C is finite, $g \geqslant f$ is a function from V to N.

We distinguish the cases $W = V$ and $W \subset V$. If $W = V$ we are done since then $h(x) = \infty$ for all $x \in V$ is the least $h \geqslant f$ that is a model of $\overline{S_C}$. Proof: if $W = V$, then $g(x) > f(x)$ and we can infer $x + g(x)$ from f, for all $x \in V$. By using the predecessor clauses we hence also infer $x + f(x) + 1$ from f for every $x \in V$. Since $\overline{S_C}$ is closed under shifting upwards, we can infer $x + f(x) + k$, and hence $x + k$, for every $k \in \mathsf{N}$ and $x \in V$.

The last case is that W is a non-empty strict subset of V, and we can apply the induction hypothesis to W to satisfy the condition of Lemma 3.3. We apply the conclusion of the lemma to g, noting that $g(V - W) = f(V - W)$ is a subset of N. Hence we can compute the least $h \geqslant g$ that is a model of $\overline{S_C}{\downarrow}W$. This function h coincides with g and f on $V - W$.

If $h(w) = \infty$ for some $w \in W$, then we simplify and apply the induction and are done as described in the first paragraphs of the step case. Otherwise we have $h : V \to \mathsf{N}$. We now make a case distinction that is decidable by Lemma 3.1, since S_C is finite. If all clauses in $\overline{S_C}$ are satisfied by h then we are done. Otherwise, we can infer in one step a value $h(y) + k + 1$ for some variable(s) y. Such y must be in $V - W$ since h is a model of $\overline{S_C}{\downarrow}W$. For every $y \in V$, let $j(y)$ be the maximum of $h(y)$ and the values $h(y) + k + 1$ that can possibly be inferred in one step. We extend W to W' with all y such that $j(y) > h(y)$ and proceed with W' and j (to keep all work done) in the same way as with W and g above. This terminates since we exhaust V. □

[1] One can take $f(v) = \infty$ to mean $\bot = v = v + 1$, which yields $\bot = v + k$ for all $k \in \mathsf{N}$.

From Theorem 3.2 we get the decidability of $\overline{S_C} \vdash_{\mathcal{H}} A \to b$ when A can be represented by a function $f : V \to \mathsf{N}$. If every variable v in V occurs in A, then $f(v)$ is simply the maximal k such that $v+k \in A$. We can then simply check $\overline{S_C} \vdash_{\mathcal{H}} A \to b$ by computing the least $g \geqslant f$ that is a model of $\overline{S_C}$ and check whether atom b is satisfied by g.

The decision method above for $\overline{S_C} \vdash_{\mathcal{H}} A \to b$ only works if every variable in V occurs in A. However, it is not difficult to extend Theorem 3.2 so that we get decidability of $\overline{S_C} \vdash_{\mathcal{H}} A \to b$ in general. Let $T = \{v^+ \to v \mid v \in V\}$. By Lemma 2.1 we see that $\overline{S_C} \vdash_{\mathcal{H}} A \to b$ is equivalent to $\overline{S_C}^+ \cup T \vdash_{\mathcal{H}} V, A^+ \to b^+$, and $\overline{S_C}^+ \cup T$ is in fact $\overline{S_{C^+}}$, where C^+ is the set of constraints $s^+ = t^+$ with $s = t$ in C. Thus we get:

Corollary 3.4. *For all A, b, s, t, $\overline{S_C} \vdash_{\mathcal{H}} A \to b$ and $C \vdash_{\mathcal{L}} s = t$ are decidable.*

Another application of Theorem 3.2 is loop checking. Given a finite semilattice presentation (V, C), a *loop* is a term t over V such that $C \vdash_{\mathcal{L}} t^+ = t$. Let L be the semilattice presented by (V, C). Let N be the semilattice with carrier N and with the usual max and successor function.

Corollary 3.5. *Exactly one of the following two decidable cases holds: (1) There is a loop; (2) There is a homomorphism $h : L \to N$.*

Proof. Let m be maximal such that $x+m$ occurs in the body of a clause in S_C. Take $f : V \to \mathsf{N}$ to be the constant m function and compute g according to Theorem 3.2. Let W be the subset of V such that $g(w) = \infty$ for all $w \in W$. Claim: if W is not empty, then we have a loop, case (1), because there exists an $n \in \mathsf{N}$ such that $\overline{S_C} \vdash_{\mathcal{H}} W + n \to w + n + 1$ for all $w \in W$.

Proof of claim. if $W = V$, then $n = m$ and we are done, otherwise take $n = \max(g(V - W)) + \mathsf{Maxgain}$. The idea of this choice of n is that variables in $V - W$ cannot play a role above n. In order to see this, define $f' : V \to \mathsf{N}$ by $f'(x) = g(x)$ if $x \in V - W$ and $f'(w) = n$ if $w \in W$. Then $g \geqslant f' \geqslant f$, so g is also the minimal model when starting from f'. Since f' and g coincide on $V - W$ we have that all clauses in $\overline{S_C}{\downarrow}(V - W)$ are satisfied by f'. By the particular choice of n, using same reasoning as in the proof of Lemma 3.3, albeit with f' instead of f, also $\overline{S_C}{\downarrow}W - \overline{S_C}|W$ is satisfied by f'. Hence the only clauses that play a role in computing g are clauses from $\overline{S_C}|W$, so we must have $\overline{S_C} \vdash_{\mathcal{H}} W + n \to w + n + 1$ for all $w \in W$. It follows that $\vee_{w \in W}\, w + n$ is a loop.

If W is empty we can construct a homomorphism $h : L \to N$, case (2). Define $h(x) = \max(g(V)) - g(x)$ for all $x \in V$. Extend h to terms over V by $h(t^+) = h(t) + 1$ and $h(s \vee t) = \max(h(s), h(t))$. We have to make sure that definition of h respects equality in L, that is, if $C \vdash_{\mathcal{L}} s = t$, then $h(s) = h(t)$. For this it suffices to show $h(s) = h(t)$ for all $s = t$ in C. This can in turn be simplified to: $h(x_1 + k_1 \vee \cdots \vee x_m + k_m) \geqslant h(y) + l$ for every $x_1 + k_1, \ldots, x_m + k_m \to y + l$ in S_C. Easy calculations show that we must prove $\min(g(x_1) - k_1, \ldots, g(x_m) - k_m) \leqslant g(y) - l$. Wlog we assume that $g(x_1) - k_1$ is the minimum on the left. Since $x_1 + k_1, \ldots, x_m + k_m \to y + l$ in S_C we know that $g(x_1) \geqslant f(x_1) \geqslant k_1$. Shifting the clause upwards by $g(x_1) - k_1$ we get the clause $x_1 + g(x_1), \ldots, x_m + k_m + g(x_1) - k_1 \to y + l + g(x_1) - k_1$ in S_C^+. Due to the assumption that $g(x_1) - k_1$ is minimal, the body of this clause is satisfied by g. Since g is a model of S_C^+ by Theorem 3.2, the conclusion is also satisfied by g, that is, $g(y) \geqslant l + g(x_1) - k_1$, so $g(y) - l \geqslant g(x_1) - k_1$. This completes the proof that h respects equality in L.

It should be clear that (1) and (2) exclude each other. □

4. Complexity analysis

All proofs in this note are constructive, so that they in fact contain algorithms. In this section we shall show that these algorithms are polynomial. The small-model property in Corollary 4.2 below, a refinement of the small-model property in [2], will be instrumental.

Let's define the input size. The size $|E|$ of logical expression E is the number of logical symbols in E. The size $|f|$ of a function $f : V \to \mathsf{N}^\infty$ is taken to be the *maximum of all its values* $< \infty$. This choice for the size of f implies that the complexity of some algorithms depending on f is weakly polynomial. However, for the important Corollary 3.4 and 3.5 the algorithms are strongly polynomial, that is, in PTIME.

Our algorithms are essentially performing just forward reasoning. However, since we have an infinite language, one has to take care to terminate, which is explained in the inductive proofs. Moreover, termination should happen in a polynomial number of reasoning steps, each taking at most polynomial time. We prepare by the following lemma.

Lemma 4.1. *Let (V, C) be a finite semilattice presentation and $\mathsf{Maxgain} \geqslant 0$ the smallest number such that each clause in S_C has gain at most $\mathsf{Maxgain}$. Let $f : V \to \mathsf{N}^\infty$ be a model of $\overline{S_C}$ such that V can be partitioned as $V = L \cup H \cup I$ with $I = \{v \in V \mid f(v) = \infty\}$ and $f(x) - f(y) > \mathsf{Maxgain}$ for all $x \in H$ and $y \in L$. Then $g : V \to \mathsf{N}^\infty$ defined by $g(x) = f(x) - 1$ for all $x \in H$ and $g(y) = f(y)$ for all $y \in L \cup I$ is also a model of $\overline{S_C}$.*

Proof. We only have to check clauses with conclusion over H. Let $y_1 + k_1, \ldots, y_m + k_m \to x + l$ be a clause in $\overline{S_C}$ with $y_i \in V$ and $x \in H$. If the premiss is satisfied in g, and some $y_i \in L$, then

$$g(x)+1 = f(x) > f(y_i) + \mathsf{Maxgain} = g(y_i) + \mathsf{Maxgain} \geqslant k_i + \mathsf{Maxgain} \geqslant l.$$

Hence $g(x) \geqslant l$, so also the conclusion holds in g. If no $y_i \in L$, then we use that any clause $A \to b$ that only mentions variables in $H \cup I$ is satisfied in g when $A^+ \to b^+$ is satisfied in f. $\square$

We immediately get the following small-model property.

Corollary 4.2. *For any f, the least $g \geqslant f$ that is a model of $\overline{S_C}$ satisfies $|g| \leqslant |f| + |V| * \mathsf{Maxgain}$.*

We now analyse the complexity of various results point for point.

- The complexity of the test in Lemma 3.1 is clearly polynomial in $|A \to b|$ and $|f|$.
- In Theorem 3.2 our choice for the size of f becomes clear: with only one clause $x, y \to y+1$ in S_C and $f(y) = 0$, forward reasoning takes $f(x)+1$ steps to arrive at the model $f(y) = f(x)+1$. This is polynomial in the value of $f(x)$, but not in its binary representation.
  The proof of Theorem 3.2 is intertwined with the proof of Lemma 3.2. In view of Corollary 4.2 it suffices that there is a polynomial bound on the work done for each forward inference and that there is steady progress in the global state encoded by functions $f : V \to \mathsf{N}^\infty$, until the algorithm terminates. Both are easily verified by inspection of the proofs, using that the test in Lemma 3.1 is polynomial.
- In the statement of Corollary 3.4 and 3.4 there is no function $f : V \to \mathsf{N}^\infty$. However, the proofs apply Theorem 3.2 with such a function, satisfying $|f| \leqslant |A|$ and $|f| \leqslant |S_C|$, respectively. Hence both corollaries yield algorithms that are polynomial in the input size.

One may wonder what is the role of the assumption that the endomorphism is inflationary, leading to the predecessor clauses in $\overline{S_C}$. A first answer is that models of the predecessor clauses are downward closed, leading to an efficient representation of models by functions $f : V \to \mathsf{N}^\infty$. Moreover we have the following example.

Let p_i be the i-th prime number and consider clauses $x_i \to x_i + p_i$, and $x_1+1, \ldots, x_n+1 \to y+1$, and $y+1 \to y$ as the only predecessor clause. Include all the upward shifts of these clauses. Then we can infer $x_1, \ldots, x_n \to y$, but forward reasoning takes exponentially many steps.

5. Discussion

5.1. Motivation

The motivation for this problem comes from dependent type theory, where the relevant operations on universe levels are to take the supremum of two levels, and to increment a level.

In order to avoid universe inconsistencies in type theory, it has been suggested in [6], [4], [8] to use constraints on universe levels. In [6], [4] these constraints are linear inequalities between universe levels. In [8] also the maximum of two universe levels is used. A typing would then only be valid if its constraints can be inferred from the set of constraints in the context. Moreover, the latter set should be consistent in the sense that there are no loops. As defined above, a loop is a semilattice term that is equal to its successor; a good intuition is that loops lead to universe inconsistencies comparable to the paradoxes in set theory.

In type theory it is important that typing checking is decidable. The results of this note show that having typings depend on a set of universe level constraints preserves the decidability of type checking.

Since dependent type theory is meant to be a foundation of mathematics, we want to make minimal mathematical assumptions about the universe levels. For example, we don't say that they are natural numbers with a zero, successor and a maximum function (like Voevodsky in [8], referring to Presburger Arithmetic for decidability). Such assumptions would weaken, at least philosophically, the foundational claim of dependent type theory: natural numbers are introduced as an inductive type at some later point in the development. For similar reasons we don't assume that universe levels are totally ordered, nor that the endomorphism is injective.

5.2. Example

Let S_C consist of the clauses $a, b \to b+1$; $b \to c+3$; $c+1 \to d$; $b, d+2 \to e$. We shall show how the proof of Theorem 3.2 works to find the minimal model above the function that is constant 0. Sets of variables will be denoted by a string, e.g., $V = abcde$. We denote functions with domain V by a string of values, e.g., 00000. (Digits will suffice in this example.) We have $\mathsf{Maxgain} = 3$.

First we compute the function $g_0 = 01300$ with the maximal values that can be obtained *in one step* from 00000. We have $W_0 = bc$. (We give indices to W, g in the third paragraph of the induction step in the proof of Theorem 3.2, since we need to iterate the induction step.) The proof of Theorem 3.2 now invokes Lemma 3.3 to compute the minimal model above

g_0 of all clauses in S with conclusion over W_0, $a, b \to b+1$ and $b \to c+3$. Also the proof of Lemma 3.3 is inductive, but we immediately get that this minimal model is $h_0 = 01400$.

We now check whether h_0 is a model of S_C. It is not: (only) the clause $c+1 \to d$ is not satisfied, and the maximal value for d is 3. So we continue with $g_1 = 01430$ and $W_1 = bcd$ and compute the minimal model h_1 of the clauses with conclusion over W_1, which happens to be g_1 itself.

One more, very similar round yields $g_2 = 01431$ and $W_2 = bcde$ and $h_2 = g_2$, which satisfies all clauses of S_C, and is the minimal model starting from the function that is constant 0.

An interesting variation would be to add the clause $e \to a$ and to see that the algorithm then detects the loop. This is indeed the case because $e \to a$ is not satisfied by h_2 and so $g_3 = 11431$ and $W_3 = abcde$ are computed. Now $V = W_3$ and the loop has been detected, the minimal model is the function that is constant ∞.

5.3. Related work

Voevodsky remarked in [8] that universe level expressions are exactly 'linear' functions of universe level variables in the tropical (max-plus) semiring, even though he referred to Presburger Arithmetic for decidability. For tropical semirings, see the book [3].

Some problems in tropical semirings are formulated in the same language as in this note, but interpreted in the integer, rational, or real numbers. The latter are totally ordered, which makes these problems different. Consider, for example, the constraint $x \vee y = x^+$. When the ordering is total, this constraint implies $y = x^+$. However, if x and y are incomparable, there are models of $x \vee y = x^+$ in which $y = x^+$ is false. The simplest such countermodel has three elements. The model of $x \vee y = x^+$ as described in the proof of Theorem 2.2 is based on terms over $V = \{x, y\}$, modulo the congruence $\equiv$. This congruence gives $x+k \vee y+l \equiv x+l+1$ if $k \leqslant l$ (use $x, y \to x+1$), and $x+k \vee y+l \equiv x+k$ if $k > l$ (use $x+1 \to y$). In this model $y = x^+$ is false ($y \to x+1$ cannot be derived).

Another connection is with work on uniform word problems with endomorphisms. Both [1,5] show that this problem is decidable but EXPTIME-complete. We describe a PTIME algorithm in a special case: for only *one* endomorphism which is moreover *inflationary*. It seems to be an open problem whether there is a PTIME algorithm for one endomorphism without any extra assumption. A similar question can be asked in the case of finitely many endomorphisms that are all inflationary.

It is easy to describe Hofmann's [5] algorithm with our notation. This is an EXPTIME algorithm for the case without the predecessor clauses $a+1 \to a$. The idea is to provide a complete cut-free derivation system. Given a set of clauses R, the derivations rules are the following.

1. $A \to a$ if a is in A
2. $A, (B+1) \to b+1$ if $B \to b$
3. $A \to a$ if there is a rule $c_1, \ldots, c_n \to d$ in R such that $A \to c_1, \ldots, A \to c_n$ and $A, d \to a$.

With the subformulas of an atom $a+k$ being all $a+i$ with $i \leqslant k$, one sees that the subformula property holds, since only rules $c_1, \ldots, c_n \to d$ in R are allowed in 3. A general cut rule is admissible. Deciding whether $A \to a$ follows from R can then be done by a top-down search of (cut-free) proofs in this system.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We are grateful to the anonymous referees for constructive criticism and pointers to related work. We would also like to thank Franz Baader and Viorica Sofronie-Stokkermans for updating us on related decidability results for description logics.

References

[1] Franz Baader, Sebastian Brandt, Carsten Lutz, Pushing the EL envelope, in: IJCAI 2005, 2005, pp. 364–369.
[2] Marc Bezem, Robert Nieuwenhuis, Enric Rodríguez, The max-atom problem and its relevance, in: I. Cervesato, H. Veith, A. Voronkov (Eds.), Proceedings LPAR-15, in: LNAI, vol. 5330, Springer-Verlag, Berlin, 2008, pp. 47–62.
[3] Peter Butkovič, Max-Linear Systems: Theory and Algorithms, Springer-Verlag, 2010.
[4] Robert Harper, Robert Pollack, Type checking with universes, Theor. Comput. Sci. 89 (1991) 107–136.
[5] Martin Hofmann, Proof-theoretic approach to description logic, in: LICS 2005, 2005, pp. 229–237.
[6] Gérard Huet, Extending the calculus of constructions with Type:Type, Unpublished manuscript, April 1987.
[7] Paul Lorenzen, Algebraische und logistische Untersuchungen über freie Verbände, J. Symb. Log. 16 (2) (1951) 81–106, English translation by Stefan Neuwirth: https://arxiv.org/abs/1710.08138.
[8] Vladimir Voevodsky, A universe polymorphic type system, manuscript, http://www.math.ias.edu/Voevodsky/voevodsky-publications_abstracts.html#UPTS, 2014.