Study of the Subtyping Machine of Nominal Subtyping with Variance*

ORI ROTH, Technion—IIT, Israel

This is a study of the computing power of the subtyping machine behind Kennedy and Pierce's nominal subtyping with variance. We depict the lattice of fragments of Kennedy and Pierce's type system and characterize their computing power in terms of regular, context-free, deterministic, and non-deterministic tree languages. Based on the theory, we present Treetop—a generator of C[#] implementations of subtyping machines. The software artifact constitutes the first feasible (yet POC) fluent API generator to support context-free API protocols in a decidable type system fragment.

CCS Concepts: • **Software and its engineering** → **General programming languages; Inheritance; API languages; Software development techniques.**

Additional Key Words and Phrases: subtyping, variance, metaprogramming, fluent API, DSL

ACM Reference Format:

Ori Roth. 2021. Study of the Subtyping Machine of Nominal Subtyping with Variance. *Proc. ACM Program. Lang.* 5, OOPSLA, Article 137 (October 2021), 27 pages. <https://doi.org/10.1145/3485514>

1 INTRODUCTION

Object-oriented programming is all about inheritance or subtyping, and how software can and cannot be properly engineered using it. In the eyes of the theoretical computer scientist or the type theorist, the question is what (compile-time) computations can and cannot be done with subtyping. The question becomes challenging when generics, i.e., parametric polymorphism, interact with subtyping, raising the question of when a certain generic taking type argument t_1 is a subtype of the same generic taking type argument t_2 . In the invariant regime of subtyping with generics, the answer is negative (unless $t_1 = t_2$); in the covariant regime, the answer is positive precisely when t_1 is a subtype of t_2 ; in contrast, in the contravariant regime, when t_2 is a subtype of t_1 .

Consider, for example, the C[#] [Hejlsberg et al. 2003] code in List. 1 (unless said otherwise, all code examples in this paper are written in C[#]):

```
1 interface a<out x> {} interface b<out x> {} interface E {}
2 class v0<x> : a<v0<a<x>>>, a<a<x>>, a<x>, b<v0<b<x>>>, b<b<x>>, b<x> {}
3 a<b<b<a<b<b<a<E>>>>>> w1 = new v0<E>(); // Compiles: abbabba is a palindrome
4 a<b<b<a<b<a<a<E>>>>>> w2 = new v0<E>(); // Does not compile: abbaaba is not a palindrome
```

Listing 1. C[#] subtyping machine recognizing palindromes over $\{a, b\}$

*See a full version of this article [Roth 2021a] and the accompanied artifact [Roth 2021b].

Author's address: Ori Roth, Technion—IIT, Haifa, Israel, soriroth@cs.technion.ac.il.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/10-ART137

<https://doi.org/10.1145/3485514>

The program in [List. 1](#) defines four classes (lines 1–2): generic (polymorphic) classes `a` and `b` employing a covariant type parameter `x`, as indicated by the keyword `out`; non-generic class `E`; and generic class `v0` that inherits from the six supertypes written after the colon. Below the classes are two variable declarations (lines 3–4). The type of variable `w1` is `a<b<b<a<b<b<a<E>>>>>>>`, encoding the word $w_1 = abbabba$, and the type of variable `w2` is `a<b<b<a<b<a<a<E>>>>>>>`, encoding the word $w_2 = abbabaa$. [List. 1](#) demonstrates the encoding of ℓ , the non-deterministic context-free formal language of palindromes over alphabet $\{a, b\}$, in the type system of $C^\#$. Each variable is assigned an instantiation of type `v0<E>`, that type checks if and only if the variable type encodes a palindrome. Thus, the assignment to variable `w1` type checks since $w_1 \in \ell$, but the assignment to `w2` fails since $w_2 \notin \ell$. The mechanics behind this “magic”, and why it matters, is revealed by our study of the computing power of subtyping in the context of the work of [Kennedy and Pierce \[2007\]](#) on the *nominal subtyping with variance* type system.

1.1 Kennedy and Pierce’s Type System and Its Decidable Fragments

Fig. 1 places our results in context. The figure shows a three-dimensional lattice, structured similarly to the λ -cube [Barendregt 1991]. The lattice is spanned by three type system features, which are explained briefly here and in greater detail in Section 3. The head of the lattice $\mathcal{S}_\top$ is the abstract type system developed by Kennedy and Pierce [2007], which employs all three features. The other lattice points are type system fragments of $\mathcal{S}_\top$ that use different feature combinations. Nodes are adorned with references to previous works (shown as citations) and to our contributions (shown as forward references).

Fig. 1. Lattice describing the eight subtyping type system fragments employing contravariance (C), expansive inheritance (X), and multiple instantiation inheritance (M) (defined in Section 3)

The lattice is spanned by three orthogonal Boolean coordinates or features denoted by C, X, and M, representing the three core features of $\mathcal{S}_\top$: *Contravariance*, *eXpansively-recursive inheritance*, and *Multiple instantiation inheritance*. The most familiar feature is the kind of variance; the horizontal direction in the figure denotes the distinction between covariant type systems (positioned on the plane on the left-hand side of the lattice) and contravariant type systems (right-hand).

Precise definitions of the three features follow in [Section 3](#); here we provide a quick review of the concepts and the intuition behind them:

(C) *Contravariance*: Type arguments assigned to contravariant type parameters can be substituted with less derived types. For example, using the “*is a*” relationship we can say that if *x* in *a*<*x*> is contravariant and *b* is *a* *c* then *a*<*c*> is an *a*<*b*>.

```

1 interface a<in x> {} // type parameter x of class a is contravariant
2 interface b : c {} // type b is a subtype of c
3 a<b> x=(a<c>) null; // type a<c> is a subtype of a<b>

```

(X) *Expansively-recursive inheritance*: The expansion of types through inheritance (as defined by Viroli [2000]) is unbounded. Intuitively, this expansion involves the *curiously recurring template pattern* [Coplien 1996], where a class appears in one of its supertypes’ generic arguments.

```

1 interface a<x, y> : b<a<b<x>, y>> {} // class a’s inheritance is expansively-recursive

```

(M) *Multiple instantiation inheritance*: Inheritance is non-deterministic [Kennedy and Pierce 2007], i.e., type arguments of supertypes are not unique.

```

1 interface a : b<c>, b<d> {} // supertype b is instantiated twice, with arguments c and d

```

A point in the lattice is denoted by the set of its features. Type system $\mathcal{S}_\top$ is denoted by the triple $\langle C, X, M \rangle$, and the notation $\langle C, X \rangle$ designates a type system that allows contravariance and expansive inheritance, but not multiple instantiation inheritance. The code examples above also demonstrate two key features ingrained within $\mathcal{S}_\top$ and its fragments: “polyadic” parametric polymorphism and multiple inheritance, allowing a class to have multiple type parameters and supertypes, respectively.

As shown in Fig. 1, the new results here, together with the proof that $\mathcal{S}_\top$ is undecidable due to Kennedy and Pierce and the subsequent proof by Grigore [2017] implying that $\langle C, X \rangle$ is undecidable, completely characterize the computing power of all eight points of the lattice. This characterization is not in terms of the Chomsky hierarchy of classes of formal languages, but rather in terms of a similar hierarchy of classes of forests, also called formal tree languages (see, e.g., [Comon et al. 2007]), as depicted by Fig. 2. The classes appearing in the figure are defined in Section 2.

Fig. 2. A Chomsky-like hierarchy of formal forest classes (defined in Section 2)

As shown in Fig. 2, the hierarchy of tree languages is a non-linear lattice, spanned by two orthogonal Boolean features abbreviated by letters in the acronyms in the figure as follows:

- *Determinism* (D < N) distinguishing between *Deterministic* forests and their super-set, *Non-deterministic* forests;
- *Method of specification* (R < C) distinguishing between *Regular* forests and their super-set, *Context-free* forests.

In Fig. 2, we see that forest class DRF is the smallest class in the diagram. Moving up the lattice, the two next classes, NRF and DCF, cannot be compared. Both are contained in NCF, which is contained in REF, the class of all recursively enumerable forests. We also make use of class $NCF_{GNF} \subset NCF$,

which includes the non-deterministic context-free forests that can be specified by a tree grammar in Greibach normal form (GNF). Observe that class NCF_{GNF} also contains DCF and NRF.

The theoretical contribution of our work is a characterization of the computing power of the six decidable points in the lattice of Fig. 1 in terms of the hierarchy of Fig. 2. These results are obtained by examining the *subtyping machine* behind subtyping queries—a kind of automata with finite control (dictated by the nominal inheritance table) whose auxiliary storage is not a tape, stack, or tree, but rather the contents of a specific subtyping query that changes along the run of this automaton. These machines stand behind Kennedy and Pierce’s study and Grigore’s emulation.

1.2 Fluent APIs and Treetop

Fluent API (application programming interface) [Fowler 2005] is a trendy practice for the creation of software interfaces [Nakamaru et al. 2020]. In a statically-typed, object-oriented setting, a fluent API method returns a type that receives other API methods. In this way, fluent API methods are invoked in a chain of consecutive calls:

$$\sigma_1() \cdot \sigma_2() \cdot \dots \cdot \sigma_n() \quad (1)$$

The main advantage of fluent API is its ability to enforce a protocol at compile time. If an invocation of method σ_{i+1} after a call to σ_i breaks the protocol, then the type returned from method σ_i must not declare σ_{i+1} , which would cause the expression to fail type checking.

A practical application of fluent API is embedding domain specific languages (DSLs) in a host programming language [Gil and Levy 2016; Gil and Roth 2019; Nakamaru and Chiba 2020; Nakamaru et al. 2017; Xu 2010; Yamazaki et al. 2019]. The syntax and semantics of a DSL are optimally designed for its domain; for instance, SQL [Chamberlin and Boyce 1974] is a DSL suited for composing database queries. If the protocol of the fluent API is a grammar G describing the DSL, then the chain of method calls of Eq. (1) type checks if and only if $w = \sigma_1\sigma_2 \dots \sigma_n \in L(G)$, i.e., the expression encodes a legal DSL program. For example, with an SQL fluent API implemented in Scala [Odersky et al. 2004] we can embed database queries as Scala expressions: `select("id").from(grades).where(_.grade>90)`. The expression `select("name").where(_.length==4)`, however, does not type check, as it represents an illegal SQL query, **SELECT** name **WHERE** length=4.

The academic study of fluent APIs seeks new ways to encode increasingly complex DSL classes. Results in the field are given in the form of *generators*: software artifacts that convert a formal grammar into a fluent API encoding its language. Fling [Gil and Roth 2019] and TypeLevelLR [Yamazaki et al. 2019] are two prominent fluent API generators that support all deterministic context-free languages (DCFLs). Grigore [2017] implemented a generator that accepts context-free grammars describing general context-free DSLs—a proper super-set of DCFLs. Specifically, his implementation coded a CYK parser in an intermediate simple, yet high level language, and then translated the code in this language into a Turing machine, being emulated in a subtyping query. Unfortunately, the resulting APIs are unusable in practice, as they lead to extremely long compilation times and crash the compiler.

Examining Fig. 1 we see that the *contravariant* regime for combining genericity with subtyping is essential for the undecidability result of Kennedy and Pierce and its use by Grigore to encode Turing machines with Java [Arnold et al. 2006] generics. Contravariance is also essential to Grigore’s implementation of a fluent API generator for *all* context-free languages. Grigore’s result is noteworthy since the best previous fluent API generators [Gil and Levy 2016; Gil and Roth 2019; Yamazaki et al. 2019], relying on invariance, were only able to encode the deterministic subset of context-free languages (languages with an LR grammar). As shown in Fig. 1, our results prove that

the covariant regime, exemplified by the simple phrase “*a bag of apples is bag of fruit*”, is weaker than its contravariant counterpart.

Establishing that covariance is weaker than contravariance and noticing that it is stronger than invariance, it is natural to ask whether the entire class of context-free languages can be encoded in a type system with nominal subtyping combined with only covariant genericity. Conversely, the question is motivated by the argument that there is no theoretical difficulty in recognizing context-free languages in a Turing-complete type system, but it might be a challenge to recognize these in a decidable type system.

Unfortunately, the code produced by Grigore’s fluent API generator collapses under its own weight (measured in gigabytes) and crashes rugged modern compilers. Another natural question to ask is: Is the (computationally) weaker artillery of covariance also lighter in weight (in terms of code size)?

Treetop is a product of our theoretical study that answers both questions affirmatively. It is a proof of concept (POC) generator of “covariant generic subtyping” machines that are implemented as C[#] code. Given G , a CFG defining a formal language ℓ (such as the language of palindromes in the example of [List. 1](#)), Treetop generates a library in C[#] that defines a fluent API of ℓ . The library uses expansive inheritance (the X feature in the lattice of [Fig. 1](#)) and multiple instantiation inheritance (M in this lattice), but not contravariance (C). Thus, Treetop not only supports general context-free grammars, but it also does this in a decidable type system fragment. We demonstrate in an experiment non-trivial subtyping machines generated by Treetop that compile relatively fast, yet we also identify at least one example which induces impractically long compilation times using the CSC C[#] compiler.

Outline. [Sections 2](#) and [3](#) elaborate on Kennedy and Pierce’s work on subtyping, to refresh readers’ understanding of concepts such as formal forests and tree grammars, and set up our notations. Subtyping machines are then defined in [Section 4](#) as computational devices that model nominal subtyping. [Section 5](#) derives the results of [Fig. 1](#). Treetop, our API generator for context-free protocols, is the focus of [Section 6](#). [Section 7](#) discusses the results of our study.

Note on bold notation. The bold version of a symbol denotes a (possibly empty) sequence of instances of the non-bold version of this symbol, e.g., $\mathbf{x}$ stands for $x_1, x_2, \dots, x_k$ for some $k \geq 0$. This notation is abused when the intent is clear from the context, e.g.,

$$\tau[\mathbf{x} \leftarrow \mathbf{t}]$$

abbreviates

$$\tau_1[x_1 \leftarrow t_1, \dots, x_k \leftarrow t_k], \dots, \tau_n[x_1 \leftarrow t_1, \dots, x_k \leftarrow t_k],$$

where τ , $\mathbf{x}$, and $\mathbf{t}$ are the three sequences $\tau = \tau_1 \dots, \tau_n$, $\mathbf{t} = t_1 \dots, t_n$, and $\mathbf{x} = x_1, \dots, x_n$.

2 PRELIMINARIES: FORMAL FORESTS AND THEIR GRAMMARS

We assume that the reader is familiar with the classical theory of automata and formal languages (e.g., [[Hopcroft et al. 2007](#)]). Here we provide a quick review of the standard generalization of the theory to *forests*, also called (formal) *tree languages*. (For a more detailed review, see standard references such as [[Comon et al. 2007](#); [Gécseg and Steinby 1997](#); [Osterholzer 2018](#)].)

We use parenthetical notation for terms (trees), e.g.,

$$t = \sigma_1(\sigma_2(\sigma_3()), \sigma_4()) \tag{2}$$

Rank $n \in \mathbb{N}$ of symbol σ , denoted $\text{rank}(\sigma)$, is the number of its children. If $n = 0$, then σ is called a leaf, and if $n = 1$, then σ is monadic. In both cases, σ ’s parentheses may be omitted, e.g., t in [Eq. \(2\)](#) becomes $\sigma_1(\sigma_2\sigma_3, \sigma_4)$. The height of a tree t is denoted $\text{height}(t)$, e.g., for t in [Eq. \(2\)](#), $\text{height}(t) = 3$.

Fig. 3. The derivation of the term in Eq. (4), encoding the list $\langle 2, 1, 0 \rangle$, from variable $List$ according to the productions of Eq. (3)

The set of trees over ranked alphabet Σ (i.e., an alphabet of symbols with ranks) is denoted Σ^Δ , generalizing the Kleene closure K^* for forests.

Consider the following definitions, adapted from Comon et al. [2007, §2.1.1]:

$$\begin{aligned} \text{(a) } Nat &::= z & \text{(b) } Nat &::= s(Nat) \\ \text{(c) } List &::= nil & \text{(d) } List &::= cons(Nat, List) \end{aligned} \quad (3)$$

These are tree grammar productions deriving Peano numbers (productions (a,b)) and cons-lists of those integers (productions (c,d)). An integer $n \in \mathbb{N}$ is encoded by the monadic term $s^n z$. Instead of strings, tree grammar productions derive terms: The left-hand side of a production is a variable, and the right-hand side is a term with which the variable can be substituted. As in the string case, terminal (ground) terms are derived from variables by repeatedly substituting the latter according to the given productions. For example, we encode the list $\langle 2, 1, 0 \rangle$ using the term

$$cons(ssz, cons(sz, cons(z, nil))) \quad (4)$$

This term can be derived from variable $List$ from Eq. (3), as illustrated in Fig. 3. Fig. 3 uses arrows to depict variable substitution; an arrow label references the production from Eq. (3) used to apply the derivation.

Formally, a forest grammar $G = \langle \Sigma, V, v_0, R \rangle$ employs a finite ranked alphabet Σ , a finite ranked set of variables V , a designated initial leaf $v_0 \in V$, and a finite set of productions R . To generate a terminal tree, grammar G repeatedly applies R productions to the initial leaf v_0 , as explained below. Intermediate trees, members of $(\Sigma \cup V)^\Delta$, are called tree forms, similar to sentential forms of string grammars.

Let X be a finite set of *parameter* leaf symbols, disjoint to sets Σ and V . Terms over $\Sigma \cup V$, denoted t , are ground, while terms over $\Sigma \cup V \cup X$, denoted τ , are unground. Unground terms are used as tree patterns, where the X parameters can be substituted with terms (usually ground). The application of substitution $s = \{x \leftarrow t\} = \{x_1 \leftarrow t_1, \dots, x_k \leftarrow t_k\}$ to unground term τ , denoted $\tau[s]$, yields a term where every occurrence of parameter x_i is substituted with term t_i . (A substitution's curly brackets may be omitted if written inside square brackets.) A term τ matches term τ' if and only if there exists a substitution s such that $\tau'[s] = \tau$, denoted $\tau \sqsubseteq_s \tau'$. Let $\text{params}(\tau) \subseteq X$ denote the set of parameters appearing in term τ .

A production $\rho = \tau ::= \tau' \in R$ describes a tree transformation. If tree t matches τ with substitution s , $t \sqsubseteq_s \tau$, the application of ρ to t , denoted $t[\rho]$, yields τ' substituted by s ,

$$t \sqsubseteq_s \tau \Rightarrow t[\tau ::= \tau'] = \tau'[s].$$

We require that every parameter $x \in X$ on the right-hand side of production $\rho \in R$ be found on the left-hand side, so that all the applications of ρ to ground trees are ground.

Note the distinction between variables and parameters. Variables are found in tree forms and constitute the nodes remaining to be derived. Parameters, on the other hand, appear only in productions and are used as templates in tree patterns.

We say that tree form t derives tree form t' , denoted $t \rightarrow t'$ (or $t \rightarrow_G t'$), if t can be transformed into t' by a single application of an R production:

$$t \rightarrow t' \Leftrightarrow \exists \rho \in R, t[\rho] = t'.$$

The tree language (forest) of grammar G , $L(G)$, is the set of terminal trees that can be derived from the initial variable v_0 in any number of steps:

$$L(G) = \{t \in \Sigma^\Delta \mid v_0 \rightarrow_G^* t\}.$$

Each class (family) of tree grammars is defined by a set of restrictions imposed on the shape of productions. In regular tree grammars that define regular forests (NRF), productions are of the form $v ::= \sigma(v)$, where $\sigma \in \Sigma$ and $v, v \in V$; a more relaxed definition simply restricts variables to be leaves. Observe that the productions of Eq. (3) belong to a regular grammar.

Productions of context-free tree grammars (CFTG), yielding context-free forests (NCF), are of the form $v(x) ::= \tau$ (the parameters in x are unique). For example, the following productions belong to a CFTG:

$$(a) v_0 ::= v_1(\text{leaf}) \quad (b) v_1(x) ::= v_1(\text{node}(x, x)) \quad (c) v_1(x) ::= x \quad (5)$$

These productions derive complete binary trees over $\{\text{leaf}, \text{node}\}$. The binary trees are derived recursively. The child of variable v_1 , captured by parameter x , describes a tree of height $n \geq 1$; production (b) is then used to replace x with $\text{node}(x, x)$, the complete binary tree of height $n + 1$. The recursion terminates by production (c), which simply returns the constructed tree.

In CFTG in Greibach normal form (GNF), the roots of the right-hand side of productions are terminal, $v(x) ::= \sigma(\tau)$ [Guessarian 1983]. Observe that none of the productions in Eq. (5) are in GNF: The roots of the right-hand sides are variables in productions (a,b) and a parameter in production (c), both forbidden in GNF. In contrast, the following productions do belong to a CFTG in GNF:

$$(a) v_0x ::= av_0ax \quad (b) v_0x ::= aax \quad (c) v_0x ::= ax \\ (d) v_0x ::= bv_0bx \quad (e) v_0x ::= bbx \quad (f) v_0x ::= bx \quad (6)$$

Variable v_0 derives palindromes over $\{a, b\}$. For instance, the palindrome *abbabba* is derived from the tree form v_0E (E is a dummy leaf) as follows:

$$v_0E \xrightarrow{(a)} av_0aE \xrightarrow{(d)} abv_0baE \xrightarrow{(d)} abbv_0bbaE \xrightarrow{(c)} abbabbaE$$

(The production from Eq. (6) used in each derivation is written above the arrow.) The productions in Eq. (6) are context-free, as v_0 employs a parameter, and they are in GNF, as the root of the right-hand side of each production is a terminal (either a or b).

Notice that all regular grammars are context-free and in GNF.

In deterministic context-free tree grammars in GNF, including deterministic regular tree grammars, if both $\rho_1 = v ::= \sigma(\tau)$ and $\rho_2 = v ::= \sigma(\tau')$ are productions in R , then $\tau = \tau'$. Deterministic regular grammars define deterministic regular forests (DRF), a proper subset of regular forests, also called path-closed forests [Comon et al. 2007]. Likewise, deterministic context-free tree grammars define deterministic context-free forests (DCF). In contrast to context-free string grammars, not every context-free tree grammar has a GNF—but every deterministic tree grammar does [Guessarian 1983]. The class of context-free forests that can be described by a grammar in GNF is denoted as NCF_{GNF} .

In this paper, we use an extended version of context-free tree grammars where the initial variable v_0 is replaced by an initial tree.

Definition 2.1 (ECFTG). Extended context-free tree grammars (ECFTGs) generalize standard CFTGs by allowing initialization by any tree t_0 . Otherwise, the semantics are the same.

A standard CFTG is extended by definition, with $t_0 = v_0$. We show that extended grammars are as expressive as standard grammars:

LEMMA 2.2. *For every ECFTG (in GNF), there exists a CFTG (in GNF) that defines the same forest.*

PROOF. Let $G = \langle \Sigma, V, t_0, R \rangle$ be an ECFTG. If the initial tree t_0 is a variable, $t_0 = v \in V$ and we are done. Otherwise, let CFTG $G' = \langle \Sigma, V', v_0, R' \rangle$, introducing the initial variable $v_0 \notin V$, $V' = V \cup \{v_0\}$, and employing the original productions, $R \subseteq R'$. If the initial tree is terminal, $t_0 = \sigma_0 \in \Sigma$, simply add the rule $v_0 ::= \sigma_0$ to R' . Else t_0 is not a leaf: Let $t_0 = \xi_0(t_0)$, where the root $\xi_0 \in \Sigma \cup V$ is of rank one or above. For every production $\xi_0(x) ::= \tau \in R$ if ξ_0 is a variable, or production $\xi_0(x) ::= \tau$, $\tau = \xi_0(x)$ if ξ_0 is terminal, add production $v_0 ::= \tau[x \leftarrow t_0]$ to R' .

Derivations of the initial variable v_0 expand it into the initial tree t_0 and apply an original R derivation in a single step. After v_0 is derived, the resulting tree form is identical to the one derived from t_0 ,

$$v_0[v_0 ::= \tau[x \leftarrow t_0]] = t_0[\xi_0(x) ::= \tau],$$

thus they derive the same forest.

This construction preserves GNF. Newly introduced rules derive tree $\tau[x \leftarrow t_0]$, whose root is terminal given that either (i) the root of t_0 is terminal, or otherwise (ii) τ is the right-hand side of a rule in R , starting with a terminal given that the original grammar is in GNF. $\square$

3 PRELIMINARIES: NOMINAL SUBTYPING WITH VARIANCE

This work focuses on Kennedy and Pierce's [2007] abstract type system of *nominal subtyping with variance*. This type system supports nominal subtyping with declaration-site variance as found in, e.g., C[#] and Scala. Fig. 4 describes the abstract syntax (sub-figures (a,b)) and semantics (sub-figures (c,d,e)) of the type system, both adopted from Kennedy and Pierce [2007]. Note that we reuse notations introduced in Section 2 in the context of formal forests, foreshadowing the connection between the two formalisms.

The full syntax of the type system is given in Fig. 4 (a) and a nomenclature is available in Fig. 4 (b). A program consists of a class table Δ and a subtyping query q . A class table over a finite set of class names Γ is a finite set of inheritance rules δ . An inheritance rule $\delta = \gamma(\diamond x) : \tau$ declares a parameterized class (polytype) named γ and its inheritance from (super-) types τ ; to highlight that a sequence of supertypes τ is empty, we write $\emptyset$ after the colon. We require supertypes to reference only class parameters, $\text{params}(\tau) \subseteq x$. Although Kennedy and Pierce included mixin inheritance in their type system, allowing a class to inherit one of its parameters, we chose to exclude it; accordingly, $\tau \cap x = \emptyset$ (nevertheless, the implications of mixin inheritance are discussed in Section 7). Variance is denoted by a diamond $\diamond$, which can be covariant, $\diamond = +$, contravariant, $\diamond = -$, or invariant, $\diamond = \circ$. The variance of the i^{th} parameter of class γ is denoted $\diamond_i^\gamma$, and the vector describing all γ parameter variances is denoted $\diamond^\gamma$. A subtyping query $t_\perp <: t_\top$ (also written $t_\perp <:_\Delta t_\top$), for subtype $t_\perp$ and supertype $t_\top$, both ground, type checks if and only if $t_\perp$ is a subtype of $t_\top$ (considering class table Δ).

$P ::= \Delta q$ $\Delta ::= \delta_1 \dots \delta_n \quad n \geq 0$ $\delta ::= \gamma(\diamond x) : \tau \quad \begin{cases} \text{params}(\tau) \subseteq x \\ \tau \cap x = \emptyset \end{cases}$ $x ::= \diamond x_1, \dots, \diamond x_k \quad k \geq 0$ $\diamond ::= \circ \mid + \mid -$ $\tau ::= \tau_1, \dots, \tau_k \quad k \geq 0$ $\tau ::= \gamma(\tau) \mid x$ $q ::= t_{\perp} <: t_{\top}$ $t ::= \gamma(t)$ $t ::= t_1, \dots, t_k \quad k \geq 0$	P program δ class declaration Δ class table; a set of class declarations γ class name drawn from finite set Γ τ unground type τ sequence of unground types x type parameter drawn from finite set X disjoint from Γ $\diamond x$ type parameter with variance $\diamond x$ sequence of type parameters with variances $\circ$ invariance $+$ covariance $-$ contravariance q subtyping query t ground type t sequence of ground types
(a) Syntax	(b) Nomenclature
$\frac{t \sqsubseteq_s \gamma(x) \quad \gamma(x) : \tau}{t <:: \tau[s]}$	$\begin{aligned} &(\text{VAR}) \frac{\forall i, t_i <:_{\circ_i} t'_i}{\gamma(t) <: \gamma(t')} \quad \frac{t <: t'}{t <:_{+} t'} \quad \frac{t <: t'}{t <:_{\circ} t} \quad \frac{t >: t'}{t <:_{-} t'} \\ &(\text{SUPER}) \frac{t <:: t' \quad t' <: t''}{t <: t''} \end{aligned}$
(c) Inheritance	(d) Subtyping
$\begin{aligned} &(\text{ACYCLIC INHERITANCE}) \quad \gamma(\tau) <::^+ \gamma'(\tau') \Rightarrow \gamma \neq \gamma' \\ &(\text{VALID VARIANCE}) \quad \neg \diamond = \begin{cases} - & \diamond = + \\ \circ & \diamond = \circ \\ + & \diamond = - \end{cases} \quad \frac{\diamond_i \in \{\circ, +\}}{\diamond x \vdash x_i} \quad \frac{\forall i, \begin{cases} \diamond_i^{\gamma} \in \{\circ, +\} \Rightarrow \diamond x \vdash \tau_i \\ \diamond_i^{\gamma} \in \{\circ, -\} \Rightarrow \neg \diamond x \vdash \tau_i \end{cases}}{\diamond x \vdash \gamma(t)} \quad \frac{\diamond x \vdash \tau}{\gamma(\diamond x) : \tau \checkmark} \end{aligned}$	
(e) Well-Formedness	

Fig. 4. Kennedy and Pierce's type system (without mixin inheritance)

Example. For example, we can write the C[#] program of List. 1 using abstract syntax as follows (recall that we use an abbreviated notation of monadic terms):

$$\Delta \left\{ \begin{array}{l} a(+x) : \emptyset \\ b(+x) : \emptyset \\ E : \emptyset \\ v_0(\circ x) : av_0ax, aax, ax, bv_0bx, bbx, bx \end{array} \right. \quad (7)$$

$$q \left\{ \begin{array}{l} v_0E <: abbabbaE \end{array} \right.$$

Class table Δ has four classes, a , b , E , and v_0 . In C[#], type parameters are invariant by default, covariance is denoted by the keyword **out**, and contravariance by the keyword **in**. Therefore, the type parameters of a and b are covariant ($\diamond_1^a = \diamond_1^b = +$), and the type parameter of v_0 is invariant ($\diamond_1^{v_0} = \circ$). Query q is the subtyping query invoked by the variable assignment in line 3 of List. 1.

According to the type inheritance relation $t <:: t'$, defined in Fig. 4 (c), type t inherits type t' if t' matches one of the supertypes of t 's root. Unlike Kennedy and Pierce, we distinguish between the type inheritance relation and class inheritance, $\gamma : \gamma'$, found in class declarations. We also use the transitive closures of these:

$$\frac{\gamma(x) : \gamma'(\tau) \quad \gamma'(x') :^* \tau}{\gamma(x) :^* \gamma[x' \leftarrow \tau]} \quad (8)$$

The same rules apply for transitive type inheritance. The asterisk is replaced with a plus $+$ in the non-reflexive variants $:^+$ and $<::^+$.

The rules of subtyping are detailed in Fig. 4 (d). The SUPER rule establishes that a class is a subtype of its supertypes. The VAR rule handles the type parameters, depending on their variance: Statement $\gamma(t) <: \gamma(t')$ implies $t <: t'$ if γ 's type parameter is covariant, $t = t'$ if invariant, and $t' <: t$ if contravariant. If γ accepts multiple type parameters, then each may have its own variance. We also define transitive subtyping:

$$\frac{}{t <::^* t} \quad \frac{t <: t' \quad t' <::^* t''}{t <::^* t''} \quad (9)$$

For example, we can use the SUPER and VAR rules to prove the subtyping query of Eq. (7):

$$\begin{aligned} & v_0 E <: \text{abbabba} E \\ \Leftarrow & av_0 a E <: \text{abbabba} E \quad (\text{SUPER}, v_0(\circ x) : av_0 a x) \\ \Leftarrow & v_0 a E <: \text{bbabba} E \quad (\text{VAR}, \diamond_1^a = +) \\ \Leftarrow & bv_0 ba E <: \text{bbabba} E \quad (\text{SUPER}, v_0(\circ x) : bv_0 bx) \\ \Leftarrow & v_0 ba E <: \text{babba} E \quad (\text{VAR}, \diamond_1^b = +) \\ \Leftarrow & v_0 bba E <: \text{babba} E \quad (\text{SUPER}, v_0(\circ x) : bv_0 bx) \\ \Leftarrow & v_0 bba E <: \text{abba} E \quad (\text{VAR}, \diamond_1^b = +) \\ \Leftarrow & \text{abba} E <: \text{abba} E \quad \checkmark \quad (\text{SUPER}, v_0(\circ x) : ax) \end{aligned}$$

Note that this subtyping query is recursive because the type parameters of classes a and b are covariant. If the parameters were contravariant, then the query would reverse ($>$) after applying the VAR rule.

A class table is well-formed if it complies with the rules in Fig. 4 (e). First, inheritance must be acyclic, i.e., a class may not (transitively) inherit itself. Thus, the class inheritance closure $:^*$ is finite. Kennedy and Pierce also explain that transitive subtyping and semantic soundness rely on certain restrictions to class inheritance, $\gamma(\diamond x) : \tau$.

To demonstrate the problem, consider the following *invalid* class declarations:

$$\begin{aligned} a(-x) &: \emptyset \\ b(+y) &: a(y) \end{aligned}$$

Type parameter y of class b is covariant, but it also substitutes for a 's contravariant parameter x . This means that y can be used both covariantly and contravariantly. For instance, if $c <:: d$, then

$$b(c) <: b(d)$$

but also

$$b(d) <: a(c)$$

One way to solve this issue is by changing b 's supertype to $a(a(y))$:

$$b(+y) : a(a(y))$$

Once done, y behaves covariantly, as expected:

$$\begin{aligned} & bc <: aad \\ \Leftarrow & aac <: aad \quad (\text{SUPER}, b(+y) : aay) \\ \Leftarrow & ac >: ad \quad (\text{VAR}, \diamond_1^a = -) \\ \Leftarrow & c <: d \quad \checkmark \quad (\text{VAR}, \diamond_1^a = -) \end{aligned}$$

Intuitively, one may think of a covariant type parameter as defining a “positive position”, and a contravariant one a “negative position”. A covariant type parameter $x \in \mathbf{x}$ can be placed in supertype τ only in positive positions, and if it is contravariant, only in negative positions. A

sub-tree τ' of τ in a negative position reverses its polarity: Positive positions become negative, and negative positions become positive. The initial position of τ is positive. In this context, an invariant position is both positive and negative, and so are invariant type parameters. Thus, invariant type parameters can appear in any position, while a type in an invariant position contains only invariant type parameters, because its positions are both positive and negative.

For example, consider the following class table:

$$\begin{array}{ll} a(+x) : \emptyset & d(+x_1) : ada x_1 \\ b(-x) : \emptyset & e(-x_2) : bbab x_2 \\ c(\circ x) : \emptyset & f(-x_3, \circ x_4, +x_5) : af(x_4, ax_4, x_4) \end{array} \quad (10)$$

This class table, ranging over classes a, b, c, d, e , and f , is well-formed. The table has no inheritance cycles, e.g., although type d appears in the type it inherits, we consider only the roots of supertypes in checks of circularity. Restrictions on the use of type parameters are also respected. Covariant type parameter x_1 in type $ada x_1$ is in a positive position, as both a and d are covariant. Contravariant type parameter x_2 in type $bbab x_2$ is in a negative position: Type b reverses the polarity an odd number of times and type a does not change it, so the final position is negative. Invariant type parameter x_4 can be placed in all the positions of f in type $af(x_4, ax_4, x_4)$, and it is the only parameter suitable for the invariant (second) position. The C[#] programming language enforces the rules of Fig. 4 (e); in List. 2, we encode the class table of Eq. (10) as C[#] interface declarations that compile as expected.

```

1 interface a<out x> {} interface d<out x1> : a<d<a<x1>>> {}
2 interface b<in x> {} interface e<in x2> : b<b<a<b<x2>>>> {}
3 interface c<x> {} interface f<in x3, x4, out x5> : a<f<x4, a<x4>, x4>> {}

```

Listing 2. The class table of Eq. (10) encoded in C[#]

Let the type system featured in Fig. 4 be denoted as $\mathcal{S}_\top$. Kennedy and Pierce [2007] proved that $\mathcal{S}_\top$ is undecidable. Nevertheless, they introduced three restrictions on $\mathcal{S}_\top$, each defining a decidable type system fragment of $\mathcal{S}_\top$. Each restriction removes a certain feature of the type system: contravariance (C), expansive inheritance (X), or multiple instantiation inheritance (M). We denote a type system by the set of features used by its programs, e.g., $\langle C, X, M \rangle$ is an alias of $\mathcal{S}_\top$, and its subset containing only programs without contravariance is denoted by $\langle X, M \rangle$. Each element of the power set of $\{C, X, M\}$ defines a type system. We organize the eight type systems in a lattice, depicted in Fig. 1, ordered with respect to inclusion (of features and programs). The bottom of the lattice is $\mathcal{S}_\perp$, bare of any special features. All these type systems employ polyadic parametric polymorphism, i.e., where a class can have multiple type parameters.

The first restriction of Kennedy and Pierce forbids contravariant type parameters, $-x$, yielding type system $\langle X, M \rangle$. The second restriction disallows multiple instantiation inheritance, making inheritance deterministic: If $\gamma(x) :^+ \gamma'(\tau)$ and $\gamma(x) :^+ \gamma'(\tau')$, then necessarily $\tau = \tau'$. For instance, the class table of Eq. (7) does not conform to this restriction, as class v_0 extends both $a(x)$ and $a(ax)$. Kennedy and Pierce suspected that single instantiation inheritance, featured in $\langle C, X \rangle$, is not enough to achieve decidability. Grigore [2017] verified the suspicion by simulating Turing machines with subtyping queries in Java, which employs deterministic inheritance. To make single inheritance decidable, it had to be further restricted; the resulting type system will not be discussed in this paper.

The third restriction, due to Viroli [2000], limits the growth of types during subtyping. Consider, for example, the class table depicted in Eq. (10). Given type dt , for some type argument t , on either side of a subtyping query, it can change to $adat$ by applying the SUPER rule and then to dat by

VAR, yielding the type dt' for $t' = at$. Type dt can, therefore, expand indefinitely during subtyping. Kennedy and Pierce named this kind of inheritance *recursively expansive* and showed its connection to infinite subtyping proofs.

Originally, Viroli illustrated this growth of types by a graph, depicting the dependencies between the different type variables employed by the class table. In this paper, however, we use a more approachable definition, following an equivalence result of Kennedy and Pierce. Consider the following definition:

Definition 3.1 (Inheritance and decomposition closure). Set T of types is closed under inheritance and decomposition if for every type $\gamma(t)$ in T , its sub-trees are in T , $t \subset T$, and so are its supertypes, $\gamma(t) <:: t \Rightarrow t \in T$. The minimal super-set of set S closed under inheritance and decomposition is denoted $\text{cl}(S)$. (The curly braces of set S may be omitted inside $\text{cl}(\cdot)$, e.g., $\text{cl}(t_1, t_2)$.)

With the class table of Eq. (10), for instance, $\text{cl}(dt) \supset \{da^n t \mid n \geq 0\}$, as explained above. Kennedy and Pierce [2007] noted that the set of types involved in a subtyping query $t_\perp <: t_\top$ is contained in their inheritance and decomposition closure, $\text{cl}(t_\perp, t_\top)$, following the SUPER (inheritance) and VAR (decomposition) rules. If a closure $\text{cl}(t_\perp, t_\top)$ is finite, then the subtyping query $t_\perp <: t_\top$ is decidable, as we can trace infinite subtyping cycles. Therefore, if the closure $\text{cl}(t, t')$ of any two types t and t' is finite, subtyping is decidable, and the class table is called finitary. Kennedy and Pierce proved that subtyping with non-expansive inheritance is decidable, by showing that a class table is non-expansive if and only if it is finitary. Thus, we can use Definition 3.1 to define non-expansive inheritance directly:

Definition 3.2 (Non-expansive class table). Class table Δ is non-expansive if and only if the inheritance and decomposition closure of any set of types T is finite, $|\text{cl}(T)| < \infty$.

We denote the type system fragment where all the class tables are non-expansive as $\langle C, M \rangle$.

If a class table Δ conforms to type system $\mathcal{S}$, we write $\mathcal{S} \vdash \Delta$. For example, class table Δ of Eq. (10) conforms to $\mathcal{S}_\top$, written $\mathcal{S}_\top \vdash \Delta$, as it is well-formed, but not to $\langle C, M \rangle$, written $\langle C, M \rangle \not\vdash \Delta$, as its inheritance is expansively recursive.

4 SUBTYPING MACHINES AS FOREST RECOGNIZERS

A subtyping machine is a set of type system definitions that utilize subtyping and its rules to perform a computation. Grigore [2017], who coined the term, used subtyping machines to simulate Turing machines in Java, thus proving its type system is undecidable.

This section formalizes subtyping machines as general computation devices, following a standard approach in fluent API research [Gil and Roth 2020, 2019; Yamazaki et al. 2019]. Subtyping machines recognize formal forests, i.e., sets of trees or terms. In Section 5, we use this formalism to classify the computational power of decidable subtyping in its various forms.

This section's running example is the program shown in List. 1. We focus on the type declaration in the first two lines of the code:

```
1 interface a<out x> {} interface b<out x> {} interface E {}
2 class v0<x> : a<v0<a<x>>>, a<a<x>>>, a<x>, b<v0<b<x>>>, b<b<x>>>, b<x> {}
```

Let us name this program Pali. In Section 1, we introduced Pali as an example of a subtyping machine that recognizes palindromes, but did not explain what “recognize” means. Our goal is, therefore, to formalize an interpretation of Pali as a formal forest of palindromes.

In our type system, presented in Fig. 4, a program $P = \Delta q$ comprises a set of declarations Δ constituting a class table and a query (or assertion) $q = t_\perp <: t_\top$. The query succeeds (compiles) if and only if type $t_\perp$ is a subtype of type $t_\top$ considering class table Δ and the subtyping semantics

of the type system, $t_{\perp} <_{\Delta} t_{\top}$. Therefore, a given class table defines the set of queries that compile against it; we call this set the *extent* of the class table.

The extent of Pali's class table Δ (written using abstract syntax in Eq. (7)) contains, for instance, the queries

$$v_0E <: aE, \quad v_0E <: abaE, \quad \text{and} \quad v_0E <: abbabbaE,$$

because they type check against Δ . In general, the extent of Δ contains all the queries of the form

$$v_0E <: pE$$

where p is a palindrome over $\{a, b\}$ (we prove this in Section 5.2). Δ 's extent, however, also contains queries that do not include palindromes at all, such as

$$E <: E, \quad v_0aE <: baE, \quad \text{and} \quad av_0E <: abE.$$

To correlate Pali's extent with a forest of palindromes, we first need to restrict the contents of its queries.

Recall that formal grammars distinguish between terminal and non-terminal (variable) symbols, drawn from sets Σ and V , respectively. This dichotomy allows grammars to use variables as auxiliary symbols, which cannot appear as part of derived words. To make our definition of subtyping machines effective, we introduce an analogous separation in Γ , the set of class names in Δ . Let sub-alphabet $\Sigma_{\perp} \subseteq \Gamma$ and super-alphabet $\Sigma_{\top} \subseteq \Gamma$ be the domains of subtypes $t_{\perp}$ and supertypes $t_{\top}$ in subtyping queries, i.e., types $t_{\perp}$ and $t_{\top}$ consist of type names drawn from designated Γ subsets $\Sigma_{\perp}$ and $\Sigma_{\top}$. We amend Kennedy and Pierce's type system, defined in Fig. 4, to include the sub- and super-alphabets:

$$q ::= t_{\perp} <: t_{\top} \quad t_{\perp} \in \Sigma_{\perp}^{\Delta}, t_{\top} \in \Sigma_{\top}^{\Delta} \quad (\Sigma_{\perp}^{\Delta}, \Sigma_{\top}^{\Delta} \subseteq \Gamma) \quad (11)$$

From this point on, we discuss programs that include the amendment presented in Eq. (11). We define the extent of class table Δ as the subset of $\Sigma_{\perp}^{\Delta} \times \Sigma_{\top}^{\Delta}$:

Definition 4.1 (The extent of a class table). Let Δ be a class table, as described in Fig. 4. The extent of Δ , denoted $L(\Delta)$, contains all the pairs of types $\langle t_{\perp}, t_{\top} \rangle$ such that the subtyping query $t_{\perp} <: t_{\top}$ type checks against Δ :

$$L(\Delta) = \{ \langle t_{\perp}, t_{\top} \rangle \mid t_{\perp} \in \Sigma_{\perp}^{\Delta}, t_{\top} \in \Sigma_{\top}^{\Delta}, t_{\perp} <_{\Delta} t_{\top} \} \quad \Sigma_{\perp}, \Sigma_{\top} \subseteq \Gamma$$

Definition 4.1 gives us more control over the expressiveness of programs. The extent of a class table, however, does not yet describe a formal forest. Although types are homomorphic to trees, an extent is a set of type pairs $\langle t_{\perp}, t_{\top} \rangle$, rather than a set of lone types. To solve this issue, we fix the subtype $t_{\perp}$ to some constant type that is no longer a variable part of the subtyping query. Then, a class table Δ defines a (tree) language of $t_{\perp}$ supertypes; this set is the projection of type $t_{\perp}$ onto extent $L(\Delta)$:

Definition 4.2 (The language of a class table). Let Δ be a class table and $t_{\perp}$ be a type, as described in Fig. 4. The language of Δ over $t_{\perp}$, denoted $L_{t_{\perp}}(\Delta)$, contains all the supertypes of $t_{\perp}$ against Δ ,

$$L_{t_{\perp}}(\Delta) = \{ t_{\top} \mid \langle t_{\perp}, t_{\top} \rangle \in L(\Delta) \} = \{ t_{\top} \mid t_{\top} \in \Sigma_{\top}^{\Delta}, t_{\perp} <_{\Delta} t_{\top} \} \quad \Sigma_{\top} \subseteq \Gamma$$

We can now use Definition 4.2 to describe Pali's language over subtype $t_{\perp} = v_0E$ as a forest of palindromes:

$$L(\text{Pali})_{v_0E} = \{ pE \mid p \text{ is a palindrome} \} \quad (\Sigma_{\top} = \{a, b, E\}) \quad (12)$$

In other words, Pali is a subtyping machine that recognizes palindromes as supertypes of type $t_{\perp} = v_0E$. We prove Eq. (12) in Section 5.2.

Considering [Definition 4.2](#), we see that class tables describe formal languages (forests), as grammars do, and their compilers serve as language (forest) recognizers, similarly to grammar parsers. Correspondingly, a type system is a family of programs that recognizes a class of languages:

Definition 4.3 (The class of languages of a type system). The computational class (expressiveness) of subtyping-based type system $\mathcal{S}$, denoted $\mathcal{L}(\mathcal{S})$, is the set of languages described by its programs:

$$\mathcal{L}(\mathcal{S}) = \{L_{t_\perp}(\Delta) \mid t_\perp \in \Sigma_\perp^\Delta, \mathcal{S} \vdash \Delta\} \quad \Sigma_\perp \subseteq \Gamma$$

(Recall that $\mathcal{S} \vdash \Delta$ means class table Δ is well-formed and conforms to the restrictions of type system $\mathcal{S}$.)

5 EXPRESSIVENESS OF DECIDABLE SUBTYPING

Recall that the type system of Kennedy and Pierce $\mathcal{S}_\top$ features contravariance (C), expansive inheritance (X), and multiple instantiation inheritance (M). Subset S of $\{C, X, M\}$ describes the type system fragment of $\mathcal{S}_\top$ endorsing only the features included in S . Subtyping in $\mathcal{S}_\top$ and $\langle C, X \rangle$ is known to be undecidable, but it is decidable in the other six fragments [[Grigore 2017](#); [Kennedy and Pierce 2007](#)]. In [Section 4](#), we described class tables as subtyping machines that recognize formal forests ([Definition 4.2](#)) and defined the expressiveness of a type system as the class of forests that its tables recognize ([Definition 4.3](#)). Now we use these formalizations to describe the computational power of Kennedy and Pierce's decidable type system fragments.

Given type system $\mathcal{S}$, we show $\mathcal{S}$ is at least as complex as class of languages $\mathcal{L}$ by proving that any language ℓ in $\mathcal{L}$ is described by some class table Δ and subtype $t_\perp$ of $\mathcal{S}$:

$$(\forall \ell \in \mathcal{L}, \exists \Delta, t_\perp, \mathcal{S} \vdash \Delta, \ell = L_{t_\perp}(\Delta)) \Rightarrow \mathcal{L} \subseteq \mathcal{L}(\mathcal{S})$$

To show that type system $\mathcal{S}$ is not more complex than class of languages $\mathcal{L}$, we prove that any class table Δ and subtype $t_\perp$ describe an $\mathcal{L}$ language:

$$(\forall \Delta, t_\perp, \mathcal{S} \vdash \Delta, \exists \ell \in \mathcal{L}, \ell = L_{t_\perp}(\Delta)) \Rightarrow \mathcal{L}(\mathcal{S}) \subseteq \mathcal{L}$$

If the expressiveness of $\mathcal{S}$ is bounded by $\mathcal{L}$ from both directions, then it is exactly $\mathcal{L}$:

$$\mathcal{L} \subseteq \mathcal{L}(\mathcal{S}) \subseteq \mathcal{L} \Rightarrow \mathcal{L}(\mathcal{S}) = \mathcal{L}$$

In all the following proofs, formal languages are specified by their corresponding grammars.

5.1 Non-Expansive Subtyping Is Regular

[Greenman et al. \[2014\]](#) raised a concern about non-expansive inheritance restricting some conventional class table use-cases. We go one step further by showing that the expressiveness of non-expansive class tables is regular, i.e., of the lowest tier.

THEOREM 5.1. *Any regular forest can be encoded in a class table using only multiple instantiation inheritance:*

$$\text{NRF} \subseteq \mathcal{L}(\langle M \rangle)$$

Intuition. We want to simulate the derivation of a regular tree grammar with subtyping. In the first step, we encode every terminal node σ of rank k by class σ that has k type parameters. Thus, terminal trees (derived from the grammar) and types (in subtyping queries) coincide, as both use the same term notation and base alphabet, e.g., $\sigma_1(\sigma_2, \sigma_3\sigma_4)$ is both a tree and a type. Next, every regular derivation rule $v ::= \sigma(v)$ is directly encoded by a corresponding inheritance rule $v : \sigma(v)$. With these definitions, the subtyping query $v <: t$ simulates the grammar derivation $v \rightarrow^* t$: Let $t = \sigma(t)$. Grammar variable v derives t if v derives σ , $v ::= \sigma(v)$ and then, recursively, variables v derive σ 's children, $v \rightarrow^* t$. Correspondingly, the subtyping query $v <: t$ type checks if class v

inherits type σ , $v : \sigma(v)$ and then, recursively, types v are subtypes of t , $v <: t$. The recursion of subtyping originates from the VAR rule, under the condition that the type parameters of σ are covariant.

Example. Consider the regular tree grammar in Eq. (3) in Section 2, describing lists of Peano numbers. This grammar can be encoded in a C# program with neither contravariance nor expansive recursive inheritance, as demonstrated in List. 3.

```

1  interface z {} // Terminal leaf z
2  interface s<out x> {} // Terminal node s(x)
3  interface nil {} // Terminal leaf nil
4  interface cons<out x1, out x2> {} // Terminal node cons(x1, x2)
5  interface Nat: // Variable Nat
6      z, // Production Nat ::= z
7      s<Nat> {} // Production Nat ::= s(Nat)
8  interface List: // Variable List
9      nil, // Production List ::= nil
10     cons<Nat, List> {} // Production List ::= cons(Nat, List)
11 // Subtyping query List <: cons(ssz, cons(sz, cons(z, nil)))
12 cons<s<s<z>>, cons<s<z>>, cons<z, nil>>> t=(List) null;
```

Listing 3. C# program encoding the grammar of Eq. (3)

The grammar terminals s , z , cons , and nil are encoded by interfaces with covariant type parameters, and the grammar variables Nat and List are encoded by interfaces whose supertypes are the possible productions of each variable. The last line of the listing contains a variable assignment that ensures type $v_0 = \text{List}$ is a subtype of the C# type encoding the list $\langle 2, 1, 0 \rangle$ (described in Eq. (4)). This assignment compiles, as expected.

PROOF OF THEOREM 5.1. A given regular forest is described by a regular tree grammar $G = \langle \Sigma, V, v_0, R \rangle$. We encode G as class table Δ , $\langle M \rangle \vdash \Delta$, such that the language of Δ (over $t_\perp = v_0$) is exactly the language described by G , $L(G)$.

Construction. The class table Δ employs type names drawn from $\Gamma = \Sigma \cup V$. The number of type parameters of each type $\sigma \in \Sigma$ is identical to its rank in G , and all type parameters are covariant (nodes $v \in V$ are leaves). For each regular production rule $v ::= \sigma(v) \in R$, set type $\sigma(v)$ as a super-type of type v , $v : \sigma(v)$. We set the subtype to $t_\perp = v_0$, and the super-alphabet to $\Sigma_\top = \Sigma$.

Correctness of the construction. Subtyping against Δ directly emulates derivations of grammar G , where the left-hand side is the current variable (tree form), starting with the initial variable v_0 , and the right-hand side is the derived tree. We show that a tree t is produced by variable v if and only if type t is a supertype of type v ,

$$v \rightarrow_G^* t \Leftrightarrow v <: t \quad (13)$$

by induction on the height of t , $n = \text{height}(t)$.

If $n = 0$, then tree t is a leaf, $t = \sigma$, $\sigma \in \Sigma$. Variable v derives σ only if production $v ::= \sigma$ is in R . In that case, type σ is a supertype of v , $v : \sigma \in \Delta$, so v is a subtype of t by inheritance:

$$v \rightarrow_G^* \sigma \Leftrightarrow v ::= \sigma \in R \Leftrightarrow v : \sigma \in \Delta \Leftrightarrow v <: \sigma$$

Suppose Eq. (13) holds for $n \geq 0$, and let t be a tree of height $n+1$, $t = \sigma(t)$, $\sigma \in \Sigma$, $t \in \Sigma^\Delta$, where the heights of sub-trees t are n at most. If variable v derives t , then there is a production $v ::= \sigma(v)$ in R , for which every child variable $v_i \in v$ derives the corresponding child tree t_i , $v \rightarrow_G^* t$. This production ensures type v inherits type $\sigma(v)$. As the heights of t are no higher than n , we can use the inductive assumption to deduce $v <: t$. Because the type parameters of σ are covariant,

we can show that v is a subtype of t by applying inheritance, decomposition, and the inductive assumption:

$$\begin{aligned} v <: \sigma(\mathbf{t}) &\Leftarrow v : \sigma(\mathbf{v}) \wedge \sigma(\mathbf{v}) <: \sigma(\mathbf{t}) \\ \sigma(\mathbf{v}) <: \sigma(\mathbf{t}) &\Leftarrow \mathbf{v} <: \mathbf{t} \wedge \diamond^\sigma = + \end{aligned}$$

Overall, we have now proved [Eq. \(13\)](#):

$$v \rightarrow_G^* t \Leftrightarrow v ::= \sigma(\mathbf{v}) \wedge \mathbf{v} \rightarrow_G^* \mathbf{t} \Leftrightarrow v : \sigma(\mathbf{v}) \wedge \mathbf{v} <: \mathbf{t} \Leftrightarrow v <: t$$

By applying [Eq. \(13\)](#) to type $t_\perp = v_0$, we get

$$v_0 \rightarrow_G^* t \Leftrightarrow t_\perp <: t,$$

i.e., $L(G) = L_{t_\perp}(\Delta)$ by definition.

Class table Δ conforms to the restrictions of the type system, $\langle M \rangle \vdash \Delta$. First, the class table is well-formed. Only types $v \in V$ inherit types of the form $\sigma(\mathbf{v})$ for $\sigma \in \Sigma$, so there are no inheritance cycles and no type parameters in the wrong position. Second, Δ does not employ contravariance and is non-expansive: The structure of inheritance rules guarantees that any type in $\text{cl}(t)$ is higher than t by one level at the most. $\square$

THEOREM 5.2. *Non-expansive class tables describe regular forests:*

$$\mathcal{L}(\langle C, M \rangle) \subseteq \text{NRF}$$

Intuition. A grammar variable $v \in V$ encodes a set of queries Q , $v = Q$. These queries restrict the set of types (trees) that v can derive. Each query $q \in Q$ is of the form $t \sim x$, where (i) t is a type, (ii) relation $\sim$ is subtyping ($\sim = <:$), suptyping (the opposite of subtyping, $\sim = >$), or equivalence ($\sim = =$), and (iii) symbol x is a placeholder for the type to be derived. If query $q = t \sim x$ is included in variable v , then our construction ensures that the type derived by v satisfies q :

$$t \sim x \in v, v \rightarrow^* t' \Rightarrow t \sim t'$$

For example, the initial variable v_0 includes a single query $t_\perp <: x$, so each type t' derived by v_0 satisfies $t_\perp <: t'$ —the set of these types is, by [Definition 4.2](#), exactly the language of the class table over $t_\perp$.

As the class table is non-expansive, both sets of variables V and productions R are finite and computable. If the initial query $t_\perp <: x$ is reduced to some other query $t \sim x$, then, as observed by [Kennedy and Pierce \[2007\]](#), type t is included in $t_\perp$'s inheritance and decomposition closure $\text{cl}(t_\perp)$ (see [Definition 3.1](#)). Because the class table is non-expansive, this set must be finite (see [Definition 3.2](#)). Thus, for a general query $t \sim x$, there are $|\text{cl}(t_\perp)|$ choices for t , three choices for $\sim$, and, consequently, $3|\text{cl}(t_\perp)|$ possible queries overall. As variables are sets of queries, we get that the set of variables V is finite, $|V| \leq 2^{3|\text{cl}(t_\perp)|}$. A variable v can derive terminal γ , $v ::= \gamma(\mathbf{v})$, if and only if the queries in v can be reduced to the queries in variables $\mathbf{v}$; variables $\mathbf{v}$ then recursively simulate the subtyping algorithm with γ 's children. In the proof below we show that as the class table is non-expansive, it is possible to compute the derived queries $\mathbf{v}$ in finite time.

Example. Consider the following non-expansive class table Δ , first presented in [Kennedy and Pierce \[2007\]](#):

$$\begin{aligned} N(-x) &: \emptyset \\ C &: NNC \end{aligned}$$

Let subtype $t_{\perp} = C$ and super-alphabet $\Sigma_{\top} = \{N, C\}$; then, class table Δ is encoded by the following grammar G :

$$\begin{array}{ll} v_0 = \{C <: x\} & v_0 \rightarrow C \\ v_1 = \{NC >: x\} & v_0 \rightarrow Nv_1 \\ & v_1 \rightarrow Nv_0 \end{array}$$

(Definitions not reachable from the initial variable are omitted.) The resulting language is indeed regular: $L(G) = \{N^{2k}C \mid k \geq 0\}$. The grammar productions simulate subtyping against Δ during derivation:

- Variable $v_0 = C <: x$ derives C because if $x = C$, $C <: C$ holds.
- Variable $v_0 = C <: x$ derives Nv_1 because if $x = Nx'$ for some type x' , the query in v_0 holds if x' satisfies the query in v_1 :

$$C <: Nx' \Leftarrow NNC <: Nx' \Leftarrow NC >: x' = v_1.$$

- Similarly, variable $v_1 = NC >: x$ derives Nv_0 because if $x = Nx'$ for some type x' , the query in v_1 holds if x' satisfies the query in v_0 :

$$NC >: Nx' \Leftarrow C <: x' = v_0.$$

- Nevertheless, $v_1 = NC >: x$ cannot derive C , because the query $NC >: C$ leads to an infinite proof:

$$NC >: C \Leftarrow NC >: NNC \Leftarrow C <: NC \Leftarrow \dots \quad (14)$$

As the infinite proof in the latter case is cyclic, it can be identified during the grammar construction.

Additional notations. A query of the form $t \sim x$, where t is a type, x is a type parameter, and $\sim \in \{<:, =, >:\}$, is denoted by q . A set of queries is denoted by Q . The set of Q queries where parameter x is substituted with term τ is denoted by $Q[\tau]$:

$$Q[\tau] = \{t \sim \tau \mid t \sim x \in Q\}$$

The conjunction of the queries in Q is denoted by $\bigwedge Q$. We write $\bigwedge Q \vdash \bigwedge Q'$ to say that the queries in Q imply the queries in Q' , i.e., $\bigwedge Q \Rightarrow \bigwedge Q'$. If the queries in Q are true, we write $\vdash \bigwedge Q$. A set of queries is minimal if it does not contain redundant queries, e.g., including $t <: x$ in addition to $t = x$ although the latter implies the former.

PROOF OF THEOREM 5.2. Let Δ be a class table of type system $\langle C, M \rangle$, and consider subtyping-queries between type $t_{\perp}$ and types over super-alphabet $\Sigma_{\top}$. We show that the set of $t_{\perp}$ supertypes is a regular forest, described by a regular tree grammar $G = \langle \Sigma, V, v_0, R \rangle$.

Construction. Let $\Sigma = \Sigma_{\top}$. The set of variables V contains all sets of queries over $\text{cl}(t_{\perp})$:

$$V = \{Q \mid \forall (t \sim x) \in Q, t \in \text{cl}(t_{\perp}), \sim \in \{<:, =, >:\}\}$$

Let $v_0 = \{t_{\perp} <: x\}$. For every variable $v = Q \in V$ and terminal $\sigma \in \Sigma$, $\text{rank}(\sigma) = n$, add derivation rule $v ::= \sigma(v)$ to R if and only if there exists a sequence of variables $v = Q$, such that the sets in Q are minimal and $\bigwedge Q' \vdash \bigwedge Q''$, for set Q' and sets Q'' ,

- (1) $Q' = Q[\sigma(x)] = \{t \sim \sigma(x_1, \dots, x_n) \mid (t \sim x) \in Q\}$
- (2) $\forall i = 1, \dots, n, Q'_i = Q_i[x_i] = \{t \sim x_i \mid (t \sim x) \in Q_i\}$

Expressed in words, parameter x in v is replaced with the term $\sigma(x_1, \dots, x_n)$, $n = \text{rank}(\sigma)$; the type parameters x in v are enumerated according to their position in the sequence; then we require that the queries of v prove all the queries in v . As the inheritance and decomposition closure of $t_{\perp}$ is always finite, variables v can be found by traversing every possible proof of $t \sim \sigma(x)$.

Correctness of the construction. Observe that the grammar is well-defined. As Δ is non-expansive, $\text{cl}(t_{\perp})$ is finite, thus V is finite. For every R production $v ::= \sigma(v)$, types t on the left-hand side

of v queries are drawn from $\text{cl}(t_\perp)$, so their sub-queries contain types also drawn from $\text{cl}(t_\perp)$ by definition of the closure. Therefore, $v \subseteq V$, and thus R is finite.

We show that variable $v = Q$ derives tree t , $v \rightarrow_G^* t$ if and only if substituting parameter x in Q 's queries with type t satisfies them:

$$v \rightarrow_G^* t \Leftrightarrow \vdash \bigwedge v[t] \quad (15)$$

Eq. (15) is proved by induction on the height of t , $n = \text{height}(t)$.

If $n = 0$, then t is a leaf, $t = \sigma$. Variable v , describing a set of queries Q , produces σ if and only if $v ::= \sigma \in R$. For nodes of rank 0, the construction introduces transition $v ::= \sigma$ if and only if the queries in set $Q' = \{t \sim \sigma \mid (t \sim x) \in Q\}$ are satisfied (by empty Q'), i.e., if replacing parameter x in Q with $t = \sigma$ satisfies its queries:

$$v \rightarrow_G^* \sigma \Leftrightarrow v ::= \sigma \in R \Leftrightarrow \vdash \bigwedge v[\sigma]$$

Suppose Eq. (15) holds for $n \geq 0$, and let t be a tree of height $n + 1$, $t = \sigma(t)$, $\sigma \in \Sigma$, $t \in \Sigma^\Delta$, where the heights of t are, at most, n . Variable v derives t , $v \rightarrow_G^* t$ if and only if v derives σ with production $v ::= \sigma(v) \in R$, where variables v derive σ 's children, $v \rightarrow_G^* t$. According to the construction, $v ::= \sigma(v) \in R$ if and only if $\bigwedge v[x] \vdash \bigwedge v[\sigma(x)]$. By the inductive assumption, $\vdash \bigwedge v[t]$ if and only if $v \rightarrow_G^* t$, as the heights of types in t are not higher than n . Thus, by assigning t to x we get that $\vdash \bigwedge v[t] \vdash \bigwedge v[\sigma(t)]$, or $\vdash \bigwedge v[\sigma(t)]$ in short:

$$v \rightarrow_G^* \sigma(t) \Leftrightarrow v ::= \sigma(v) \in R \wedge v \rightarrow_G^* t \Leftrightarrow \vdash \bigwedge v[\sigma(t)]$$

By applying Eq. (15) to $v_0 = \{t_\perp <: x\}$, we get that

$$v_0 \rightarrow_G^* t \Leftrightarrow \vdash \bigwedge \{t_\perp <: t\} \Leftrightarrow t_\perp <: t,$$

and, therefore, $L(G) = L_{t_\perp}(\Delta)$ by definition. $\square$

From Theorems 5.1 and 5.2, it follows that type systems $\langle M \rangle$ and $\langle C, M \rangle$ are precisely as complex as regular forests.

COROLLARY 5.3. *Class tables of non-expansive type systems describe regular forests:*

$$\mathcal{L}(\langle M \rangle) = \mathcal{L}(\langle C, M \rangle) = \text{NRF}$$

PROOF. Consider that (i) $\text{NRF} \subseteq \mathcal{L}(\langle M \rangle)$ by Theorem 5.1, (ii) $L(\langle M \rangle) \subseteq \mathcal{L}(\langle C, M \rangle)$ as any $\langle M \rangle$ class table conforms to $\langle C, M \rangle$, $\langle M \rangle \subset \langle C, M \rangle$, and that (iii) $\mathcal{L}(\langle C, M \rangle) \subseteq \text{NRF}$ by Theorem 5.2; thus,

$$\text{NRF} \subseteq \mathcal{L}(\langle M \rangle) \subseteq \mathcal{L}(\langle C, M \rangle) \subseteq \text{NRF} \Rightarrow \mathcal{L}(\langle M \rangle) = \mathcal{L}(\langle C, M \rangle) = \text{NRF} \quad \square$$

When both expansive inheritance and multiple instantiation inheritance are taken away, the complexity of the type system is trimmed down to deterministic forests.

PROPOSITION 5.4. *Class tables of non-expansive type systems with single instantiation inheritance describe deterministic regular forests:*

$$\mathcal{L}(\mathcal{S}_\perp) = \mathcal{L}(\langle C \rangle) = \text{DRF}$$

PROOF. By repeating the constructions in the proofs of Theorem 5.1 and Theorem 5.2 without multiple instantiation inheritance, we get the corresponding results with deterministic regular forests (and grammars).

In the proof of [Theorem 5.1](#), a production $v ::= \sigma(v)$ is directly encoded as the inheritance rule $v : \sigma(v)$. Thus, the determinism of the grammar coincides with the class table's single instantiation inheritance:

$$\frac{v ::= \sigma(v_1) \wedge v ::= \sigma(v_2) \Leftrightarrow v : \sigma(v_1), \sigma(v_2) \quad v ::= \sigma(v_1) \wedge v ::= \sigma(v_2) \Rightarrow v_1 = v_2}{v : \sigma(v_1), \sigma(v_2) \Rightarrow v_1 = v_2}$$

If a class table employs single instantiation inheritance, then the subtyping algorithm, consisting of the VAR and SUPER rules, becomes deterministic. In the proof of [Theorem 5.2](#), a production $v ::= \sigma(v)$ is introduced if and only if the queries encoded by v prove the queries of v when node σ is encountered. As subtyping is deterministic, at most one minimal such sequence of queries v exists. $\square$

5.2 Non-Contravariant Subtyping Is Context-Free

Yet another way to make subtyping decidable is by prohibiting contravariant type parameters. Computationally, this restriction is less severe than non-expansive inheritance. We show that non-contravariant class tables are as expressive as context-free forests with grammars in GNF. Recall that Greibach productions are of the form $v(x) ::= \sigma(\tau)$, i.e., where the root of the derived tree form is terminal, $\sigma \in \Sigma$.

In this section, we use the extended variant of context-free grammars, defined in [Definition 2.1](#), where the initial tree form can be any tree—not just a variable. Recall that in [Lemma 2.2](#) we proved that extended context-free tree grammars are equivalent in expressiveness to standard context-free grammars.

THEOREM 5.5. *Any context-free tree grammar in GNF can be encoded in a non-contravariant class table:*

$$\text{NCF}_{\text{GNF}} \subseteq \mathcal{L}(\langle X, M \rangle)$$

Intuition. As in the regular case ([Theorem 5.1](#)), we want to use subtyping to simulate grammar derivation. Also, as before, we directly encode productions as inheritance rules and use covariance to enable recursive subtyping. In this case, however, we harness expansively-recursive inheritance to encode the slightly more complex structure of context-free productions.

Example. Consider our palindrome running example. The context-free grammar productions in [Eq. \(6\)](#) (shown in [Section 2](#)), together with the initial tree $t_0 = v_0 E$, derive palindromes over $\{a, b\}$ terminated by end-marker E . We already encoded this grammar as a program in abstract syntax in [Eq. \(7\)](#) (shown in [Section 3](#)) and in $C^\#$ in [List. 1](#) (shown in [Section 1](#)). Notice that the parameterized variable $v_0(x)$ is encoded by the polymorphic `interface v0<x>`.

PROOF OF THEOREM 5.5. Let $G = \langle \Sigma, V, t_0, R \rangle$ be an extended context-free tree grammar in GNF. We encode grammar G as class table Δ , $\langle X, M \rangle \vdash \Delta$, such that the language of Δ (over $t_\perp = t_0$) is exactly the extent of the grammar, $L(G)$.

Construction. Class names are drawn from $\Gamma = \Sigma \cup V$. The rank of each type $\gamma \in \Sigma \cup V$ is the same as its rank in the grammar. The parameters of class γ are covariant if $\gamma \in \Sigma$, or invariant if $\gamma \in V$. The subtype in queries is $t_\perp = t_0$, and supertype names are drawn from $\Sigma_\top = \Sigma$. For every production $v(x) ::= \sigma(\tau) \in R$, introduce inheritance rule $v(\circ x) : \sigma(\tau)$ to Δ .

Correctness of the construction. We show that tree form t derives terminal tree $t' \in \Sigma^\Delta$ if and only if type t is a subtype of type t' :

$$t \rightarrow_G^* t' \Leftrightarrow t <: t', t' \in \Sigma^\Delta \quad (16)$$

Proof is by induction on the height of t' , $n = \text{height}(t')$.

If $n = 0$, then t' is a leaf, $t' = \sigma$. Tree form t derives σ only if $t = \sigma$ (in which case, $t <: \sigma$ is immediate), or if the root of t is a variable v , and there is a production $v(x) ::= \sigma \in R$. In the latter case, class v inherits class σ , so type t is still a subtype of σ :

$$v(t) \rightarrow_G^* \sigma \Leftrightarrow v(x) ::= \sigma \in R \Leftrightarrow v(\circ x) : \sigma \in \Delta \Leftrightarrow v(t) <: \sigma$$

Assume Eq. (16) holds for all trees of height n , and let us prove this for tree $t' = \sigma(t')$ of height $n+1$. Let $t = \xi(t)$. If the root of t is a terminal node $\xi \in \Sigma$, then it must be σ to derive t' . To complete the derivation, we require the children of t to derive the children of t' , $t \rightarrow_G^* t'$. In this case, type $\sigma(t)$ is a subtype of type $\sigma(t')$, by applying the VAR rule and σ 's covariance to get $t <: t'$, which is true by the inductive assumption, given that $\text{height}(t') \leq n$:

$$\sigma(t) \rightarrow_G^* \sigma(t') \Leftrightarrow t \rightarrow_G^* t' \Leftrightarrow t <: t' \Leftrightarrow \sigma(t) <: \sigma(t')$$

Otherwise the root of t is a variable, $\xi = v \in V$. To derive t' , there must be a production $\rho = v(x) ::= \sigma(\tau) \in R$, and its application to t has to produce t' ,

$$\tau[x \leftarrow t] \rightarrow_G^* t'.$$

Given production ρ , the construction introduces the corresponding inheritance rule $v(\circ x) : \sigma(\tau) \in \Delta$, so type t is a subtype of type t' : By applying the SUPER rule, replacing $v(t)$ with $\sigma(\tau[x \leftarrow t])$, and then the VAR rule, to reduce the query to $\tau[x \leftarrow t] <: t'$, which is true by the inductive assumption since $\text{height}(t') \leq n$:

$$\begin{aligned} v(t) \rightarrow_G^* \sigma(t') &\Leftrightarrow v(x) ::= \sigma(\tau) \wedge \tau[x \leftarrow t] \rightarrow_G^* t' \\ &\Leftrightarrow v(x) : \sigma(\tau) \wedge \tau[x \leftarrow t] <: t' \\ &\Leftrightarrow v(t) <: \sigma(t') \end{aligned}$$

By applying Eq. (16) to $t_\perp = t_0$, we get that $t_0 \rightarrow_G^* t \Leftrightarrow t_\perp <: t$; thus, $L(G) = L_{t_\perp}(\Delta)$ by definition. $\square$

THEOREM 5.6. *Non-contravariant class tables describe context-free grammars in GNF:*

$$\mathcal{L}(\langle X, M \rangle) \subseteq \text{NCF}_{\text{GNF}}$$

Intuition. For starters, assume that all the type parameters in the class table are covariant (no invariance). Consider the subtyping query $t <: t'$ where $t = \gamma(t)$ and $t' = \gamma'(t')$. To prove this query, we first use the SUPER rule to replace t with $\gamma'(\tau[x \leftarrow t])$, assuming that class $\gamma(x)$ transitively inherits ($:^*$) the type $\gamma'(\tau)$. In the next step, we use the VAR rule to recursively validate the children of γ' , $\tau[x \leftarrow t] <: t'$. To simulate subtyping with grammar derivation, we simply encode transitive inheritance as productions, i.e., if $\gamma(x) :^* \gamma'(\tau)$, $\gamma(x) ::= \gamma'(\tau)$. Then, the query $t <: t'$ is true if and only if $t \rightarrow^* t'$ because, as in the subtyping case, $t = \gamma(t)$ derives $\gamma'(\tau[x \leftarrow t])$ and then $\tau[x \leftarrow t]$ recursively derive t' :

$$\begin{aligned} \gamma(t) <: \gamma'(t') &\Leftrightarrow \gamma(x) :^* \gamma'(\tau) \wedge \tau[x \leftarrow t] <: t' \\ \gamma(t) \rightarrow^* \gamma'(t') &\Leftrightarrow \gamma(x) ::= \gamma'(\tau) \wedge \tau[x \leftarrow t] \rightarrow^* t' \end{aligned}$$

Nonetheless, this clean correspondence between subtyping and context-free grammars collapses when invariance is re-introduced to the class table. For instance, consider the following inheritance rules:

$$\begin{aligned} a(+x, \circ y) &: \emptyset \\ b(\circ z) &: a(z, z) \end{aligned} \tag{17}$$

In the query $b(t) <: a(t_1, t_2)$, we use inheritance and decomposition to deduce both $t <: t_1$ and $t = t_2$. If type t were to be encoded directly as a tree form, as before, then it should be able to derive the two trees t_1 and t_2 , while simulating both covariance and invariance simultaneously.

To reproduce this behavior in a context-free tree grammar, we encode the kind of position of type t in the tree form that represents it. Let γ be the root of t . If t appears in a covariant position, then γ is encoded by variable γ_+ , or otherwise by variable $\gamma_\circ$ in an invariant position. If t is in a covariant position, $t <: t_1$, variable γ_+ derives the tree t' if and only if $t' <: t_1$; otherwise, if t is in an invariant position, $t = t_2$, variable $\gamma_\circ$ derives the tree t' if and only if $t' = t_2$.

Variable $\gamma_\circ$ simply derives the node γ . If type t is in an invariant position, $t = t_2$, we encode t by annotating each of its nodes with $\circ$, yielding the tree form $t_\circ$. This tree form derives, recursively, the tree t exactly, so

$$t_\circ \rightarrow^* t' \Leftrightarrow t' = t \quad \Rightarrow \quad t_\circ \rightarrow^* t_2 \Leftrightarrow t = t_2$$

as required. On the other hand, if t appears in a covariant position, $t <: t_1$, then its children may appear in either subtyping or equivalence queries (or even both simultaneously). Let $t = \gamma(t)$ and t_i be some child of γ . We encode t_i as two tree forms, t_{i+} and $t_{i\circ}$. Then, we can use t_{i+} whenever t_i is covariant, or otherwise $t_{i\circ}$ when t_i is invariant. Thus, type $t = \gamma(t)$ is encoded as the tree form $t_+ = \gamma_+(t_+, t_\circ)$, where children t_+ are, recursively, encoded in a similar way, and $t_\circ$ are encoded as mentioned above.

Continuing our example, we get that class b in Eq. (17) is encoded by the following productions:

$$b_+(z_+, z_\circ) ::= a(z_+, z_\circ) \quad b_\circ(z_\circ) ::= a(z_\circ, z_\circ)$$

Variable b_+ encodes class b in covariant positions, so the first argument to type a is in a covariant position, and z_+ is used, while the second position is invariant, and $z_\circ$ is picked instead. Parameter $z_\circ$ in variable $b_\circ$ fits both positions, as the initial position (of root a) is invariant.

Additional notations. To encode unground type τ as invariant tree form $\tau_\circ$, add annotation $\circ$ to all its nodes:

$$\tau_\circ = \begin{cases} x_\circ & \tau = x \\ \gamma_\circ(\tau_\circ) & \tau = \gamma(\tau) \end{cases} \quad (18)$$

We get the covariant tree form τ_+ by encoding both the covariant and invariant versions of its sub-trees:

$$\tau_+ = \begin{cases} x_+ & \tau = x \\ \gamma_+(\tau_+, \tau_\circ) & \tau = \gamma(\tau) \end{cases} \quad (19)$$

To decode an invariant tree form $\tau_\circ$ into the tree it represents τ , simply remove all the $\circ$ annotations from its nodes. If the tree form is covariant τ_+ , also remove any $+$ annotations, and trim-out half of the sub-trees:

$$\tau = \begin{cases} x & \tau_\circ = x_\circ \vee \tau_+ = x_+ \\ \gamma(\tau) & \tau_\circ = \gamma_\circ(\tau_\circ) \vee \tau_+ = \gamma_+(\tau_+, \tau_\circ) \end{cases} \quad (20)$$

Observe that if variables x_+ in covariant tree form τ_+ are substituted with covariant trees t_+ , and the invariant variables $x_\circ$ are substituted with the corresponding invariant trees $t_\circ$, then the resulting tree $\tau_+[x_+ \leftarrow t_+, x_\circ \leftarrow t_\circ]$ covariantly encodes type $\tau[x \leftarrow t]$. The same holds for invariant tree forms:

$$\tau_\circ[x_+ \leftarrow t_+, x_\circ \leftarrow t_\circ] = (\tau[x \leftarrow t])_\circ, \diamond \in \{+, \circ\} \quad (21)$$

(The proof is technical, by induction on the height of $\tau_\circ$.)

PROOF OF THEOREM 5.6. Given class table Δ , $\langle X, M \rangle \vdash \Delta$, fixed subtype $t_\perp$, and super-alphabet $\Sigma_\top$, we construct extended context-free tree grammar $G = \langle \Sigma, V, t_0, R \rangle$ that describes the same language as Δ , $L(G) = L_{t_\perp}(\Delta)$.

In this proof we use covariant and invariant tree encoding of types, defined in Eqs. (18) to (20).

Construction. Let $\Sigma = \Sigma_\top$. For every class name $\gamma \in \Gamma$, add variables γ_+ and γ_o to V , where $\text{rank}(\gamma_o) = \text{rank}(\gamma)$ but $\text{rank}(\gamma_+) = 2 \cdot \text{rank}(\gamma)$. Let $t_0 = (t_\perp)_+$. For each terminal type $\sigma \in \Sigma_\top$, add rule $\sigma_o(x_o) ::= \sigma(x_o)$ to R . For each class name $\gamma \in \Gamma$ and terminal type $\sigma \in \Sigma_\top$ of rank n , if class γ inherits (or, is) σ , $\gamma(x) :^* \sigma(\tau)$, introduce the following production to R :

$$\gamma_+(x_+, x_o) ::= \sigma(\tau_{\diamond^\sigma}) = \sigma((\tau_1)_{\diamond_1^\sigma}, \dots, (\tau_n)_{\diamond_n^\sigma})$$

(Recall that $\diamond^\sigma$ denotes the sequence of parameter variances of class σ , and $\diamond_i^\sigma$ denotes the variance of the i^{th} parameter.) As inheritance is acyclic, transitive inheritance $:^*$ is finite, so the construction of R must terminate.

Correctness of the construction. Each variable σ_o , for terminal node $\sigma \in \Sigma$, directly derives σ . Therefore, any invariant tree form t_o over Σ , $t \in \Sigma^\Delta$ derives the terminal tree t , and only if:

$$t_o \rightarrow_G^* t', t \in \Sigma^\Delta \Leftrightarrow t = t' \quad (22)$$

Conversely, let $t_+ = \gamma_+(t_+, t_o)$ be a covariant tree form. We show that t_+ derives tree t' if and only if type t is a subtype of type t' :

$$t_+ \rightarrow_G^* t' \Leftrightarrow t <: t' \quad (23)$$

Eq. (23) is proved by induction on the height of t' , $n = \text{height}(t')$.

If $n = 0$, then t' is a leaf, $t' = \sigma$. The root of t_+ , variable γ_+ , derives σ if and only if $\gamma = \sigma$, or if class γ inherits class σ , $\gamma(x) :^* \sigma$; in both cases the construction introduces production $\gamma_+(x_+, x_o) ::= \sigma$ to R :

$$\gamma_+(t_+, t_o) \rightarrow_G^* \sigma \Leftrightarrow \gamma_+(x_+, x_o) ::= \sigma \in R \Leftrightarrow \gamma(\diamond x) :^* \sigma \Leftrightarrow \gamma(t) <: \sigma$$

Assume Eq. (23) holds for $n \geq 0$, and let us prove it for tree t' of height $n + 1$, $t' = \sigma(t')$. As grammar G is in GNF, tree form t_+ can derive terminal tree t' if and only if (i) there is an R production that derives γ_+ , the root of t_+ , to σ , the root of t' , $\gamma_+(x_+, x_o) ::= \sigma(\tau_{\diamond^\sigma}) \in R$, such that (ii) the resulting sub-tree forms $\tau_{\diamond^\sigma}$, substituted with t_+ and t_o , derive σ 's children, t' :

$$\begin{aligned} \gamma_+(t_+, t_o) \rightarrow_G^* \sigma(t') &\Leftrightarrow \\ \gamma_+(x_+, x_o) ::= \sigma(\tau_{\diamond^\sigma}) \in R \wedge \tau_{\diamond^\sigma}[x_+ \leftarrow t_+, x_o \leftarrow t_o] \rightarrow_G^* t' &\quad (24) \end{aligned}$$

Now focus on tree forms $\tau_{\diamond^\sigma}[x_+ \leftarrow t_+, x_o \leftarrow t_o]$ in Eq. (24). Following the observation of Eq. (21), these forms encode the unground types τ assigned with types t , while respecting the original variances of σ , $\diamond^\sigma$:

$$\tau_{\diamond^\sigma}[x_+ \leftarrow t_+, x_o \leftarrow t_o] = (\tau[x \leftarrow t])_{\diamond^\sigma}$$

Let us then simplify Eq. (24) by writing:

$$(24) \Leftrightarrow \gamma_+(x_+, x_o) ::= \sigma(\tau_{\diamond^\sigma}) \in R \wedge (\tau[x \leftarrow t])_{\diamond^\sigma} \rightarrow_G^* t' \quad (25)$$

Focus on tree forms $(\tau[x \leftarrow t])_{\diamond^\sigma}$ in Eq. (25). If tree form $(\tau_i[x \leftarrow t])_{\diamond_i^\sigma}$ is invariant, $\diamond_i^\sigma = \circ$, then by Eq. (22),

$$(\tau_i[x \leftarrow t])_{\diamond_i^\sigma} \rightarrow_G^* t'_i \Leftrightarrow \tau_i[x \leftarrow t] = t'_i \quad (26)$$

Otherwise the tree form is covariant, $\diamond_i^\sigma = +$. As t'_i is a sub-tree of t' , its height is n at most, so we can apply the inductive assumption to get

$$(\tau_i[x \leftarrow t])_+ \rightarrow_G^* t'_i \Leftrightarrow \tau_i[x \leftarrow t] <: t'_i \quad (27)$$

Production $\gamma_+(x_+, x_o) ::= \sigma(\tau_{\diamond\sigma})$ is induced by the construction following the inheritance rule $\gamma(\diamond x) :^* \sigma(\tau)$. Thus, type $\gamma(t)$ is a subtype of $\sigma(\tau[x \leftarrow t])$ by inheritance, which is a subtype of $t' = \sigma(t')$ by decomposition. In Eqs. (26) and (27), we showed that child types $\tau[x \leftarrow t]$ and t' respect σ 's variances, so decomposition can be applied. Let us extend Eq. (25):

$$\begin{aligned} (25) &\Leftrightarrow \gamma(\diamond x) :^* \sigma(\tau) \wedge \forall i, \begin{cases} \tau_i[x \leftarrow t] = t'_i & \diamond_i^\sigma = \circ \\ \tau_i[x \leftarrow t] <: t'_i & \diamond_i^\sigma = + \end{cases} \\ &\Leftrightarrow \gamma(t) <: \sigma(t') \end{aligned} \quad (28)$$

Overall, from Eqs. (24), (25) and (28), we have

$$\gamma_+(t_+, t_o) \rightarrow_G^* \sigma(t') \Leftrightarrow \gamma(t) <: \sigma(t'),$$

concluding the induction.

Finally, we can apply Eq. (23) to $t_0 = (t_\perp)_+$ to get $t_0 \rightarrow_G^* t \Leftrightarrow t_\perp <: t$; thus, $L(G) = L_{t_\perp}(\Delta)$ by definition. $\square$

From Theorems 5.5 and 5.6, it follows that the computational class of type system $\langle X, M \rangle$ is described by context-free tree grammars in GNF.

COROLLARY 5.7. *Non-contravariant class tables describe context-free forests that have grammars in GNF:*

$$\mathcal{L}(\langle X, M \rangle) = \text{NCF}_{\text{GNF}}$$

PROOF. $\text{NCF}_{\text{GNF}} \subseteq \mathcal{L}(\langle X, M \rangle)$ by Theorem 5.5; $\mathcal{L}(\langle X, M \rangle) \subseteq \text{NCF}_{\text{GNF}}$ by Theorem 5.6; $\text{NCF}_{\text{GNF}} \subseteq \mathcal{L}(\langle X, M \rangle) \subseteq \text{NCF}_{\text{GNF}} \Rightarrow \mathcal{L}(\langle X, M \rangle) = \text{CFF}_{\text{GNF}}$. $\square$

Type system $\langle X \rangle$ is the subset of $\langle X, M \rangle$ employing single-instantiation inheritance. As before (Proposition 5.4), single instantiation inheritance reduces the complexity of the type system to its deterministic variant.

PROPOSITION 5.8. *Class tables of non-contravariant type systems with single instantiation inheritance describe deterministic context-free forests:*

$$\mathcal{L}(\langle X \rangle) = \text{DCF}$$

PROOF. Every deterministic context-free forest has a deterministic tree grammar in GNF [Guesarian 1983]. Again, by repeating the proof of Theorems 5.5 and 5.6 with single instantiation inheritance, we get the corresponding results regarding DCF.

In the proof of Theorem 5.5, we encode each production $v(x) ::= \tau$ directly as the inheritance rule $v(\circ x) ::= \tau$. Therefore, the grammar's determinism ensures the resulting class table does not use multiple instantiation inheritance:

$$\frac{\begin{aligned} v(x) ::= \sigma(\tau_1) \wedge v(x) ::= \sigma(\tau_2) &\Leftrightarrow v(\circ x) : \sigma(\tau_1), \sigma(\tau_2) \\ v(x) ::= \sigma(\tau_1) \wedge v(x) ::= \sigma(\tau_2) &\Rightarrow \tau_1 = \tau_2 \end{aligned}}{v(\circ x) : \sigma(\tau_1), \sigma(\tau_2) \Rightarrow \tau_1 = \tau_2}$$

In the proof of Theorem 5.6, we derive (i) $\gamma_\circ$ to σ , which is always deterministic, and (ii) γ_+ to $\sigma(\tau_{\diamond\sigma})$ if and only if γ inherits $\sigma(\tau)$, which is also deterministic when the class table employs single instantiation inheritance:

$$\frac{\begin{aligned} \gamma(\diamond x) :^* \sigma(\tau_1), \sigma(\tau_2) &\Leftrightarrow \gamma_+(x_+, x_o) ::= \sigma((\tau_1)_{\diamond\sigma}) \wedge \gamma_+(x_+, x_o) ::= \sigma((\tau_2)_{\diamond\sigma}) \\ \gamma(\diamond x) :^* \sigma(\tau_1), \sigma(\tau_2) &\Rightarrow \tau_1 = \tau_2 \end{aligned}}{\gamma_+(x_+, x_o) ::= \sigma((\tau_1)_{\diamond\sigma}) \wedge \gamma_+(x_+, x_o) ::= \sigma((\tau_2)_{\diamond\sigma}) \Leftrightarrow (\tau_1)_{\diamond\sigma} = (\tau_2)_{\diamond\sigma}} \quad \square$$

6 TREETOP—A CONTEXT-FREE API GENERATOR

Upon its introduction [Fowler 2005], fluent API was recognized as a promising method for embedding domain-specific languages. As fluent API methods are called in a chain, and not imperatively, they can enforce a DSL syntax, or an API protocol in general, at compile time. (See a brief introduction of the concept in Section 1.2 and more detailed discussions by Gil and Roth [2019] and Yamazaki et al. [2019].) The subject attracted academic interest, as researchers introduced new fluent API techniques supporting increasingly complex DSLs [Gil and Roth 2020; Gil and Levy 2016; Gil and Roth 2019; Nakamaru and Chiba 2020; Nakamaru et al. 2017; Xu 2010; Yamazaki et al. 2019]. Because the study of fluent API is motivated by its practical applications, results in the field are given in the form of actual tools: fluent API generators. A generator accepts a DSL grammar as input, and encodes it as a fluent API in the target programming language. The quality of a generator is determined by its domain, i.e., the class of languages (and grammars) for which it can generate fluent APIs. Generated fluent APIs also need to be practical, in terms of code size and compilation time. The last two significant works on fluent APIs, published by Gil and Roth [2019] and Yamazaki et al. [2019], presented the generators Fling and TypeLevelLR that support all deterministic context-free languages (DCFLs). (Gil and Roth [2020] later proved that the methods implemented in these generators cannot support languages beyond DCFL.) Programs generated by Fling and TypeLevelLR were shown to compile in linear time, making both tools practical solutions for embedding DSLs.

The idea of integrating subtyping into fluent APIs was first proposed by Grigore [2017]. In his design, the role of the fluent API becomes merely to encode the input DSL program in a type, so it can be accepted by a subtyping machine. Grigore’s subtyping machine implementation of the CYK parser was introduced as a part of a fluent API generator that supports all context-free languages. Moving up from DCFL to CFL is a welcome improvement: CFLs are described by general context-free grammars, which are more expressive than their deterministic variants (LR, LL, LALR, etc.) and are easier to compose, as they produce no conflicts. Unfortunately, it turned out that Grigore’s generator is not practical. The APIs it generates take a lot of time to compile, and cause the compiler to crash with stack overflow even for basic examples.

Theorem 5.5 presented in Section 5 states that any tree grammar in GNF can be recognized by a subtyping machine with expansive inheritance and multiple instantiation inheritance. As all context-free languages have grammars in GNF [Greibach 1965], this means we just found a new method for encoding CFLs using subtyping that, in contrast to Grigore’s construction, does not employ contravariance. Nevertheless, can this method be used in a real programming language, such as C[#]? Unlike Java, C[#] endorses multiple instantiation inheritance (required by our construction)—but its version of this feature is weaker than the one used in this paper. In C[#], if class γ has (transitive) supertypes τ and τ' , then there must not be a substitution s that unifies τ and τ' :

$$\frac{\tau \neq \tau' \wedge \gamma :^+ \tau \Rightarrow \nexists s, \tau[s] = \tau'[s]}{\gamma :^+ \tau' \checkmark} \quad (29)$$

(This trait of C[#] was also discussed by Kennedy and Pierce [2007].) It is possible that Theorem 5.5 cannot be utilized to its full potential in C[#], as the construction described in its proof does not conform to the restriction of Eq. (29) in the general case. Nevertheless, the construction does work for all CFLs. If we encode a context-free production $vx ::= \alpha x$ as the inheritance rule $v(x) : \alpha x$, then unification with another rule $v(x) : \alpha' x$ can only occur for $\alpha' = \alpha$. To prevent this, we simply ignore duplicated productions.

Thus, [Theorem 5.5](#) paved the way for the development of our new API generator, Treetop [[Roth 2021b](#)], used to embed CFLs in C[#]¹. Treetop is a C[#] *source generator*, a project preprocessor that creates source files, similarly to Java’s annotations processors. Source generation is an experimental Visual Studio feature for C[#] introduced in 2020². After an initial deployment, Treetop automatically transforms DSL grammars into DSL APIs during the standard compilation cycle. Given a context-free grammar G in file `MyDSL.cfg` describing a DSL, Treetop converts it into an equivalent grammar in GNF, constructs class table $\Delta \in \langle X, M \rangle$ that recognizes the DSL, $L(\Delta) = L(G)$, and encodes the result as C[#] interfaces in a new source file `MyDSLAPI.cs`. On request, Treetop also generates a subtyping-based fluent API for the DSL, following Grigore’s design.

The full version of the article [[Roth 2021a](#)] demonstrates how Treetop can be used to create a type-safe API for the *Canvas* Android library. [Ferles et al. \[2020\]](#) identified a restriction in the Canvas API protocol that can be described by a context-free grammar. The full version describes Treetop’s process flow, from encoding the Canvas grammar as a subtyping machine to integrating it into a C[#] fluent API.

In theory, the APIs generated by Treetop should be efficient in both space (code size) and time (compilation time of API invocations). The class table construction described in the proof of [Theorem 5.5](#) is straightforward, so the crucial step of the encoding is the conversion of the input grammar to GNF. This transformation is known to introduce a polynomial increase in the size of the grammar [[Blum and Koch 1999](#)], which is translated to a similar increase in code size. An API generated by Treetop is invoked in subtyping queries, whose compilation time depends on the compiler implementation. A specialized compiler can reduce Treetop’s interfaces and queries back to CFGs and strings, whose membership problem is solved in polynomial time [[Valiant 1975](#)].

The full version of the article [[Roth 2021a](#)] describes an experiment conducted to test Treetop’s performance in practice, using the CSC C[#] compiler. In the experiment, we used Treetop to compile different grammars into C[#] subtyping machines, and then we measured the compilation times of subtyping queries of various lengths. The results of the experiment show that some subtyping machines compile relatively fast even for large queries, but others lead to long compilation times.

7 DISCUSSION

A program that uses the subtyping algorithm to conduct a computation is called a subtyping machine. By formalizing subtyping machines as forest recognizers, we tie the decidable subtyping type systems of Kennedy and Pierce with familiar families of tree grammars. We proved that non-expansively-recursive subtyping machines recognize regular forests ([Corollary 5.3](#)) and non-contravariant machines recognize context-free forests with GNF ([Corollary 5.7](#)). It was also shown that multiple instantiation inheritance corresponds to grammar non-determinism ([Propositions 5.4](#) and [5.8](#)). The results are depicted in [Fig. 1](#) in [Section 1](#).

In non-expansive subtyping, all infinite proofs are cyclic, so they can be detected by recording all visited sub-queries. The construction in the proof of [Theorem 5.2](#) encodes cyclic proofs as recursive productions $v ::= v$ —which are immediately removed from the grammar. This means that infinite proofs can be intercepted in pre-computation, i.e., even before the subtyping query is inspected. Our construction, however, assumes that one of the types in the query is known and fixed, following our definition of the language of a class table ([Definition 4.2](#)). In the future, our method may be developed into a better algorithm for non-expansive subtyping that does not record sub-queries.

Context-free tree grammars (CFTGs) are recognized by pushdown tree automata (PDTA), that use an *auxiliary tree store* and accept a *tree input* [[Guessarian 1983](#)]. The proof of [Theorem 5.6](#)

¹The Treetop project is hosted at <https://github.com/OriRoth/treetop>.

²For more details, see <https://devblogs.microsoft.com/dotnet/introducing-c-source-generators/>.

describes a reduction of non-contravariant subtyping machines to CFTGs. In contrast to the regular case, here the left-hand side of the subtyping query $t_{\perp}$ is encoded into the grammar initial tree form t_0 and is not involved with grammar productions. Therefore, we can resolve non-contravariant subtyping using modified PDTAs, which accept a second input tree and use this as their initial store. Given query $t_{\perp} <: t_{\top}$, we assign $t_{\perp}$ to the tree store (possibly after the transformation of Eq. (19)) and run the PDTA on input $t_{\top}$. As our construction yields grammars in GNF, their corresponding PDTAs run in real time [Guessarian 1983]. If multiple instantiation inheritance is also restricted, then the grammars become deterministic and so does the PDTA run time. Thus, subtyping in $\langle X, M \rangle$ can be resolved by a PDTA in non-deterministic real time, and in $\langle X \rangle$ in deterministic real time.

The theoretical results are accompanied by a POC software artifact, Treetop. Through Treetop, we demonstrate that subtyping can be used to implement useful and powerful metaprogramming applications even in a decidable type system. Treetop converts context-free grammars into $C^{\#}$ subtyping machines and can be used to embed context-free DSLs as fluent APIs; its practical relevance can be argued based on the work of Ferles et al. [2020], who found cases of context-free API protocols in the wild. Comparing Treetop to the fluent API generator of Grigore [2017], we see that both support all CFGs, but Treetop does so in a decidable type system fragment. While Grigore’s generator constructs very large subtyping machines that take a very long time to compile, Treetop generates small machines that compile relatively fast (in some cases). In theory, Treetop’s subtyping machines run in polynomial time. In practice, our experiment found an example that leads to possibly exponential compilation time using the CSC compiler. In addition, Treetop’s fluent API generation does not yet match the qualities of previous practical generators (e.g., [Gil and Roth 2019; Yamazaki et al. 2019]). More research work (but also implementation) is required to support features such as auto-completion, early failure, and, run-time parsing to create an AST.

A note on mixin inheritance. Kennedy and Pierce included mixin inheritance in their abstract type system, allowing classes to inherit their type variables, $\gamma(x) : x$. Mixin inheritance was not discussed in this paper, given that not many programming languages actually support it. As it turns out, most of our theorems hold when this feature is re-introduced to the type system. The construction in the proof of Theorem 5.2 is left unchanged, since it relies on the non-expansive guarantee of inheritance in $\langle C, M \rangle$ and does not depend on the structure of inheritance declarations. In contrast, the proof of Theorem 5.6 breaks when confronted with mixin inheritance. The construction encodes (transitive) inheritance relations $\gamma(x) :^* \sigma(\tau)$ as context-free productions in GNF, $\gamma(x) ::= \sigma(\tau)$. The production corresponding to mixin inheritance is an ε -production, $\gamma(x) ::= x_i$, which does not conform to GNF. In this case the resulting grammar describes a general context-free forest, so the expressiveness of $\langle X, M \rangle$ is now bounded from above by class NCF instead of NCF_{GNF} (which is the lower bound). The computational power of type system $\langle X \rangle$, however, is already bounded from below by general deterministic NCF, so it too stays the same. Therefore, the introduction of mixin inheritance to the subtyping-based type systems in Fig. 1 only adds a loose bound to $\langle X, M \rangle$,

$$NCF_{\text{GNF}} \subseteq \mathcal{L}(\langle X, M \rangle) \subseteq \text{NCF}.$$

REFERENCES

- Ken Arnold, James Gosling, and David Holmes. 2006. *The Java Programming Language* (4 ed.). Addison-Wesley, Upper Saddle River, NJ.
- Henk Barendregt. 1991. Introduction to generalized type systems. *J. Functional Prog.* 1 (1991), 125–154. Issue 2. <https://doi.org/10.1017/s0956796800020025>.
- Norbert Blum and Robert Koch. 1999. Greibach Normal Form Transformation Revisited. *Information and Computation* 150, 1 (1999), 112 – 118. <https://doi.org/10.1006/inco.1998.2772>
- Donald D. Chamberlin and Raymond F. Boyce. 1974. SEQUEL: A Structured English Query Language. In *Proc. ACM SIGFIDET Work. Data Description, Access & Control* (Ann Arbor, Michigan) (SIGFIDET ’74). ACM, NY/USA, 249–264.

<https://doi.org/10.1145/800296.811515>

- H. Comon, M. Dauchet, R. Gilleron, C. Löding, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. 2007. Tree Automata Techniques and App. Available on: <http://tata.gforge.inria.fr/> (accessed Oct. 2020).
- James O. Coplien. 1996. *Curiously Recurring Template Patterns*. SIGS Publications, Inc., USA, 135–144.
- Kostas Ferles, Jon Stephens, and Isil Dillig. 2020. Verifying Correct Usage of Context-Free API Protocols (Extended Version). *arXiv:2010.09652* [cs.PL]
- Martin Fowler. 2005. FluentInterface. <https://www.martinfowler.com/bliki/FluentInterface.html> (accessed August 12, 2020).
- Ferenc Gécseg and Magnus Steinby. 1997. *Tree Languages*. Springer-Verlag, Berlin, Heidelberg, 1–68.
- Joseph Gil and Ori Roth. 2020. Ties between Parametrically Polymorphic Type Systems and Finite Control Automata. *arXiv preprint arXiv:2009.04437* (2020). *arXiv:2009.04437* [cs.PL]
- Yossi Gil and Tomer Levy. 2016. Formal language recognition with the Java type checker. In *30th Europ. Conf OO Prog. (ECOOP 2016) (Leibniz International Proceedings in Inf. (LIPIcs), Vol. 56)*, Shriram Krishnamurthi and Benjamin S. Lerner (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 10:1–10:27. <https://doi.org/10.4230/LIPIcs.ECOOP.2016.10>
- Yossi Gil and Ori Roth. 2019. Fling—a fluent API generator. In *33rd Europ. Conf OO Prog. (ECOOP 2019) (Leibniz International Proceedings in Inf. (LIPIcs), Vol. 134)*, Alastair F. Donaldson (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 13:1–13:25. <https://doi.org/10.4230/LIPIcs.ECOOP.2019.13>
- Ben Greenman, Fabian Muehlboeck, and Ross Tate. 2014. Getting F-Bounded Polymorphism into Shape. *SIGPLAN Not.* 49, 6 (June 2014), 89–99. <https://doi.org/10.1145/2666356.2594308>
- Sheila A. Greibach. 1965. A New Normal-Form Theorem for Context-Free Phrase Structure Grammars. *J. ACM* 12, 1 (Jan. 1965), 42–52. <https://doi.org/10.1145/321250.321254>
- Radu Grigore. 2017. Java Generics Are Turing Complete. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (Paris, France) (POPL 2017)*. Association for Computing Machinery, New York, NY, USA, 73–85. <https://doi.org/10.1145/3009837.3009871>
- Irène Guessarian. 1983. Pushdown tree automata. *Math. Syst. Theory* 16, 1 (1983), 237–263.
- Anders Hejlsberg, Scott Wiltamuth, Peter Golde, and Mads Torgersen. 2003. *The C# Programming Language* (2nd ed.). Addison-Wesley, Reading, MA, USA.
- John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. 2007. *Introduction to automata theory, languages, and computation* (3 ed.). Pearson Addison Wesley, Boston, MA.
- Andrew Kennedy and Benjamin Pierce. 2007. On decidability of nominal subtyping with variance. In *Int. Work. Found. & Devel. OO Lang. (Nice, France) (FOOL/WOOD’07)*. ACM, Nice, France, Article 5, 12 pages. <http://foolwood07.cs.uchicago.edu/program/kennedy-abstract.html>
- Tomoki Nakamaru and Shigeru Chiba. 2020. Generating a generic fluent API in Java. *Art, Sci., & Eng. Prog.* 4, Article 9 (Feb. 2020), 23 pages. Issue 3. <https://doi.org/10.22152/programming-journal.org/2020/4/9>
- Tomoki Nakamaru, Kazuhiro Ichikawa, Tetsuro Yamazaki, and Shigeru Chiba. 2017. Silverchain: a fluent API generator. In *Proc. 16th ACM SIGPLAN Int. Conf Generative Prog. (GPCE’17)*. ACM, Vancouver, BC, Canada, 199–211.
- Tomoki Nakamaru, Tomomasa Matsunaga, Tetsuro Yamazaki, Soramichi Akiyama, and Shigeru Chiba. 2020. An Empirical Study of Method Chaining in Java. In *Proceedings of the 17th International Conference on Mining Software Repositories (Seoul, Republic of Korea) (MSR ’20)*. Association for Computing Machinery, New York, NY, USA, 93–102. <https://doi.org/10.1145/3379597.3387441>
- Martin Odersky, Philippe Altherr, Vincent Cremet, Burak Emir, Sebastian Maneth, Stéphane Micheloud, Nikolay Mihaylov, Michel Schinz, Erik Stenman, and Matthias Zenger. 2004. *An Overview of the Scala Programming Language*. Technical Report IC/2004/64. EPFL Lausanne, Switzerland.
- Johannes Osterholzer. 2018. *New Results on Context-Free Tree Languages*. Ph. D. Dissertation. Technische Universität Dresden, Zellescher Weg 18 01069 Dresden. <https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa-235434>
- Ori Roth. 2021a. Study of the Subtyping Machine of Nominal Subtyping with Variance (full version). *arXiv:2109.03950* [cs.PL]
- Ori Roth. 2021b. Treemap: A Context-Free Fluent API Generator for C#. <https://doi.org/10.5281/zenodo.5091711>
- Leslie G. Valiant. 1975. General context-free recognition in less than cubic time. *J. Comput. System Sci.* 10, 2 (1975), 308 – 315. [https://doi.org/10.1016/S0022-0000\(75\)80046-8](https://doi.org/10.1016/S0022-0000(75)80046-8)
- Mirko Viroli. 2000. *On the recursive generation of parametric types*. Technical Report. DEIS-LIA-00-002, Università di Bologna.
- Hao Xu. 2010. EriLex: an embedded domain specific language generator. In *Objects, Models, Components, Patterns*, Jan Vitek (Ed.). Springer, Berlin, Heidelberg, 192–212.
- Tetsuro Yamazaki, Tomoki Nakamaru, Kazuhiro Ichikawa, and Shigeru Chiba. 2019. Generating a fluent API with syntax checking from an LR grammar. *Proc. ACM Program. Lang.* 3, Article 134 (Oct. 2019), 24 pages. <https://doi.org/10.1145/3360560>