Efficient type reconstruction in the presence of inheritance

Marcin Benke

December, 1994

Abstract

The complexity of type reconstruction for simply-typed lambda calculus with subtype relation resulting from single inheritance (i.e. being a disjoint union of tree-like posets) is analyzed. As a result a class of posets including (but not restricted to) trees is defined, for which the said problem is solvable in polynomial time. Moreover, for this class of posets the problem of solving flat systems of inequalities is proved to be complete for NLOGSPACE.

1 Introduction

In 1984, John Mitchell [Mit84], stated the problem of type inference for the simply-typed lambda calculus with coercions between base types (we shall call this problem TRS). He proved that it is reducible to the problem of satisfiability of systems of inequalities between type variables and base type constants (called PO-SAT or FLAT-SSI), and therefore decidable. He did not though consider in depth the question of complexity.

In 1989 Wand and O'Keefe [WO89] claimed that PO-SAT is NP-complete and that TRS reduces to it in polynomial time. Regretfully, the latter proof is incorrect. In [LM92] Lincoln and Mitchell observe that algorithm for trees works in exponential rather than polynomial time.

TRS reduces to the problem of solving system of inequalities involving not only constants and variables but also an arrow — a binary functor (this problem, called PO-SAT will be specified in detail later). Tiuryn [Tiu92] proved that SSI is PSPACE-hard.

Further Wand and O'Keefe observe that if partial order is known to be a tree, then the satisfiability problem is solvable in polynomial time. In fact, their algorithm for trees works in exponential rather than polynomial time.

Main result In this paper we prove a theorem from which it follows that the SSI problem for trees is indeed satisfiable in polynomial time. Therefore the type reconstruction problem for such posets is in PTIME.

Why trees are important The original motivation for studying orders being trees was that trees correspond to the coercion structure of single-inheritance object systems. Although traditionally, recursive types are used for inheritance mechanism, recent works show that it is possible to construct type systems for inheritance which do not rely on recursive types (cf. e.g. [PT92]).

On the other hand trees are present even if we consider simple coercion. Perhaps the simplest possible example is a numeric or enumeration type and a couple of its subrange types.

Example One may consider types (like in the language C) *signed int* $= -2^{15}..2^{15}-1$, *unsigned int* $= 0..2^{16}-1$, *long int* $= -2^{31}..2^{31}-1$ and coercions between them based on respective inclusions. Then the poset would look like

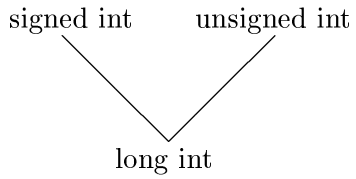

Figure 1: An example subtyping relation involving subrange types

Why trees aren't trivial The main problem is that since we consider non-flat inequalities, any result for trees actually has to cover a wider range of posets, since any algorithm working for given poset, must obviously work for poset formed of terms of given shape over this poset. Therefore the problem is in finding a class of posets wide enough to facilitate inductive proof, and narrow enough to fulfill satisfactory thesis.

Example Consider set of types from previous exmaple

$$C = \{signed\ int,\ unsigned\ int,\ long\}$$

and let

$$C^{\star\to\star} = \{\rho \to \tau : \rho, \tau \in C\}$$

then coercions between elements of $C^{\star\to\star}$ form a poset

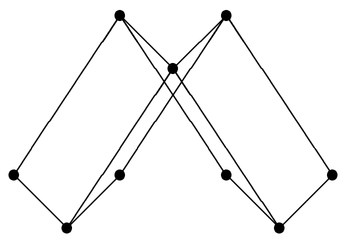

Figure 2: The poset $C^{\star\to\star}$

This poset seems all but trivial. In fact it bear some resemblance to the posets for which the type reconstruction is NP-hard [LM92] and the SSI problem is PSPACE-hard [Tiu92] — crowns

In fact, because of this resemblance it had been for some time conjectured that the SSI problem for trees is NP-hard. However, as this work proves, said resemblance was deceiving.

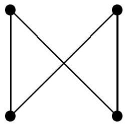

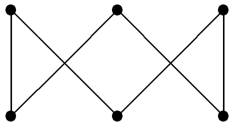

Figure 3: (a) 2-crown (b) 3-crown

Related work A lot of effort has been devoted to studying type reconstruction recently. Bruce and Mitchell [BM92] present models of subtyping, recursive types and higher order polymorphism with types interpreted as partial equivalence relations over a D_∞ lambda model. Tiuryn and Wand [TW93] study complexity of type reconstruction with subtyping and recursive types. Kozen, Palsberg and Schwartzbach [KPS92] analyze subtype relations for inferring partial types, where all types regardless of their rank are in relation with Ω, representing the 'unknown' type.

2 Definitions

2.1 The problem

Let $\langle Q, \leq \rangle$ be a finite poset. Let $\to$ be binary infix operation symbol, and let $\mathcal{T}_Q$ be the term algebra over $(Q, \to)$. Its carrier consist of finite terms without variables and is partially ordered by extending the order from Q to all terms by the following rules

$$\frac{r_1 \leq t_1 \quad t_2 \leq r_2}{(t_1 \to t_2) \leq (r_1 \to r_2)}$$

That means that we can think of "$\to$" as of an operation which is antimonotonic in the first argument and monotonic in second, whereas product is monotonic in both arguments.

A system Σ of inequalities is a finite set of formulas of the form

$$\Sigma = \{\tau_1 \leq \rho_1, \ldots, \tau_n \leq \rho_n\}$$

where τ's and ρ's belong to $\mathcal{T}_Q(V)$

Definition 1 *Σ is said to be* flat *if every term in it is of size 1 i.e. is either a constant or a variable. Note that if Σ is not flat then by its size we mean the total size of terms it contains, rather than the number of inequalities.*

Definition 2 *Σ is said to be* satisfiable *in $\mathcal{T}_Q$ if there is a valuation $v : V \to \mathcal{T}_Q$ such that $\tau_i[v] \leq \rho_i[v]$ holds in $\mathcal{T}_Q$ for all i.*

There are actually two problems which are addressed in this paper:

(Q-SSI) *Given a system Σ of inequalities, is it satisfiable in $\mathcal{T}_Q$?*

(Q-FLAT-SSI) *Given a flat system Σ of inequalities, is it satisfiable in $\mathcal{T}_Q$?*

2.2 Weak satisfiability

Definition 3 *Let $Q_1, \dots, Q_n$ be the connected components of Q. Let $\mathcal{T}_\star$ be the set of terms without variables over the set of constants $\{\star_1, \dots, \star_m\}$ and binary operation symbol $\rightarrow$. The elements of $\mathcal{T}_\star$ we call* shapes. *We use a canonical map $(\cdot)_\star : \mathcal{T}_Q(V) \rightarrow \mathcal{T}_\star(V)$ which maps all constants to $\star$ but leaves the variables untouched. We call $(\tau)_\star$ the shape of τ. It is defined as follows:*

$$(c)_\star = \star_i \quad \text{if } c \in Q_i$$

$$(v)_\star = v \quad \text{if } v \in V$$

$$(\tau \rightarrow \rho)_\star = (\tau)_\star \rightarrow (\rho)_\star$$

For $\sigma \in \mathcal{T}_\star$ let

$$Q^\sigma = \{\rho \in \mathcal{T}_Q : (\rho)_\star = \sigma\}$$

A system of inequalities $\Sigma = \{\tau_1 \leq \rho_1, \dots, \tau_n \leq \rho_n\}$ is said to be weakly satisfiable *if $\Sigma_\star = \{(\tau_1)_\star = (\rho_1)_\star, \dots, (\tau_n)_\star = (\rho_n)_\star\}$ is satisfiable in $\mathcal{T}_\star$.*

Weak satisfiability is clearly a necessary condition for satisfiability. It is decidable in (and in fact complete for) polynomial time since it is an instance of the unification problem.

2.3 Distances and disks in a poset [1]

Definition 4 *By an* up-fence *of length n from x to y we mean a set $\{x_0, \dots, x_n\}$, such that $x_0 = x, x_n = y$, and $x_{2i} \leq x_{2i+1} \geq x_{2i+2}$*

The up-distance from x to y in Q, *$d^{+1}(x,y)$ is the smallest number n such that there exists an up-fence of length n from x to y. If such n does not exist, we assume $d^{+1}(x,y) = +\infty$*

Definition 5 *By a* down-fence *of length n from x to y we mean a set $\{x_0, \dots, x_n\}$, such that $x_0 = x, x_n = y$, and $x_{2i} \geq x_{2i+1} \leq x_{2i+2}$*

The down-distance from x to y in Q, *$d^{-1}(x,y)$ is the smallest number n such that there exists an up-fence of length n from x to y. If such n does not exist, we assume $d^{-1}(x,y) = +\infty$*

In the sequel, we shall use distances introduced above in the context

$$Q \models d^{\pm 1}(x,y) \leq n$$

by which we mean that the respective distance in Q does not exceed n.

Definition 6 *By a* disk in Q *with the center a and radius (r^{+1}, r^{-1}) we shall mean the set*

$$D_Q(a, (r^{+1}, r^{-1})) = \{x \in Q : Q \models d^\varepsilon(a,x) \leq r^\varepsilon\}, \quad \varepsilon \in \{-1, 1\}$$

[1] The definitions and facts in this and the following section come from [NR85]

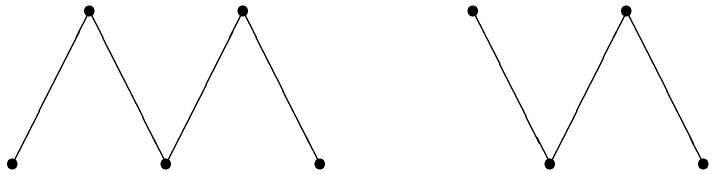

Figure 4: Fences: (a) an up-fence of length 4 (b) a down-fence of length 3

2.4 Retractions and absolute retracts

We say that $P \supseteq Q$ retracts to Q ($P \rhd Q$) if there exists an order preserving and idempotent map $f : P \to Q$. Nevermann and Rival [NR85] show a class of posets for which exists a simple condition which is necessary and sufficient for retractibility.

Definition 7 *A set*

$$\mathcal{H} = \{(v_t, \delta_t^{+1}, \delta_t^{-1}) : t \in T\}$$

where T is an index set, $v_t \in Q$, is a hole in Q if

$$\bigcap_{t \in T} D_Q(v_t, \delta_t^{+1}, \delta_t^{-1}) = \emptyset$$

and $\mathcal{H}$ contains no proper and nonempty subfamily of disks with empty intersection.

We say that this hole is separated in $P \supseteq Q$ if

$$\bigcap_{t \in T} D_P(v_t, \delta_t^{+1}, \delta_t^{-1}) = \emptyset$$

Definition 8 *A set Q is called an absolute retract if it is a retract of any ordered set in which all its holes are separated.*

Definition 9 *A partially ordered set satisfies the the two disk property (alias Helly property) if, for each family $\mathcal{D}$ of disks*

$$\bigcap \mathcal{D} \neq \emptyset$$

whenever

$$D_1 \cap D_2 \neq \emptyset$$

for each D_1, D_2 in $\mathcal{D}$.

Proposition 1 *[NR85] Every Helly poset is an absolute retract. Helly posets are closed under products, retractions and disjoint unions.*

Corollary 2 *If Q is a Helly poset, then so is Q^σ for any shape σ.*

It is not difficult to see that trees (and disjoint unions of them) posess the two disk property and are therefore absolute retracts. On the other hand, crowns, which are known to be intractable are not absolute retracts.

2.5 Ground consistency

Given system of ineqaualities we introduce a deduction system which allows to infer inequalities that must be satisfied by any solution of Σ [Tiu92].

$$\Sigma \vdash_s \tau \leq \rho \quad \text{ for } \tau \leq \rho \in \Sigma \tag{1}$$

$$\frac{\Sigma \vdash_s \tau \leq \rho \quad \Sigma \vdash_s \rho \leq \sigma}{\Sigma \vdash_s \tau \leq \sigma} \tag{2}$$

$$\frac{\Sigma \vdash_s (\tau_1 \to \tau_2) \leq (\rho_1 \to \rho_2)}{\Sigma \vdash_s \rho_1 \leq \tau_1, \quad \tau_2 \leq \rho_2} \tag{3}$$

$$\tag{4}$$

Definition 10 *We call* Σ ground consistent *whenever for each* $c_1, c_2 \in Q$ *if* $\Sigma \vdash_s c_1 \leq c_2$*, then* $Q \models c_1 \leq c_2$.

Theorem 3 *Ground consistency is complete for NLOGSPACE wrt NC1 reductions*

Proof: First we are going to prove that ground consistency is co-NLOGSPACE hard. But, as observed by Immerman,

$$co - NLOGSPACE = NLOGSPACE$$

Thus ground consistency is NLOGSPACE-hard.

The proof goes by reduction of the reachability in directed graph. Given a graph $G =< V, E >$, and a pair of vertices $a, b \in V$, take $Q = \{a, b\},\ a \leq b$. All other vertices are treated as constants. For each $(x, y) \in E$ put an inequality $x \leq y$ in Σ. Thus defined, Σ is ground consistent iff there is no path from b to a in g.

Note that the reduction describede above can be in fact carried out in NC1. Thus ground consistency is NLOGSPACE-hard with regard to NC1 reductions

To see that ground consistency is indeed in NLOGSPACE consider the algorithm which we exhibit in more detail for the flat case:

- nondeterministically choose a pair of constants a, b such that $a \leq b$ does not hold in Q
- check that $\Sigma \vdash_s a \leq b$ by nondeterministically choosing inequalities from Σ that form a chain fom a to b:

  $$a \leq \alpha_1, \alpha_1 \leq \alpha_2, \alpha_n \leq_b,$$

  where $\alpha_i \in Q \cup V$ for $i = 1, \ldots, n$

Indeed, logarithmic space is needed to record consecutive α's, since we only need to record the current element, rather than the whole chain.

To extend this algorithm for general case, we should systematically number subterms of Σ, then guess the pair of constants and nondeterministically detect the inference that leads to an inconsistency.

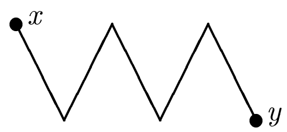

Figure 5: An example for the [Symm] rule: $d^{+1}(y,x) \leq 5$, $d^{-1}(x,y) \leq 5$

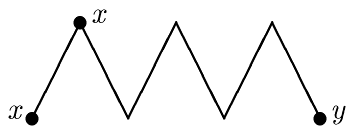

Figure 6: An example for the [Dual] rule: $d^{-1}(x,y) \leq 5$, but since $x \leq x$, we have $d^{+1}(x,y) \leq 6$

2.6 Distance consistency

Given system of inequalities we introduce a deduction system which allows to infer facts about distances between elements that must be satisfied by any solution of Σ. (It is easy to prove that the system is sound; cf. figures 2.6 through 2.6 for intuitions behind the rules)

$$\text{[Ax]} \qquad \Sigma \vdash_d d^{+1}(x,y) \leq 1 \quad \text{for } x \leq y \in \Sigma \tag{5}$$

$$\text{[Symm]} \qquad \frac{\Sigma \vdash_d d^{\varepsilon}(x,y) \leq n}{\Sigma \vdash_d d^{(\varepsilon(-1)^n)}(y,x) \leq n} \tag{6}$$

$$\text{[Dual]} \qquad \frac{\Sigma \vdash_d d^{\varepsilon}(x,y) \leq n}{\Sigma \vdash_d d^{(-\varepsilon)}(x,y) \leq n+1} \tag{7}$$

$$\text{[Join]} \qquad \frac{\Sigma \vdash_d d^{\varepsilon}(x,y) \leq n,\ d^{\varepsilon(-1)^{n-1}}(y,z) \leq k}{\Sigma \vdash_d d^{\varepsilon}(x,z) \leq n-1+k} \tag{8}$$

$$\text{[Subterm]} \qquad \frac{\Sigma \vdash_d d^{\varepsilon}(x_1 \to x_2, y_1 \to y_2) \leq n}{\Sigma \vdash_d d^{-\varepsilon}(x_1,y_1) \leq n, \Sigma \vdash_d d^{\varepsilon}(x_2,y_2) \leq n} \tag{9}$$

Definition 11 *We call* Σ distance consistent *whenever for each* $c_1, c_2 \in Q$ *and* $n \in N$ *if* $\Sigma \vdash_d d^{\varepsilon}(c_1,c_2) \leq n$, *then* $Q \models d^{\varepsilon}(c_1,c_2) \leq n$.

Distance consistency is decidable in polynomial time. We can construct two $O(|\Sigma|) \times O(|\Sigma|)$ arrays, for d^{+1} and d^{-1} respectively in such a way that in the array for d^{ε}, in the cell indexed by τ and ρ there is the least number n such that $\Sigma \vdash_d d^{\varepsilon}(\tau,\rho) \leq n$ holds, or (a special symbol for) $+\infty$ if no such n exists.

In fact, even the following holds:

Proposition 4 *distance consistency is NLOGSPACE complete*

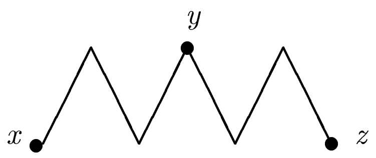

Figure 7: An example for the [Join] rule: $d^{+1}(x,y) = 3$, $d^{-1}(y,z) = 3$, joined they yield $d^{+1}(x,z) = 6$

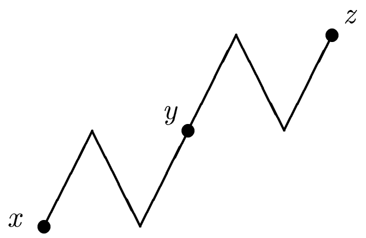

Figure 8: Another example for the [Join] rule: $d^{+1}(x,y) \leq 3$, $d^{+1}(y,z) \leq 3$, yield when joined $d^{+1}(x,z) \leq 5$

Proof: Hardness is obvious, since ground consistency is but a special case of distance consistency. On the other hand, an algorithm to check the latter can be constructed easily, following the one descri-bed in the proof of Theorem 3. The only difference is that the direction of inequalities may vary throughout the chain; we must maintain a counter for the number of such alternations. For this counrter actually only a constant space is needed, since maximum distance is fixed with the given poset, and we are looking for fences *shorter* than given distance.

3 The result

The main result presented in this paper is that a system of inequalities is satisfiable over a poset satisfying the Helly property is satisfiable iff it is weakly satisfiable and distance consistent. Thus satisfiability of such systems is decidable in polynomial time.

Lemma 5 $\Sigma \vdash_s x \leq y$ *iff* $\Sigma \vdash_d d^{+1}(x,y) \leq 1$

Proof: The left-to-right implication follows from axiom 5. The opposite one can be proved by simple induction on inference of $\Sigma \vdash_d d^{+1}(x,y) \leq 1$. □

Corollary 6 *If Σ is distance consistent then it is also ground consistent.*

Let Σ be ground consistent. Let us define the following equivalence relation $\simeq$ on $var(\Sigma)$:

$$x \simeq y \iff \Sigma \vdash_s x \leq y \wedge \Sigma \vdash_s y \leq x$$

By an *extension of Q generated by Σ* we shall mean the set

$$Q_\Sigma = Q \cup var(\Sigma)/\simeq$$

with the order defined as the transitive closure of the sum of the order on Q and the order on $Q \cup var(\Sigma)/\simeq$, induced by $\Sigma \vdash_s$.

Lemma 7 *A flat system of inequalities Σ is satisfiable if and only if Σ is ground consistent and $Q_\Sigma \rhd Q$*

Proof: Each retraction $f : Q_\Sigma \to Q$ determines a solution of Σ:

$$\sigma(x) = f([x]_\simeq)$$

On the other hand, if σ is a solution of Σ, then

$$x \simeq y \Rightarrow \sigma(x) = \sigma(y)$$

$\Box$

Lemma 8 *Let Σ — ground consistent. Then:*

(a) *If $Q_\Sigma \models [x] \leq c$ for some $x \in var(\Sigma)$, $c \in Q$ then there exists $c' \in Q$ such that*

$$\Sigma \vdash_s x \leq c' \quad Q \models c' \leq c$$

(b) *If $Q_\Sigma \models c \leq [x]$ for some $x \in var(\Sigma)$, $c \in Q$ then there exists $c' \in Q$ such that*

$$\Sigma \vdash_s c' \leq x \quad Q \models c \leq c'$$

(c) *If $Q_\Sigma \models [x_1] \leq [x_2]$ for some $x_1, x_2 \in var(\Sigma)$ then either $\Sigma \vdash_s x_1 \leq x_2$ or there exist $c_1, c_2 \in Q$ such that*

$$\Sigma \vdash_s x_1 \leq c_1 \quad Q \models c_1 \leq c_2 \quad \Sigma \vdash_s c_2 \leq x_2$$

Proof: (a) Since order on Q_Σ is the transitive closure of the sum of order on Q and the one induced by $\Sigma \vdash_s$, there exist $p_0, p_1, \ldots, p_n$ such that $p_0 = [x]$, $p_n = c$ and for all $1 \leq i \leq n$ either $p_{i-1} = [w_{i-1}]$, $p_i = [w_i]$, $\Sigma \vdash_s w_{i-1} \leq w_i$ or $Q \models p_{i-1} \leq p_i$. Let k be the smallest such that $p_k \in Q$. From ground consistency of Σ it follows that $p_k \leq c$ holds in Q. On the other hand for all $0 \leq i \leq k-1$, $p_i = [w_i]$ for some $w_i \in var(\Sigma)$ and thus by the rule 2 we have $\Sigma \vdash_s x \leq p_k$ $\Box$

The proof of (b) is analogous, and to prove (c) we should use the same argument twice. It is worth mewntioning that we cannot get rid of the second part of the alternative in (c), i.e. it is not always true that if $Q_\Sigma \models [x_1] \leq [x_2]$ then $\Sigma \vdash_s x_1 \leq x_2$. $\blacksquare$

Lemma 9 *If Σ is distance consistent than we have for any $c_1, c_2 \in Q$*

$$Q_\Sigma \models d^\varepsilon(c_1, c_2) \leq n \iff Q \models d^\varepsilon(c_1, c_2) \leq n$$

Proof: The implication ($\Leftarrow$) is obvious. To prove ($\Rightarrow$) assume the contrary, i.e. that

$$Q_\Sigma \models d^\varepsilon(c_1, c_2) \leq n$$

and

$$\neg Q \models d^\varepsilon(c_1, c_2) \leq n.$$

We show that there exist $d_1, d_2 \in Q$, $m \leq 0$ such that

$$\Sigma \vdash_d d^\varepsilon(d_1, d_2) \leq m$$

and

$$\neg Q \models d^\varepsilon(d_1, d_2) \leq m.$$

The proof is by induction on n. The basis of induction follows from the lemma 8 (or from the dual lemma for d^{-1}). The induction step goes as follows

Let $c_1 = x_0, x_1, \ldots, x_n = c_2$ be a fence in Q_Σ, connecting c_1 with c_2. If $x_1 \in Q$, then we may use the induction hypothesis since

$$Q_\Sigma \models d^{-\varepsilon}(x_1, c_2) \leq n - 1$$

$$\neg Q \models d^{-\varepsilon}(x_1, c_2) \leq n - 1$$

On the other hand, if $x_1 = [v]$, $v \in var(\Sigma)$, then choose largest i such that $\Sigma \vdash_d d^{-\varepsilon}(x_1, x_i) \leq i - 1$ (or $i = 1$, if no such exists). From the previous lemma it follows that there exists $d_1 \in Q$ such that $Q \models d^\varepsilon(c_1, d_1) \leq 1$ and $\Sigma \vdash_d d^\varepsilon(d_1, x_1) \leq 1$. Similarly, if $x_1 \notin Q$, then there exists $d_2 \in Q$ such that $\Sigma \vdash_d d^{\varepsilon(-1)^i}(x_i, d_2) \leq 1$.

Now, one of the following must hold:

$$\neg Q \models d^\varepsilon(d_1, d_2) \leq i$$

or

$$Q_\Sigma \models d^{\varepsilon(-1)^i}(d_2, c_2) \leq n - i - 1,$$

$$\neg Q \models d^{\varepsilon(-1)^i}(d_2, c_2) \leq n - i - 1$$

The first case obviously contradicts the distance consistency of Σ and the second one allows us to use the induction hypothesis $\square$

Theorem 10 *A flat system of inequalities over a poset satisfying the Helly property is satisfiable if and only if it is distance consistent.*

Proof: Again, the left-to-right implication is pretty obvious. By the lemma 7, in order to prove the opposite one, it suffices to prove that $Q_\Sigma \rhd Q$. Furthermore, since Q satisfies the Helly property and thus is an absolute retract, all we have to prove that all holes in Q of cardinality 2 are separated in Q_Σ. Let us assume the contrary, i.e. that there exists a pair of disks which is a hole in Q that is not separated in Q_Σ.

So we may chose $a_1, a_2 \in Q, r_1^{+1}, r_1^{-1}, r_2^{+1}, r_2^{-1} \in N, x \in Q_\Sigma$ such that

$$D_Q(a_1, r_1^{+1}, r_1^{-1}) \cap D_Q(a_2, r_2^{+1}, r_2^{-1}) = \emptyset$$

$$x \in D_{Q_\Sigma}(a_1, r_1^{+1}, r_1^{-1}) \cap D_{Q_\Sigma}(a_2, r_2^{+1}, r_2^{-1})$$

We can further assume that the sum $r_1^{+1} + r_1^{-1} + r_2^{+1} + r_2^{-1}$ is minimal.

If there is $x \in Q$ meeting assumptions, we can use Lemma 9 to show that Σ is not distance consistent. Otherwise $x = [v]$. Now, if $\Sigma \vdash_d d^\varepsilon(a_1, x) \leq r_1^\varepsilon$ and $\Sigma \vdash_d d^\theta(a_2, x) \leq r_2^\theta$ and by the rules [Join] and [Symm] rules we can infer that $\Sigma \vdash_d d^\varepsilon(a_1, a_2) \leq r$ while $Q \models d^\varepsilon(a_1, a_2 \geq r)$, where $r = r_1^\varepsilon + r_2^\theta$ if $(-1)^{r_1^\varepsilon} = (-1)^{r_1^\theta}$, and $r = r_1^\varepsilon + r_2^\theta$ otherwise. That contradicts the distance consistency of Σ.

If the opposite holds, we can assume without loss of generality that

$$\Sigma \not\vdash d^\varepsilon(a_1, v) \leq r_1^\varepsilon$$

Let $a_1 = x_0, x_1, \ldots, x_n = [v]$ be a fence in Q_Σ, connecting a_1 with $[v]$, and i be the largest integer such that

$$\Sigma \not\vdash d^\varepsilon(x_i, [v]) \leq r_1^\varepsilon - i$$

Note that $i < r_1^\varepsilon$.

In other words we have

$$\Sigma \vdash_d d^\alpha(x_{i+1}, v) \leq r_1^\varepsilon - i - 1$$

$$\Sigma \not\vdash d^{-\alpha}(x_i, x_{i+1}) \leq 1$$

$$Q_\Sigma \models d^{-\alpha}(x_i, x_{i+1}) \leq 1$$

If $i = 0$, then by Lemma 8 there exists $a' \in Q$ such that $Q \models d^\varepsilon(a_1, a') \leq 1$ and $\Sigma \vdash_d d^\varepsilon(a, x_1') \leq 1$. Thus we can consider a' instead of a_1, and $i = 1$ instead of 0.

Neither x_i nor x_{i+1} can be a constant from Q, since then we could place the center of disk in x_i, thus obtaining a pair of disks, fulfilling assumptions but with smaller sum of radiuses, which would contradict minimality. But if $x_{i+1} = [w]$, by the Lemma 3 there exists constant c such that $\Sigma \vdash_d d^\alpha(d, w) \leq 1$, which leads again to the contradiction with minimality. □

Corollary 11 *FLAT-SSI is NLOGSPACE complete for Helly posets.*

This result can be lifted to arbitrary systems of inequalities leading us to the following

Theorem 12 *A system of inequalities Σ over a poset satisfying the Helly property is satisfiable iff it is weakly satisfiable and distance consistent.*

Proof: [2]

The $(\Rightarrow)$ implication is obvious. The opposite implication is proved by induction on the number of equivalence classes of $\sim$ defined on $var(\Sigma)$ as follows

$$x \sim y \quad \text{iff} \quad \Sigma_* \models x = y$$

Suppose first that the quotient set $var(\Sigma)/\sim$ has only one element. Then we consider the set

$$\hat{\Sigma} = \{\xi \leq \xi' : \Sigma \vdash_s \xi \leq \xi' \text{ and } |\xi| = 1 \text{ or } |\xi'| = 1\}$$

It is easy to prove that Σ is satisfiable iff $\hat{\Sigma}$ is satsfiable. In fact above two systems have the same solutions. It follows that if τ is a term of depth greater than 1 which occurs in $\hat{\Sigma}$ then it must be

[2]The proof deliberately follows the one in [Tiu92] almost to the letter where possible.

a term without variables. Thus Σ can be viewed as a flat system of inequalities. Since Σ is weakly satisfiable and distance consistent it follows that $\hat{\Sigma}$ is too. Hence it has a solution and so has Σ.

Now let us assume that $var(\Sigma)/\sim$ has $n+1$ elements and let $y \in var(\Sigma)$ be such that for no $z \in var(\Sigma)$ there is a term τ with $|\tau| > 1$, $z \in var(\tau)$ and $\Sigma_* \models \tau = z$. This would contradict weak satisfiability of Σ

Let

$$[y] = \{z \in var(\Sigma) | z \sim y\} = \{y_1, \ldots, y_k\}$$

and consider the set

$$\hat{\Sigma} = \{\xi \leq \xi' : \Sigma \vdash_s \xi \leq \xi' \text{ and } |\xi| \in [y] \text{ or } |\xi'| \in [y]\}$$

It follows from the choice of y that if τ is a non variable term which occurs in $\hat{\Sigma}$, then it contains no variables. Moreover all such terms must be of the same shape, say σ. If there are no such terms then we choose shape σ arbitrarily.

It follows that $\hat{\Sigma}$ can be viewed as a flat system of inequalities over $\mathcal{T}_\sigma$ which satisfies the Helly property. It is weakly satisfiable and distance consistent. Let $\hat{v} : [y] \to \mathcal{T}_\sigma$ be a solution of $\hat{\Sigma}$ and let

$$\Sigma_1 = \hat{v}(\Sigma) = \{\hat{v}(\tau) \leq \hat{v}(\rho) : \tau \leq \rho \in \Sigma\}$$

In the above definition $\hat{v}$ acts as an identity on variables other than those in $[y]$. Let $\sim_1$ be the equivalence relation associated with Σ_1. We claim

$$|var(\Sigma_1)/\sim_1| = n$$

$$\Sigma_1 \text{ is weakly satisfiable}$$

These are easy to prove since $(\Sigma_1)_* = \Sigma_*[\sigma/y_1, \ldots, \sigma/y_k]$, and if $(\Sigma_1)_* \models \tau = \rho$, then $\Sigma_* \models \tau = \rho$.

To complete the proof we need to prove that

$$\Sigma_1 \text{ is distance consistent}$$

Let τ be a term such that it contains no constants, every variable of τ occurs exactly once, $var(\tau) \cap var(\Sigma) = \emptyset$, and for some function $\tilde{v} : var(\tau) \to \{*_1, \ldots, *_m\}$, $\tilde{v}(\tau) = \sigma$. Let $\tau_1, \ldots, \tau_k$ be terms obtained from τ by renaming its variables so that they have pairwise disjoint variables and disjoint from $var(\Sigma)$. Let $\Sigma_2 = \Sigma[\tau_1/y_1, \ldots, \tau_k/y_k]$ and let $\eta : \bigcup_{i=1}^{k} var(\tau_i) \to Q$ be a function such that $\eta(\tau_i) = \hat{v}(y_i)$, for $i = 1, \ldots, k$. Such a function is uniquely determined by the above conditions.

In order to show distance consistency of Σ_1 we need one more definition.

Definition 12 *A sequence $\xi_1, \ldots, \xi_{2\ell}$ of terms in $\mathcal{T}_C(X)$ is called an η-*chain *(of fences) if the following conditions hold*

$$\Sigma_2 \vdash_d \; d^{\varepsilon_i}(\xi_i, \xi_{i+1}) \leq n_i, \quad \textit{for even } i < 2\ell \tag{10}$$

$$\eta(\xi_i) = \eta(\xi_{i+1}), \quad \textit{for odd } i < 2\ell \tag{11}$$

ℓ is called the size *of the chain, while its length is computed according to the rule [Join], i.e. adding the lengths of its links and subtracting the number of joints where no direction alteration occurs (cf. figures 2.6 and 2.6*

Equality in (11) represents syntactic identity of terms. Again, we assume that in the above formula η acts as identity on variables other than those in $\bigcup_{i=1}^{k} var(\tau_i)$. Now, distance consistency of Σ_1 follows immediately from the following two claims.

$$\begin{array}{c}\text{For all terms } \tau, \rho, \text{ if } \Sigma_1 \vdash_d d^{\varepsilon}(\tau,\rho) \leq n, \text{ then there} \\ \text{exists an } \eta\text{-chain of length } n \text{ from } \tau \text{ to } \rho \text{ with } \varepsilon_2 = \varepsilon\end{array} \tag{12}$$

$$\begin{array}{c}\text{For all constants } a, b \in Q, \text{ if there is an } \eta\text{-chain of} \\ \text{length } n, \text{ with } \varepsilon_2 = \varepsilon \text{ from } a \text{ to } b, \text{ then } Q \models d^{\varepsilon}(a,b) \leq n\end{array} \tag{13}$$

The proof of (13) is by an obvious induction on the size of η-chain. We sketch the proof of (12). We use induction on the length of derivation of $d^{\varepsilon}(\tau,\rho) \leq n$ from Σ_1.

If $\tau \leq \rho$ is in Σ_1, then there is an inequality $\tau' \leq \rho'$ in Σ_2 such that $\eta(\tau') = \tau$ and $\eta(\rho') = \rho$. Thus τ, τ', ρ', ρ is an η-chain with both length and size of 1.

The η-chain for the rule [Symm] is obtained simply by chasing the chain from its premise backwards. For the [Dual] rule we simply apply the same rule to the first chain link in η-chain, obtaining dual η-chain of length greater by one than the original one:

$$\Sigma_2 \vdash_d \ d^{\varepsilon_2}(\xi_2,\xi_3) \leq n_2$$

then by [Dual]

$$\Sigma_2 \vdash_d \ d^{-\varepsilon_2}(\xi_2,\xi_3) \leq n_2 + 1$$

If the last rule in the derivation of $\Sigma_1 \vdash_d d^{\varepsilon}(\tau,\rho) \leq n$ was [Join], then we have $\Sigma_1 \vdash_d d^{\varepsilon}(\tau,\xi) \leq n$ and $\Sigma_1 \vdash_d d^{\varepsilon(-1)^{n-1}}(\xi,\rho) \leq k$ We have two η-chains: one from τ to ξ and another from ξ to ρ. Their concatenation yields an η-chain from τ to ρ. Its length is as stated in the [Join] rule.

Finally, if the last rule was [Subterm], then $\Sigma_1 \vdash_d d^{\varepsilon}(\tau,\rho) \leq n$ must have been obtained from, say $\Sigma_1 \vdash_d d^{\theta}(\xi,\xi') \leq n$. Let $\xi_1,\ldots,\xi_{2\ell}$ be an η-chain of length n from ξ to ξ' and assume that it is one of mimimal size. If all terms in that chain are of depth greater than 1, then by applying the rule [Subterm] throughout the chain of fences we get an η-fence from τ to ρ. Otherwise, let ξ_j be a term of depth 1. It cannot be a constant nor a variable in $\bigcup_{i=1}^{k} var(\tau_i)$ since $|\xi| > 1$. Hence $\xi_j \in var(\Sigma) - [y]$. If j is odd then since $|\xi| > 1$ and $|\xi'| > 1$, it follows that $1 < j < 2\ell - 1$ and

$$\Sigma_2 \vdash_d \ d^{\varepsilon_{j-1}}(\xi_{j-1},\xi_j) \leq n_i, \quad \text{and} \quad \Sigma_2 \vdash_d \ d^{\varepsilon_{j+1}}(\xi_{j+1},\xi_{j+2}) \leq n_i$$

Moreover, $\eta(\xi_j) = \eta(\xi_{j+1})$. Thus ξ_j and ξ_{j+1} are identical variables and we obtain a contradiction since the η-chain can be shortened (by removing one η-link). The case when j is even is argued similarly. This proves that all terms in the η-chain must be of depth greater than 1, thereby completing the proof of (12). ∎

4 Further work

Nevermann and Rival [NR85] claim (without giving examples) that Helly property is not a necessary condition for being an absolute retract. Although it is easily seen that FLAT-SSI for absolute retracts is in PTIME, it remains open whether so does SSI. My conjecture is that it is indeed so.

Furthermore the problem of further classification of posets presents na interesting field for study: it seems that the posets can be divided into two classes: tractable posets, for which flat problem is NLOGSPACE complete and the general case is P-complete, and the intractable ones (like crowns), for which the flat problem is NP-complete and the general is PSPACE-hard. There are no known examples of posets falling "between" this classes, and our conjecture is that there are in fact no such posets.

References

[Ben93] Marcin Benke. Efficient type reconstruction in the presence of inheritance (extended abstract). In S. Sokołowski A. M. Borzyszkowski, editor, *MFCS'93: Mathematical Foundations of Computer science, Proc. 18th Intern. Symp., Gdansk 1993*, pages 272–280. Springer-Verlag, 1993. LNCS 711.

[BM92] Kim Bruce and John C. Mitchell. Per models of subtyping, recursive types and higher-order polymorphism. In *Conf. Rec. ACM Symp. Principles of Programming Languages*, 1992.

[KPS92] Dexter Kozen, Jens Palsberg, and Michael I. Schwartzbach. Efficient inference of partial types. Technical Report DAIMI PB-394, Computer science Dept., Aarhus University, April 1992.

[LM92] Patrick Lincoln and John C. Mitchell. Algorithmic aspects of type inference with subtypes. In *Conf. Rec. ACM Symp. Principles of Programming Languages*, pages 293–304, 1992.

[Mit84] John C. Mitchell. Coercion and type inference. In *Conf. Rec. ACM Symp. Principles of Programming Languages*, pages 175–185, 1984.

[NR85] P. Nevermann and I. Rival. Holes in ordered sets. *Graphs and Combinatorics*, (1):339–350, 1985.

[PT92] Benjamin C. Pierce and David N. Turner. Object-oriented programming without recursive types. Technical Report ECS-LFCS-92-225, LFCS, University of Edinburgh, August 1992.

[Tiu92] Jerzy Tiuryn. Subtype inequalities. In *Proc. 7th IEEE Symp. Logic in Computer Science*, pages 308–315, 1992.

[TW93] J. Tiuryn and M. Wand. Type reconstruction with recursive types and atomic subtyping. In M.-C. Gaudel and J.-P. Jouannaud, editors, *TAPSOFT'93: Theory and Practice of Software Development, Proc. 4th Intern. Joint Conf. CAAP/FASE*, pages 686–701. Springer-Verlag, 1993. LNCS 668.

[WO89] M. Wand and Patrick O'Keefe. On the complexity of type inference with coercion. In *Proc. ACM Conf. Functional Programming and Computer Architecture*, 1989.