Featherweight Ownership and Immutability Generic Java - Technical Report

Yoav Zibin yoav.zibin@gmail.com

1 Introduction

This technical report contains proofs that were omitted from our paper entitled "Ownership and Immutability in Generic Java". Please read the paper first, and only then proceed to reading this technical report, because this technical report is not self contained. We only include a summary of the syntax (Fig. 1), subtyping rules (Fig. 2), expression typing rules (Fig. 3), and reduction rules (Fig. 4).

We begin with some definitions. A *rule* has the form $\frac{A}{B}$, where A is the *assumption* and B is the *conclusion*. If A is empty, we also call the rule an *axiom*. An *instance* of a rule/assumption/conclusion is any substitution of variables in the rule. A *derivation sequence* is a sequence of elements (each element is an instance of a conclusion), where the assumptions needed for each element appear as previous elements in the sequence.

An expression/type is called *closed* if it does not contain any free variables (such as wildcards, `this`, `I`, `O`, or `This`).

The length of a sequence $\overline{\mathtt{x}}$ is written $\#(\overline{\mathtt{x}})$,

To define *ftype*`(e,f,T)` and *mtype*`(e,f,T)`, we use the auxiliary function *substitute*:

$$\mathit{substitute}(\mathtt{e},\mathtt{C<MO,IP>},\mathtt{T}) = \begin{cases} \mathtt{error} & O(\mathtt{T}) = \mathtt{This} \text{ and } \mathtt{z} = \mathtt{error} \\ [\mathtt{z}/\mathtt{This},\mathtt{MO}/\mathtt{O},\mathtt{IP}/\mathtt{I}]\mathtt{T} & \text{otherwise} \end{cases} \qquad \mathtt{z} = \begin{cases} \mathtt{l} & \mathtt{e} = \mathtt{l} \\ \mathtt{This} & \mathtt{e} = \mathtt{this} \\ \mathtt{error} & \text{otherwise} \end{cases}$$

Formally, *ftype*`(e,f,C<MO,IP>)` = *substitute*$(\mathtt{e},\mathtt{C<MO,IP>},\mathit{ftype}(\mathtt{f},\mathtt{C}))$, and similarly for *mtype*.

Given an expression `e`, we define $K(\mathtt{e})$ to be the set of all ongoing constructors in `e`, i.e., all locations in subexpressions `e;return l`. Formally,

$$K(\mathtt{e}) = \begin{cases} K(\mathtt{e}') \cup \{l\} & \text{if } \mathtt{e} = (\mathtt{e}';\mathtt{return\ l}) \\ K(\mathtt{e}') & \text{if } \mathtt{e} = (\mathtt{e}'.\mathtt{f}) \\ K(\mathtt{e}') \cup K(\mathtt{e"}) & \text{if } \mathtt{e} = (\mathtt{e}'.\mathtt{f=e"}) \\ \cup K(\overline{\mathtt{e}'}) & \text{if } \mathtt{e} = (\mathtt{new\ N}(\overline{\mathtt{e}'})) \\ K(\mathtt{e"}) \cup K(\overline{\mathtt{e}'}) & \text{if } \mathtt{e} = (\mathtt{e".m}(\overline{\mathtt{e}'})) \end{cases}$$

A heap H is *well-typed for* K if it satisfies two conditions: (i) Each field location is a subtype (using K,Γ_H) of the declared field type, i.e., for every location `l`, where $H[\mathtt{l}] = \mathtt{C<NO,NI>}(\overline{\mathtt{v}})$ and $\mathit{fields}(\mathtt{C}) = \overline{\mathtt{f}}$, and for every field $\mathtt{f}_i$, we have that either $\mathtt{v}_i = \mathtt{null}$ or $K,\Gamma_H \vdash \Gamma_H(\mathtt{v}_i) \le \mathit{ftype}(\mathtt{l},\mathtt{f}_i,\mathtt{C<NO,NI>})$. (ii) There is a linear order $\preceq^T$ over $\mathrm{dom}(H)$ such that for every location `l`, $\theta_H(\mathtt{l}) = \mathtt{World}$ or $\theta_H(\mathtt{l}) \prec^T \mathtt{l}$, and $I_H(\mathtt{l}) = \mathtt{Mutable}$ or $\kappa_H(\mathtt{l}) \preceq^T \mathtt{l}$.

2 Subtyping

First we prove some lemmas regarding subtyping.

Lemma 2.1. *If* $K,\Gamma \vdash \mathtt{C<MO,IP>} \le \mathtt{C'<MO',IP'>}$, *then*
(i) $\mathtt{MO'} \neq \mathtt{?} \Rightarrow \mathtt{MO} = \mathtt{MO'}$,
(ii) $(\mathtt{IP'} \neq \mathtt{Immut_l}$ *or* $(\mathtt{IP'} = \mathtt{Immut_l}$ *and* $(\mathtt{l} \not\prec_\theta \mathtt{MO'}$ *or* $\mathtt{l} \notin K)) \Rightarrow K,\Gamma \vdash \mathtt{IP} \le \mathtt{IP'}$,
(iii) `C` *is a subclass of* `C'`,
(iv) $K,\Gamma \vdash \mathtt{D<MO,IP>} \le \mathtt{D<MO',IP'>}$ *for any class* `D`,
(v) $K,\Gamma \vdash \mathtt{D<l,IP>} \le \mathtt{D<l,IP'>}$ *for any class* `D` *and* $\mathtt{MO'} \preceq_\theta \mathtt{l}$.

FT ::= C<FO, IP>	Field (and method return) Type.
T ::= C<MO, IP>	Type.
N ::= C<NO, NI>	Non-variable type (for objects).
NO ::= World | l	Non-variable Owner parameter (for objects).
FO ::= NO | This | O	Field Owner parameter.
MO ::= FO | ?	Method Owner parameter (including generic wildcard).
NI ::= Mutable | $\texttt{Immut}_{\texttt{l}}$	Non-variable Immutability parameter (for objects).
VI ::= NI | Immut | I	Variable Immutability for new.
IP ::= ReadOnly | VI	Immutability Parameter.
IG ::= ReadOnly | Immut | Mutable | Raw	Immutability method Guard.
M ::= <I extends IG>? FT m($\overline{\texttt{T}}$ $\overline{\texttt{x}}$) { return e; }	Method declaration.
L ::= class C<O, I> extends C′<O, I>{ $\overline{\texttt{FT}}$ $\overline{\texttt{f}}$; $\overline{\texttt{M}}$ }	cLass declaration.
v ::= null | l	Values: either null or a location l.
e ::= v | x | e.f | e.f = e | e.m($\overline{\texttt{e}}$) | new C<FO, VI>($\overline{\texttt{e}}$) | e ; return l	Expressions.

Figure 1: FOIGJ Syntax. The terminals are null, owner parameters (O, This, World), and immutability parameters (I, ReadOnly, Mutable, Raw, Immut). Given a location l, $\texttt{Immut}_{\texttt{l}}$ represents an immutable object with cooker l. The program source code cannot contain the grayed elements (locations are only created during execution/reduction in R-New of Fig. 4).

$$\frac{}{K,\Gamma \vdash \texttt{I} \leq \Gamma(\texttt{I})}\ (\text{S1}) \qquad \frac{}{K,\Gamma \vdash \texttt{T} \leq \texttt{T}}\ (\text{S2}) \qquad \frac{K,\Gamma \vdash \texttt{S} \leq \texttt{T} \quad K,\Gamma \vdash \texttt{T} \leq \texttt{U}}{K,\Gamma \vdash \texttt{S} \leq \texttt{U}}\ (\text{S3}) \qquad \frac{\texttt{class C<O, I> extends C}'\texttt{<O, I>}}{K,\Gamma \vdash \texttt{C<MO, IP>} \leq \texttt{C}'\texttt{<MO, IP>}}\ (\text{S4})$$

$$\frac{}{K,\Gamma \vdash \texttt{Mutable} \leq \texttt{Raw}}\ (\text{S5}) \qquad \frac{}{K,\Gamma \vdash \texttt{Raw} \leq \texttt{ReadOnly}}\ (\text{S6}) \qquad \frac{}{K,\Gamma \vdash \texttt{Immut} \leq \texttt{ReadOnly}}\ (\text{S7})$$

$$\frac{K,\Gamma \vdash \texttt{IP} \leq \texttt{IP}'}{K,\Gamma \vdash \texttt{C<MO, IP>} \leq \texttt{C<MO, IP}'\texttt{>}}\ (\text{S8}) \qquad \frac{}{K,\Gamma \vdash \texttt{C<MO, IP>} \leq \texttt{C<?, IP>}}\ (\text{S9}) \qquad \frac{\texttt{l} \in K}{K,\Gamma \vdash \texttt{Immut}_{\texttt{l}} \leq \texttt{Raw}}\ (\text{S10})$$

$$\frac{\texttt{l} \notin K}{K,\Gamma \vdash \texttt{Immut}_{\texttt{l}} \leq \texttt{Immut}}\ (\text{S11}) \qquad \frac{\texttt{l} \notin K}{K,\Gamma \vdash \texttt{Immut} \leq \texttt{Immut}_{\texttt{l}}}\ (\text{S12}) \qquad \frac{\texttt{l} \prec_\theta \texttt{NO}}{K,\Gamma \vdash \texttt{C<NO, Immut>} \leq \texttt{C<NO, Immut}_{\texttt{l}}\texttt{>}}\ (\text{S13})$$

Figure 2: FOIGJ Subtyping Rules ($K,\Gamma \vdash \texttt{T} \leq \texttt{T}'$). Rule S13 shows the connection between cooker l and owner NO.

Proof. In this proof we omit $K,\Gamma \vdash$ because the context K,Γ is clear. First note that due to S13, it is not the case that IP ≤ IP′. Therefore, we need parts (ii) and (iv)–(v), which connects IP and IP′ in other ways.

(i) All subtyping rules maintain the same owner parameter, except S9, thus if MO′ ≠ ?, then the owner must be preserved.

(ii) All subtyping rules maintain that IP ≤ IP′ except S13. If l $\not\prec_\theta$ MO′ then we cannot use S13. If l ∉ K then (from S12) Immut ≤ IP′, and if the proof used S13 then C<MO, IP> ≤ C<MO, Immut>, thus IP ≤ Immut ≤ IP′.

(iii) All rules maintain the same class, and S4 permits subclassing.

(iv) We create a new derivation sequence where instances of rule S4 are deleted, and all occurrences of C are replaced with D.

(v) Similarly to part (iv), we delete S4 and S9, replace C with D, and replace MO and MO′ with l. The only rule where the owner matters is S13:

$$\frac{\texttt{l}' \prec_\theta \texttt{MO}'}{K,\Gamma \vdash \texttt{C<MO}', \texttt{Immut>} \leq \texttt{C<MO}', \texttt{Immut}_{\texttt{l}'}\texttt{>}}$$

and we have that MO′ $\preceq_\theta$ l, therefore l′ $\prec_\theta$ l. □

Lemma 2.2. *If* $K,\Gamma \vdash$ C<MO, IP> ≤ C′<MO, IP′>, *then for any class* D *and any owner parameter* W *such that* W = O \| World, *we have that*

$$K,\Gamma \vdash \texttt{D<}[\texttt{MO/O}]\texttt{W}, [\texttt{IP/I}]\texttt{Z>} \leq \texttt{D<}[\texttt{MO/O}]\texttt{W}, [\texttt{IP}'\texttt{/I}]\texttt{Z>}.$$

Proof. If Z ≠ I then obviously D<[MO/O]W, Z> = D<[MO/O]W, Z>. If Z = I then we need to prove that $K,\Gamma \vdash$ D<[MO/O]W, IP> ≤ D<[MO/O]W, IP′>. Because MO $\preceq_\theta$ World and W = O \| World, then MO $\preceq_\theta$ [MO/O]W. Therefore, we can apply Lem. 2.1 part (v). □

$$\frac{K \cup \{\mathtt{l}\}, \Gamma \vdash \mathtt{e} : \mathtt{T}}{K, \Gamma \vdash \mathtt{e;return\ l} : \Gamma(\mathtt{l})}\ \text{(T-RETURN)} \qquad \frac{\mathit{mtype}(\bot, \mathtt{build}, \mathtt{C<FO,VI>}) = \overline{\mathtt{T}} \to \mathtt{U} \quad K, \Gamma \vdash \overline{\mathtt{e}} : \overline{\mathtt{T'}} \quad K, \Gamma \vdash \overline{\mathtt{T'}} \leq \overline{\mathtt{T}}}{K, \Gamma \vdash \mathtt{new\ C<FO,VI>}(\overline{\mathtt{e}}) : \mathtt{C<FO,VI>}}\ \text{(T-NEW)}$$

$$\frac{}{K, \Gamma \vdash \mathtt{x} : \Gamma(\mathtt{x})}\ \text{(T-VAR)} \qquad \frac{}{K, \Gamma \vdash \mathtt{null} : \mathtt{T}}\ \text{(T-NULL)} \qquad \frac{K, \Gamma \vdash \mathtt{e} : \mathtt{C<MO,IP>} \quad \mathit{ftype}(\mathtt{e}, \mathtt{f}, \mathtt{C<MO,IP>}) = \mathtt{T}}{K, \Gamma \vdash \mathtt{e.f} : \mathtt{T}}\ \text{(T-FIELD-ACCESS)}$$

$$\frac{}{K, \Gamma \vdash \mathtt{l} : \Gamma(\mathtt{l})}\ \text{(T-LOCATION)} \qquad \frac{\begin{array}{c} K, \Gamma \vdash \mathtt{e.f} : \mathtt{T} \quad K, \Gamma \vdash \mathtt{e'} : \mathtt{T'} \quad K, \Gamma \vdash \mathtt{T'} \leq \mathtt{T} \quad K, \Gamma \vdash \mathtt{e} : \mathtt{C<MO,IP>} \\ K, \Gamma \vdash \mathtt{IP} \leq \mathtt{Raw} \quad \mathit{isTransitive}(\mathtt{e}, \Gamma, \mathtt{C<MO,IP>}) \quad \mathtt{MO} \neq \mathtt{?} \end{array}}{K, \Gamma \vdash \mathtt{e.f} = \mathtt{e'} : \mathtt{T'}}\ \text{(T-FIELD-ASSIGNMENT)}$$

$$\frac{\begin{array}{c} K, \Gamma \vdash \mathtt{e_0} : \mathtt{C<MO,IP>} \quad \mathit{mtype}(\mathtt{e_0}, \mathtt{m}, \mathtt{C<MO,IP>}) = \overline{\mathtt{T}} \to \mathtt{T"} \quad K, \Gamma \vdash \overline{\mathtt{e}} : \overline{\mathtt{T'}} \quad K, \Gamma \vdash \overline{\mathtt{T'}} \leq \overline{\mathtt{T}} \quad \mathit{mguard}(\mathtt{m}, \mathtt{C}) = \mathtt{IG} \\ K, \Gamma \vdash \mathtt{IP} \leq \mathtt{IG} \quad \mathtt{IG} = \mathtt{Raw} \Rightarrow \mathit{isTransitive}(\mathtt{e_0}, \Gamma, \mathtt{C<MO,IP>}) \quad \mathit{mtype}(\mathtt{m}, \mathtt{C}) = \overline{\mathtt{U}} \to \mathtt{V} \quad O(\mathtt{T}_i) = \mathtt{?} \Rightarrow O(\mathtt{U}_i) = \mathtt{?} \end{array}}{K, \Gamma \vdash \mathtt{e_0.m}(\overline{\mathtt{e}}) : \mathtt{T"}}\ \text{(T-INVOKE)}$$

Figure 3: FOIGJ Expression Typing Rules ($K, \Gamma \vdash \mathtt{e} : \mathtt{T}$).

$$\frac{\mathtt{l} \notin \mathrm{dom}(H) \qquad \mathtt{VI'} = \begin{cases} \mathtt{Immut_l} & \text{if } \mathtt{VI} = \mathtt{Immut} \text{ or } (\mathtt{VI} = \mathtt{Immut_C} \text{ and } \mathtt{c} \notin K) \\ \mathtt{VI} & \text{otherwise} \end{cases}}{K \vdash H, \mathtt{new\ C<NO,VI>}(\overline{\mathtt{v}}) \to H[\mathtt{l} \mapsto \mathtt{C<NO,VI'>}(\overline{\mathtt{null}})], \mathtt{l.build}(\overline{\mathtt{v}})\mathtt{;return\ l}}\ \text{(R-NEW)}$$

$$\frac{K \cup \{\mathtt{l}\} \vdash H, \mathtt{e} \to H', \mathtt{e'}}{K \vdash H, \mathtt{e;return\ l} \to H', \mathtt{e'}\mathtt{;return\ l}}\ \text{(R-C1)} \qquad \frac{H[\mathtt{l}] = \mathtt{C<NO,NI>}(\overline{\mathtt{v}}) \quad \mathit{fields}(\mathtt{C}) = \overline{\mathtt{f}}}{K \vdash H, \mathtt{l.f}_i \to H, \mathtt{v}_i}\ \text{(R-FIELD-ACCESS)}$$

$$\frac{H[\mathtt{l}] = \mathtt{C<NO,NI>}(\overline{\mathtt{v}}) \quad \mathit{fields}(\mathtt{C}) = \overline{\mathtt{f}} \quad \mathtt{NI} = \mathtt{Mutable} \text{ or } \kappa_H(\mathtt{l}) \in K \quad \mathtt{v'} = \mathtt{null} \text{ or } \mathtt{l} \preceq_\theta \theta_H(\mathtt{v'})}{K \vdash H, \mathtt{l.f}_i = \mathtt{v'} \to H[\mathtt{l} \mapsto \mathtt{C<NO,NI>}([\mathtt{v'}/\mathtt{v}_i]\overline{\mathtt{v}})], \mathtt{v'}}\ \text{(R-FIELD-ASSIGNMENT)}$$

$$\frac{}{K \vdash H, \mathtt{v;return\ l} \to H, \mathtt{l}}\ \text{(R-RETURN)} \qquad \frac{H[\mathtt{l}] = \mathtt{C<NO,NI>}(\ldots) \quad \mathit{mbody}(\mathtt{m}, \mathtt{C}) = \overline{\mathtt{x}}.\mathtt{e'}}{K \vdash H, \mathtt{l.m}(\overline{\mathtt{v}}) \to H, [\overline{\mathtt{v}}/\overline{\mathtt{x}}, \mathtt{l}/\mathtt{this}, \mathtt{l}/\mathtt{This}, \mathtt{NO}/\mathtt{O}, \mathtt{NI}/\mathtt{I}]\mathtt{e'}}\ \text{(R-INVOKE)}$$

Figure 4: FOIGJ Reduction Rules ($K \vdash H, \mathtt{e} \to H', \mathtt{e'}$), excluding all congruence rules except R-C1.

Lemma 2.3. *If* $K, \Gamma \vdash \mathtt{C<MO,IP>} \leq \mathtt{C'<MO,IP'>}$*, both types are closed,* $\mathit{isTransitive}(\bot, \Gamma, \mathtt{C'<MO,IP'>})$ *and* $\mathtt{IP'} \leq \mathtt{Raw}$*, then* $\mathtt{IP} = \mathtt{IP'}$.

Proof. If $\mathtt{IP'} = \mathtt{Mutable}$ then obviously $\mathtt{IP} = \mathtt{Mutable}$. Otherwise $\mathtt{IP'} = \mathtt{Immut_l}$ and $\mathtt{l} \in K$ (because $\mathtt{IP'} \leq \mathtt{Raw}$). From definition of $\mathit{isTransitive}(\bot, \Gamma, \mathtt{C'<MO,IP'>})$, $\mathtt{l} \not\preceq_\theta \mathtt{MO}$, thus S13 cannot be applied, and that is the only applicable rule where $\mathtt{Immut_l}$ appears as a supertype (because S12 cannot be applied since $\mathtt{l} \in K$), thus $\mathtt{IP} = \mathtt{IP'}$. □

The next lemma shows that if $\mathtt{e'}$ was reduced to $\mathtt{e}$ (therefore the type of $\mathtt{e}$ is a subtype of $\mathtt{e'}$), then $\mathtt{e}$ can call any method that $\mathtt{e'}$ could. Phrased differently, if $\mathtt{e'}$ satisfies a method's guard then $\mathtt{e}$ would as well.

Lemma 2.4. *If* $K, \Gamma \vdash \mathtt{C<MO,IP>} \leq \mathtt{C'<MO,IP'>}$*, and both types are closed, then*
(i) $K, \Gamma \vdash \mathtt{IP'} \leq \mathtt{Mutable} \Rightarrow K, \Gamma \vdash \mathtt{IP} \leq \mathtt{Mutable}$,
(ii) $K, \Gamma \vdash \mathtt{IP'} \leq \mathtt{Immut} \Rightarrow K, \Gamma \vdash \mathtt{IP} \leq \mathtt{Immut}$,
(iii) $\mathit{isTransitive}(\bot, \Gamma, \mathtt{C'<MO,IP'>})$ *and* $K, \Gamma \vdash \mathtt{IP'} \leq \mathtt{Raw} \Rightarrow K, \Gamma \vdash \mathtt{IP} \leq \mathtt{Raw}$,
(iv) if $\mathtt{IG} = \mathtt{Raw} \Rightarrow \mathit{isTransitive}(\bot, \Gamma, \mathtt{C'<MO,IP'>})$*, where* $\mathtt{IG} = \mathtt{ReadOnly} \mid \mathtt{Immut} \mid \mathtt{Mutable} \mid \mathtt{Raw}$*, then* $K, \Gamma \vdash \mathtt{IP'} \leq \mathtt{IG} \Rightarrow K, \Gamma \vdash \mathtt{IP} \leq \mathtt{IG}$.

Proof. (i) Trivial because we must have that $\mathtt{IP'} = \mathtt{IP} = \mathtt{Mutable}$ or ($\mathtt{IP'} = \mathtt{IP} = \mathtt{I}$ and $\mathtt{I} : \mathtt{Mutable} \in \Gamma$).

(ii) If $\mathtt{IP'} \neq \mathtt{Immut_l}$ then from Lem. 2.1 part (ii), we have $K, \Gamma \vdash \mathtt{IP} \leq \mathtt{IP'}$, and from transitivity $\mathtt{IP} \leq \mathtt{IP'} \leq \mathtt{Immut} \Rightarrow \mathtt{IP} \leq \mathtt{Immut}$. Otherwise, $\mathtt{IP'} = \mathtt{Immut_l}$, and because $\mathtt{IP'} \leq \mathtt{Immut}$, from S11 we must have that $\mathtt{l} \notin K$, and from Lem. 2.1 part (ii), we proved $K, \Gamma \vdash \mathtt{IP} \leq \mathtt{IP'}$,

(iii) From Lem. 2.3 we know that $\mathtt{IP} = \mathtt{IP'}$ thus $\mathtt{IP} \leq \mathtt{Raw}$.

(iv) Stems from parts (vi)–(viii) and the fact that for any $\mathtt{IP} \leq \mathtt{ReadOnly}$. □

If the cooker is not inside the owner, then subtypes are not over-approximation (i.e., they preserve the same cooker).

Lemma 2.5. *If* $K,\Gamma \vdash \texttt{C<MO,IP>} \leq \texttt{C'<NO,Immut}_\texttt{l}\texttt{>}$, $\texttt{l} \not\prec_\theta \texttt{NO}$, *and* $\texttt{l} \in K$, *then* $\texttt{IP} = \texttt{Immut}_\texttt{l}$.

Proof. The only subtyping rule where the supertype has a cooker $\texttt{Immut}_\texttt{l}$ are rules S12 and S13, and they can't be applied because we assumed that $\texttt{l} \not\prec_\theta \texttt{NO}$ and $\texttt{l} \notin K$. Thus only the reflexivity rule can be applied. Note that rule $K,\Gamma \vdash \texttt{I} \leq \Gamma(\texttt{I})$ cannot be applied because we assume that $\Gamma(\texttt{I})$ is a method guard IG, and $\texttt{IG} \neq \texttt{Immut}_\texttt{l}$ because according to our syntax

$$\texttt{IG} = \texttt{ReadOnly} \mid \texttt{Immut} \mid \texttt{Mutable} \mid \texttt{Raw}.$$

□

Next we prove a substitution lemma: substituting I, O, or This, does not change the subtype relation.

Lemma 2.6. *If* $K,\Gamma \vdash \texttt{T} \leq \texttt{T}'$ *then for every* $\texttt{IP},\texttt{MO},\texttt{MO}'$ *such that* $\texttt{IP} \leq \Gamma(\texttt{I})$, *we have that*

$$K,\Gamma \vdash [\texttt{IP}/\texttt{I},\texttt{MO}/\texttt{O},\texttt{MO}'/\texttt{This}](\texttt{T} \leq \texttt{T}').$$

Proof. Let S denote the derivation sequence for $K,\Gamma \vdash \texttt{T} \leq \texttt{T}'$, and S_I for $\texttt{IP} \leq \Gamma(\texttt{I})$. Let S' be a new sequence in which we do the substitution $[\texttt{IP}/\texttt{I},\texttt{MO}/\texttt{O},\texttt{MO}'/\texttt{This}]$ on every element in S, and let S" be the sequence starting with S_I followed by S'. The last element in S is $K,\Gamma \vdash \texttt{T} \leq \texttt{T}'$, therefore the last element in S' and S" is what we need to prove: $K,\Gamma \vdash [\texttt{IP}/\texttt{I},\texttt{MO}/\texttt{O},\texttt{MO}'/\texttt{This}](\texttt{T} \leq \texttt{T}')$. Next we show that S" is a legal derivation sequence by showing that each element is a legal consequence of previous elements (by induction on the size of S"). Elements from S_I are of course legal. By induction we proved the first $n-1$ elements are legal, and we now prove that element n is legal (i.e., a legal consequence of previous elements). Let the corresponding element in S be $\texttt{U} \leq \texttt{U}'$, and element n is $[\texttt{IP}/\texttt{I},\texttt{MO}/\texttt{O},\texttt{MO}'/\texttt{This}](\texttt{U} \leq \texttt{U}')$. (i) If the element is an instance of rules S5–S7 or S10–S12, then substitution did not change the element. (ii) If the element is an instance of rule S1 then we replaced it with $[\texttt{IP}/\texttt{I}](\texttt{I} \leq \Gamma(\texttt{I})) = \texttt{IP} \leq \Gamma(\texttt{I})$, which is the last element of S_I. (iii) If the element is an instance of any other rule, then a simple application of the induction hypothesis proves the element is legal. □

We now prove that substitution preserves subtyping.

Lemma 2.7. *If* $K,\Gamma \vdash \texttt{T} \leq \texttt{T}'$, T *and* $\texttt{T}'$ *are closed, then for (*$\texttt{e} = \bot$ *and* $O(\texttt{T "}) \neq \texttt{This}$*) or* $\texttt{e} = \texttt{l}$, $O(\texttt{T}) \preceq_\theta \texttt{l}$, *we have that:*
(i) $K,\Gamma \vdash$ *substitute*$(\texttt{e},\texttt{T},\texttt{T"}) \leq$ *substitute*$(\texttt{e},\texttt{T}',\texttt{T"})$,
(ii) $K,\Gamma \vdash$ *ftype*$(\texttt{e},\texttt{f},\texttt{T}) \leq$ *ftype*$(\texttt{e},\texttt{f},\texttt{T}')$,
(iii) let *mtype*$(\texttt{e},\texttt{f},\texttt{T}) = \overline{\texttt{T}} \to \texttt{T}_0$ *and* *ftype*$(\texttt{e},\texttt{f},\texttt{T}') = \overline{\texttt{T}'} \to \texttt{T}'_0$, *then* $K,\Gamma \vdash \texttt{T}_i \leq \texttt{T}'_i$ *for* $i = 0,\ldots,\#(\overline{\texttt{T}})$.

Proof. From Lem. 2.1 part (i) and the fact that T and $\texttt{T}'$ are closed (no wildcards), then $O(\texttt{T}) = O(\texttt{T}')$. Let $\texttt{T} = \texttt{C<MO,IP>}$ and $\texttt{T}' = \texttt{C'<MO,IP'>}$.

Recall that *ftype*$\texttt{(e,f,C<MO,IP>)} =$ *substitute*$(\texttt{e},\texttt{C<MO,IP>},$*ftype*$(\texttt{f},\texttt{C}))$, and similarly for *mtype*. Therefore, parts (ii) and (iii) follow from part (i), and the fact that *ftype*$(\texttt{f},\texttt{C}) =$ *ftype*$(\texttt{f},\texttt{C}')$ and *mtype*$(\texttt{m},\texttt{C}) =$ *mtype*$(\texttt{m},\texttt{C}')$ (i.e., subclassing cannot change method signature or field type).

We now prove part (i). From the definition of *substitute*, and because ($\texttt{e} = \bot$ and $O(\texttt{T "}) \neq \texttt{This}$) or $\texttt{e} = \texttt{l}$, we have that *substitute*$(\texttt{e},\texttt{T},\texttt{T"}) = [\texttt{e}/\texttt{This},\texttt{MO}/\texttt{O},\texttt{IP}/\texttt{I}]\texttt{T"}$.

Let $\texttt{T"} = \texttt{D<MO",IP">}$. Because $K,\Gamma \vdash \texttt{T} \leq \texttt{T}'$, from Lem. 2.1 part (v),

$$K,\Gamma \vdash \texttt{D<l',IP>} \leq \texttt{D<l',IP'>} \text{ if } \texttt{MO} \preceq_\theta \texttt{l}' \tag{1}$$

(note that $\texttt{l}'$ can be MO or World). We need to prove that $K,\Gamma \vdash$ *substitute*$(\texttt{e},\texttt{T},\texttt{T"}) \leq$ *substitute*$(\texttt{e},\texttt{T}',\texttt{T"})$, i.e., $K,\Gamma \vdash [\texttt{e}/\texttt{This},\texttt{MO}/\texttt{O},\texttt{IP}/\texttt{I}]\texttt{T"} \leq [\texttt{e}/\texttt{This},\texttt{MO}/\texttt{O},\texttt{IP'}/\texttt{I}]\texttt{T"}$.

If $\texttt{IP"} \neq \texttt{I}$ then from reflexivity $K,\Gamma \vdash [\texttt{e}/\texttt{This},\texttt{MO}/\texttt{O}]\texttt{T"} \leq [\texttt{e}/\texttt{This},\texttt{MO}/\texttt{O}]\texttt{T"}$.

Consider now the case that $\texttt{IP"} = \texttt{I}$. We need to show that

$$K,\Gamma \vdash \texttt{D<}[\texttt{e}/\texttt{This},\texttt{MO}/\texttt{O}]\texttt{MO",IP>} \leq \texttt{D<}[\texttt{e}/\texttt{This},\texttt{MO}/\texttt{O}]\texttt{MO",IP'>} \tag{2}$$

Because $\texttt{MO"} = \texttt{O} \mid \texttt{World} \mid \texttt{This}$ and $\texttt{e} = \bot$ or $\texttt{e} = \texttt{l}$, we have that

$$\texttt{MO} \preceq_\theta [\texttt{e}/\texttt{This},\texttt{MO}/\texttt{O}]\texttt{MO"} \tag{3}$$

From (1) and (3), we proved (2).

□

Lemma 2.8. *If* $K,\Gamma \vdash$ `C<MO,IP>` $\leq$ `C'<MO,IP'>`*, and both types are closed, and*

$$\begin{array}{l} \mathit{mtype}(\bot,\mathtt{m},\mathtt{C'\texttt{<}MO,IP'\texttt{>}}) = \overline{\mathtt{T'}} \to \mathtt{U} \\ \mathit{mtype}(\bot,\mathtt{m},\mathtt{C\texttt{<}MO,IP\texttt{>}}) = \overline{\mathtt{T}} \to \mathtt{U} \\ \mathit{mguard}(\mathtt{m},\mathtt{C'}) = \mathtt{IG} \\ K,\Gamma \vdash \mathtt{IP'} \leq \mathtt{IG} \\ \mathtt{IG} = \mathtt{Raw} \Rightarrow \mathit{isTransitive}(\bot,\Gamma,\mathtt{C'\texttt{<}MO,IP'\texttt{>}}) \end{array}$$

then (i) $K,\Gamma \vdash \mathtt{T}_i \leq \mathtt{T'}_i$ *and* $K,\Gamma \vdash \mathtt{T'}_i \leq \mathtt{T}_i$*, and (ii) for any type* `T"`*, if* $K,\Gamma \vdash$ `T"` $\leq \mathtt{T'}_i$ *then* $K,\Gamma \vdash$ `T"` $\leq \mathtt{T}_i$*.*

Proof. Part (i) implies that $\mathtt{T'}_i$ and $\mathtt{T}_i$ are equivalent. Part (ii) follows directly from part (i) using transitivity rule S3.

Let $\mathit{mtype}(\mathtt{m},C) = \mathit{mtype}(\mathtt{m},C') = \overline{\mathtt{FT}} \to \mathtt{V}$, $\mathtt{FT}_i = \mathtt{D\texttt{<}FO,IP"\texttt{>}}$. Note that $\mathtt{T'}_i = [\mathtt{MO}/\mathtt{O},\mathtt{IP'}/\mathtt{I}]\mathtt{FT}_i$ and $\mathtt{T}_i = [\mathtt{MO}/\mathtt{O},\mathtt{IP}/\mathtt{I}]\mathtt{FT}_i$. Therefore, if `IP"` $\neq$ `I` then $\mathtt{T'}_i = \mathtt{T}_i$, qed.

Thus, `IP"` $=$ `I`, $\mathtt{T'}_i =$ `D<MO",IP'>`, $\mathtt{T}_i =$ `D<MO",IP>`, where `MO` $\preceq_\theta$ `MO"` (because `FO` is either `O` or `World`, but it cannot be `This`). From Lem. 2.7 part (iii), $\mathtt{T}_i \leq \mathtt{T'}_i$.

If `IG` $=$ `ReadOnly`, then FOIGJ ensures that `I` does not appear in method parameters, i.e., we must have that `IP"` $\neq$ `I`. If `IG` $=$ `Mutable`, then `IP'` $=$ `IP` $=$ `Mutable` (because `IP'` $\leq$ `IG`), thus $\mathtt{T'}_i = \mathtt{T}_i$,. If `IG` $=$ `Immut`, then `IP'` $\leq$ `Immut` (because `IP'` $\leq$ `IG`), thus from Lem. 2.4 part (ii), `IP` $\leq$ `Immut`, i.e., `IP` $\leq$ `IP'` $\leq$ `IP`. If `IG` $=$ `Raw`, then *isTransitive*$(\bot,\Gamma,$`C'<MO,IP'>`$)$, and `IP'` $\leq$ `Raw`, thus from Lem. 2.3 we have that `IP'` $=$ `IP`. □

3 Typing

We next prove that a closed expression has a closed type.

Lemma 3.1. *If* $K,\Gamma \vdash$ `e"` : `T"` *and* `e"` *is closed and* `e"` $\neq$ `null`*, then* `T"` *is closed.*

Proof. Note that `null` can have any type `T"` (even an open type) according to rule T-NULL, therefore we require that `e"` $\neq$ `null`.

We prove by induction on the structure of `e"`.

Value `e"` $=$ `v` Because `e"` $\neq$ `null`, then `v` $=$ `l`, and the type of a location is always closed `C<NO,NI>`.

Value `e"` $=$ `e;return l` Similarly, the type of a location is always closed.

Method parameter `e"` $=$ `x` We assumed `e"` is closed, thus it does not contain parameters.

Object creation `e"` $=$ `new C<FO,VI>(`$\overline{\mathtt{e}}$`)` From T-NEW, `T"` $=$ `C<FO,VI>`, and because `e"` is closed then `T"` must be closed.

Field access `e"` $=$ `e.f` Because $K,\Gamma \vdash$ `e"` : `T"`, from T-FIELD-ACCESS we have that `T"` $=$ `T` and

$$K,\Gamma \vdash \mathtt{e} : \mathtt{C\texttt{<}MO,IP\texttt{>}} \qquad \mathit{ftype}(\mathtt{e},\mathtt{f},\mathtt{C\texttt{<}MO,IP\texttt{>}}) = \mathtt{T}$$

By induction, `C<MO,IP>` is closed. Therefore `T` does not contains `O` nor `I`. Next we show it does not contain `This`. Because *ftype* did not return `error`, then either the field did not contain `This`, or it was substituted. Because `e` $\neq$ `this` then `e` $=$ `l`, and `This` was substituted with `l`.

Field assignment `e"` $=$ `e.f` $=$ `e'` Because $K,\Gamma \vdash$ `e"` : `T"`, from T-FIELD-ASSIGNMENT we have that `T"` $=$ `T'` and $K,\Gamma \vdash$ `e'` : `T'` By induction, `T'` is closed.

Method invocation `e"` $=$ `e`$_0$`.m(`$\overline{\mathtt{e}}$`)` Because $K,\Gamma \vdash$ `e"` : `T"`, from T-INVOKE we have that

$$K,\Gamma \vdash \mathtt{e}_0 : \mathtt{C\texttt{<}MO,IP\texttt{>}} \qquad \mathit{mtype}(\mathtt{e}_0,\mathtt{m},\mathtt{C\texttt{<}MO,IP\texttt{>}}) = \overline{\mathtt{T}} \to \mathtt{T"}$$

By induction `C<MO,IP>` is closed, and similar reasoning to field access concludes that `T"` is closed.

□

4 Heap

We now prove that if the heap is well-typed for K, then it is well-typed for any subset of K.

Lemma 4.1. *Given a heap H that is well-typed for K, then for any $S \subseteq K$, the heap H is well-typed for S.*

Proof. This is not trivial, because decreasing K changes the subtyping relation by turning raw objects into immutable. Recall that a well-typed heap H satisfies: (i) there is a linear order $\preceq^T$ over $\mathrm{dom}(H)$ such that for every location $\mathtt{l}$, $\theta(\mathtt{l}) = \mathtt{World}$ or $\theta(\mathtt{l}) \prec^T \mathtt{l}$, and $I(\mathtt{l}) = \mathtt{Mutable}$ or $\kappa(\mathtt{l}) \preceq^T \mathtt{l}$, and (ii) each non-null field location is a subtype of the declared field type.

First, note that the same linear order $\preceq^T$ satisfies (i) for H is well-typed for S. Suppose to the contrary that H is not well-typed for S, i.e., there is some field location $\mathtt{l.f}$ of type $\mathtt{T}$ that points to an object $\mathtt{o}$ of type $\mathtt{T"}$, and $S, \Gamma_H \not\vdash \mathtt{T"} \leq \mathtt{T}$. Because H is well-typed for K, we have that $K, \Gamma_H \vdash \mathtt{T"} \leq \mathtt{T}$. Consider the derivation sequence for $K, \Gamma_H \vdash \mathtt{T"} \leq \mathtt{T}$. We will take this sequence and transform it into a proof that $S, \Gamma_H \vdash \mathtt{T"} \leq \mathtt{T}$, which will lead to contradiction. Specifically, we replace every usage of rule S10 ($K, \Gamma_H \vdash \mathtt{Immut_l} \leq \mathtt{Raw}$) with a proof that $K, \Gamma_H \vdash \mathtt{Immut_l} \leq \mathtt{ReadOnly}$ (using either S10 or S11). We first claim that this is a valid sequence in S, Γ_H: rules S1–S9 and S13 do not use K and therefore are identical, rule S10 was removed, and rules S11–S12 is valid because $S \subset K$. We now claim that these sequence proves that $S, \Gamma_H \vdash \mathtt{T"} \leq \mathtt{T}$, i.e., that each element in the sequence has previous elements that fulfill the assumptions of the rule. Consider an element we removed $\mathtt{Immut_l} \leq \mathtt{Raw}$. Note that $\mathtt{Raw}$ does not appear in $\mathtt{T} = \mathtt{C<MO,IP>}$ because it can only appear in a method guard. The only rule where $\mathtt{Raw}$ appears as a subtype is S6, and by using transitivity (S3), we have that the only conclusion is that $\mathtt{Immut_l} \leq \mathtt{ReadOnly}$, and we added that conclusion. □

Next we prove that owner-as-dominator holds in any well-typed heap.

Lemma 4.2. *If heap H is well-typed for K, then for every location $\mathtt{l} \in \mathrm{dom}(H)$, $\mathtt{l} \mapsto \mathtt{C<NO,NI>}(\overline{\mathtt{v}})$, then either $\mathtt{v}_i = \mathtt{null}$ or $\mathtt{l} \preceq_\theta \theta(\mathtt{v}_i)$.*

Proof. Recall that heap H is *well-typed for K* if it satisfies two conditions: (i) Each field location is a subtype (using K, Γ_H) of the declared field type, i.e., for every location $\mathtt{l}$, where $H[\mathtt{l}] = \mathtt{C<NO,NI>}(\overline{\mathtt{v}})$ and $\mathit{fields}(\mathtt{C}) = \overline{\mathtt{f}}$, and for every field $\mathtt{f}_i$, we have that either $\mathtt{v}_i = \mathtt{null}$ or $K, \Gamma_H \vdash \Gamma_H(\mathtt{v}_i) \leq \mathit{ftype}(\mathtt{l}, \mathtt{f}_i, \mathtt{C<NO,NI>})$. (ii) There is a linear order $\preceq^T$ over $\mathrm{dom}(H)$ such that for every location $\mathtt{l}$, $\theta_H(\mathtt{l}) = \mathtt{World}$ or $\theta_H(\mathtt{l}) \prec^T \mathtt{l}$, and $I_H(\mathtt{l}) = \mathtt{Mutable}$ or $\kappa_H(\mathtt{l}) \preceq^T \mathtt{l}$. From part (ii), we have that the relation $\preceq_\theta$ is a tree order.

Consider some $\mathtt{v}_i \neq \mathtt{null}$, and we will prove that $\mathtt{l} \preceq_\theta \theta(\mathtt{v}_i)$. Let

$$
\begin{aligned}
H[\mathtt{v}_i] &= \mathtt{C'<NO',NI'>(...)} \\
\mathit{fields}(\mathtt{C}) &= \overline{\mathtt{f}} \\
\mathit{ftype}(\mathtt{f}_i, \mathtt{C}) &= \mathtt{C"<FO,IP>} \\
\mathit{ftype}(\mathtt{l}, \mathtt{f}_i, \mathtt{C<NO,NI>}) &= \mathtt{C"<NO",NI">}
\end{aligned}
$$

We need to prove that $\mathtt{l} \preceq_\theta \mathtt{NO'}$. Because the heap is well-typed for K, then $K, \Gamma_H \vdash \mathtt{C'<NO',NI'>} \leq \mathtt{C"<NO",NI">}$. From Lem. 2.1 part (i), we have that $\mathtt{NO'} = \mathtt{NO"}$. According to the syntax, $\mathtt{FO} = \mathtt{World} \mid \mathtt{This} \mid \mathtt{O}$ (the owner of a field cannot be $\mathtt{l}$ of course because the class declarations cannot use locations). By definition of $\mathit{ftype}(\mathtt{l}, \mathtt{f}_i, \mathtt{C<NO,NI>})$, then $\mathtt{NO"} = \mathtt{World} \mid \mathtt{l} \mid \theta(\mathtt{l})$, respectively. Thus we proved that

$$\mathtt{NO'} = \mathtt{NO"} = \mathtt{World} \mid \mathtt{l} \mid \theta(\mathtt{l})$$

Therefore, because $\mathtt{l} \preceq_\theta \mathtt{World}$ and $\mathtt{l} \preceq_\theta \mathtt{l}$ and $\mathtt{l} \preceq_\theta \theta(\mathtt{l})$, we proved that $\mathtt{l} \preceq_\theta \mathtt{NO'}$. □

5 Reduction

We consider only expressions that when reduced using the erased operational semantics, do not result in *null-pointer exceptions*. Null-pointer exceptions can be handled by adding special reduction rules that return `error`, but we prefer to leave the reduction process "stuck".

First we prove that a closed expression reduces in one step to another closed expression.

Lemma 5.1. *If $\mathtt{e}$ is closed and $K \vdash H, \mathtt{e} \to H', \mathtt{e}'$, then $\mathtt{e}'$ is closed.*

Proof. Rules R-NEW, R-FIELD-ACCESS, and R-FIELD-ASSIGNMENT result in a value, which is closed. Rule R-INVOKE results in an expression, but all free variables $\overline{\texttt{x}}, \texttt{this}, \texttt{This}, \texttt{0}, \texttt{I}$ are substituted with locations, `Immut`, $\texttt{Immut}_\texttt{l}$, or `Mutable`. The proof of the congruence rules uses the induction hypothesis. □

Lemma 5.2. *If $K \vdash H, \texttt{e} \rightarrow H$", e" then (i) $\Gamma_H \subseteq \Gamma_{H"}$, (ii) $K, \Gamma_H \vdash \texttt{e}' : \texttt{T}' \Rightarrow K, \Gamma_{H'} \vdash \texttt{e}' : \texttt{T}'$, and (iii) $K, \Gamma_H \vdash \texttt{T} \leq \texttt{T}' \Rightarrow K, \Gamma_{H'} \vdash \texttt{T} \leq \texttt{T}'$.*

Proof. Trivial. (i) None of the reduction rules removes locations or changes the type of a location, therefore H" only includes additional locations. Parts (ii) and (iii) are trivial from (i). □

Theorem 5.3. ***(Progress and Preservation)*** *For every closed expression $\texttt{e} \neq \texttt{v}$, K, and H, if $K, \Gamma_H \vdash \texttt{e} : \texttt{T}$ and H is well-typed for $K \cup K(\texttt{e})$, then there exists $H', \texttt{e}', \texttt{T}'$ such that (i) $K \vdash H, \texttt{e} \rightarrow H', \texttt{e}'$, (ii) $K, \Gamma_{H'} \vdash \texttt{e}' : \texttt{T}'$, and $K, \Gamma_{H'} \vdash \texttt{T}' \leq \texttt{T}$, (iii) H' is well-typed for $K \cup K(\texttt{e}')$, (iv) `T`, `T'`, and $\texttt{e}'$ are closed.*

Proof. Part (iv) is proved from Lem. 5.1 (we know that è is closed) and from Lem. 3.1 (we know that T" and $\grave{\texttt{T}}$ are closed).

We assume that there are no null-pointer exceptions, i.e., that for field access, assignment and method invocation, the receiver is never `null`.

It is easy to examine the reduction rules and verify there is always at most one applicable reduction rule. We will split the proof into three stages: (i) **progress:** there is exactly one applicable reduction rule (Lem. 5.4), (ii) **preservation:** $K, \Gamma_{H'} \vdash \grave{\texttt{e}} : \grave{\texttt{T}}$ and $K, \Gamma_{H'} \vdash \grave{\texttt{T}} \leq \texttt{T}$" (Lem. 5.6), and (iii) H' is well-typed for $K \cup K(\texttt{e}')$ (Lem. 5.7). □

Lemma 5.4. ***(Progress)*** *For every closed expression $\texttt{e}" \neq \texttt{v}$, H, and K, if $K, \Gamma_H \vdash \texttt{e}" : \texttt{T}"$ and and H is well-typed for $K \cup K(\texttt{e}")$, then there exists $H', \grave{\texttt{e}}$ such that $K \vdash H, \texttt{e}" \rightarrow H', \grave{\texttt{e}}$.*

Proof. We prove by examining the structure of e". Because it is closed and not a value, then according to our syntax, it must have one of the following forms:

$$\texttt{e.f} \mid \texttt{e.f} = \texttt{e} \mid \texttt{e.m}(\overline{\texttt{e}}) \mid \texttt{new C<FO,VI>}(\overline{\texttt{e}}) \mid \texttt{e;return l}$$

If the subexpressions are not all values, then we can always apply (exactly) one of the congruence rules. For example, if $\texttt{e}" = (\texttt{e;return l})$ and e is not a value, then by induction we can apply R-C1.

Therefore, e" has one of the following forms:

$$\texttt{v.f} \mid \texttt{v.f} = \texttt{v} \mid \texttt{v.m}(\overline{\texttt{v}}) \mid \texttt{new C<FO,VI>}(\overline{\texttt{v}}) \mid \texttt{v;return l}$$

Because we assumed we do not have *null pointer exceptions* (in field access, assignment or method invocation), then e" has one of the following forms:

$$\texttt{l.f} \mid \texttt{l.f} = \texttt{v} \mid \texttt{l.m}(\overline{\texttt{v}}) \mid \texttt{new C<FO,VI>}(\overline{\texttt{v}}) \mid \texttt{v;return l}$$

We will next examine the matching five reduction rules (R-FIELD-ACCESS, R-FIELD-ASSIGNMENT, R-INVOKE, R-NEW, R-RETURN) and show that their assumptions hold.

Rule R-FIELD-ACCESS

$$\frac{H[\texttt{l}] = \texttt{C<NO,NI>}(\overline{\texttt{v}}) \qquad \mathit{fields}(\texttt{C}) = \overline{\texttt{f}}}{K \vdash H, \texttt{l.f}_i \rightarrow H, \texttt{v}_i}$$

We assumed that $K, \Gamma_H \vdash \texttt{l.f}_i : \texttt{T}"$, therefore from T-FIELD-ACCESS:

$$K, \Gamma_H \vdash \texttt{l} : \texttt{C<MO,IP>} \qquad \mathit{ftype}(\texttt{l}, \texttt{f}_i, \texttt{C<MO,IP>}) = \texttt{T}"$$

From the definitions of *ftype* and *fields*, we know that $\texttt{f}_i \in \mathit{fields}(\texttt{C})$.

Rule R-FIELD-ASSIGNMENT

$$\frac{H[\texttt{l}] = \texttt{C<NO,NI>}(\overline{\texttt{v}}) \qquad \mathit{fields}(\texttt{C}) = \overline{\texttt{f}} \qquad \texttt{NI} = \texttt{Mutable} \text{ or } \kappa(\texttt{l}) \in K \qquad \texttt{v}' = \texttt{null} \text{ or } \texttt{l} \preceq_\theta \theta(\texttt{v}')}{K \vdash H, \texttt{l.f}_i = \texttt{v}' \rightarrow H[\texttt{l} \mapsto \texttt{C<NO,NI>}([\texttt{v}'/\texttt{v}_i]\overline{\texttt{v}})], \texttt{v}'}$$

Because H is well-typed for K, then from Lem. 4.2, we have that $\texttt{v}' = \texttt{null}$ or $\texttt{l} \preceq_\theta \theta(\texttt{v}')$.

We assumed that $K, \Gamma_H \vdash \texttt{l.f}_i = \texttt{v}' : \texttt{T}'$, therefore from T-Field-Assignment:

$$K, \Gamma_H \vdash \texttt{l.f}_i : \texttt{T} \qquad K, \Gamma_H \vdash \texttt{l} : \texttt{C<NO,NI>} \qquad K, \Gamma_H \vdash \texttt{NI} \leq \texttt{Raw}$$

Similarly to field access, because $K, \Gamma_H \vdash \texttt{l.f}_i : \texttt{T}$, then $\texttt{f}_i \in \mathit{fields}(\texttt{C})$. From our syntax NI is either Mutable or $\texttt{Immut}_{\texttt{l}'}$. We want to show that $\texttt{NI} = \texttt{Mutable}$ or $\kappa(\texttt{l}) \in K$. Therefore we need to show that if $\texttt{NI} = \texttt{Immut}_{\texttt{l}'}$ then $\texttt{l}' \in K$. Because $K, \Gamma_H \vdash \texttt{NI} \leq \texttt{Raw}$, it must be from S10 and therefore $\texttt{l}' \in K$.

Rule R-Invoke

$$\frac{H[\texttt{l}] = \texttt{C<NO,NI>(...)} \qquad \mathit{mbody}(\texttt{m}, \texttt{C}) = \overline{\texttt{x}}.\texttt{e}'}{K \vdash H, \texttt{l.m}(\overline{\texttt{v}}) \rightarrow H, [\overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}, \texttt{l}/\texttt{This}, \texttt{NO}/\texttt{O}, \texttt{NI}/\texttt{I}]\texttt{e}'}$$

We assumed that $K, \Gamma_H \vdash \texttt{l.m}(\overline{\texttt{v}}) : \texttt{T}$", therefore from T-Invoke we know that

$$\mathit{mtype}(\texttt{l}, \texttt{m}, \texttt{C<NO,NI>}) = \overline{\texttt{T}} \rightarrow \texttt{T}"$$

Therefore $\mathit{mbody}(\texttt{m}, \texttt{C})$ is defined.

Rule R-New

$$\frac{\texttt{l} \notin \mathrm{dom}(H) \qquad \texttt{VI}' = \begin{cases} \texttt{Immut}_{\texttt{l}} & \text{if } \texttt{VI} = \texttt{Immut} \text{ or } (\texttt{VI} = \texttt{Immut}_{\texttt{C}} \text{ and } \texttt{c} \notin K) \\ \texttt{VI} & \text{otherwise} \end{cases} \qquad H' = H[\texttt{l} \mapsto \texttt{C<NO,VI'>}(\overline{\texttt{null}})]}{K \vdash H, \texttt{new C<NO,VI>}(\overline{\texttt{v}}) \rightarrow H', \texttt{l.build}(\overline{\texttt{v}})\texttt{;return l}}$$

We assumed that $K, \Gamma_H \vdash \texttt{new C<NO,VI>}(\overline{\texttt{v}}) : \texttt{T}$", therefore from T-New we know that

$$\mathit{mtype}(\bot, \texttt{build}, \texttt{C<NO,NI>}) = \overline{\texttt{T}} \rightarrow \texttt{U}$$

Thus there is a constructor with $\#(\overline{\texttt{v}})$ of arguments.

Rule R-return Trivial because there are no assumptions for `v;return l`

□

We prove **preservation** (Lem. 5.6) for method invocation by using Lem. 5.5 that uses induction on the size of the method body.

Lemma 5.5. ***(Invocation Substitution)*** *For every heap H that is well-typed for K, location $\texttt{l} \in \mathrm{dom}(H)$, values $\overline{\texttt{v}}$, types U, and guard IG, where*

$$\begin{aligned} H[\texttt{l}] &= \texttt{C<NO,NI>} \\ K, \Gamma_H &\vdash \texttt{NI} \leq \texttt{IG} \\ K, \Gamma_H &\vdash \overline{\texttt{v}} : \overline{\texttt{U}'} \\ K, \Gamma_H &\vdash \overline{\texttt{U}'} \leq \overline{[\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{U}} \end{aligned} \tag{4}$$

Then, for any e" *such that* $\emptyset, \Gamma \vdash \texttt{e}" : \texttt{S}$, $\Gamma = \{\texttt{I} : \texttt{IG}, \overline{\texttt{x}} : \overline{\texttt{U}}, \texttt{this} : \texttt{C<O,I>}\}$, *then*

$$\begin{aligned} K, \Gamma_H &\vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}, \overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}]\texttt{e}" : \texttt{S}' \\ K, \Gamma_H &\vdash \texttt{S}' \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{S} \end{aligned}$$

Proof. We will prove by induction on the structure of e".

e" = (**e;return l**) Impossible because Γ does not contain locations.

e" = (**v**) Γ does not contain locations. Therefore, $\texttt{v} = \texttt{null}$, and we can choose $\texttt{S}' = [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{S}$.

e" = (**this**) Then $\texttt{S} = \texttt{C<O,I>}$, and

$$\begin{aligned} &[\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}, \overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}]\texttt{e}" = \texttt{l} \\ &[\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{S} = \texttt{C<NO,NI>} \\ &\texttt{S}' = \texttt{C<NO,NI>} \\ &K, \Gamma_H \vdash \texttt{l} : \texttt{S}' \\ &K, \Gamma_H \vdash \texttt{S}' \leq \texttt{C<NO,NI>} \end{aligned}$$

$\mathbf{e}" = (\mathbf{x}_i)$ Then $\mathtt{S} = \mathtt{U}_i$. We assumed that

$$K, \Gamma_H \vdash \mathtt{v}_i : \mathtt{U}_i'$$
$$K, \Gamma_H \vdash \mathtt{U}_i' \leq [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{U}_i$$

Therefore,

$$[\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}, \overline{\mathtt{v}}/\overline{\mathtt{x}}, \mathtt{l/this}]\mathtt{e}" = \mathtt{v}_i$$
$$\mathtt{S}' = \mathtt{U}_i'$$
$$\mathtt{S} = \mathtt{U}_i$$
$$K, \Gamma_H \vdash \mathtt{v}_i : \mathtt{S}'$$
$$K, \Gamma_H \vdash \mathtt{S}' \leq [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{S}$$

$\mathbf{e}" = (\mathbf{new\ D<FO,VI>(\overline{e})})$ We assumed that $K, \Gamma_H \vdash \mathtt{new\ D<FO,VI>}(\overline{\mathtt{e}}) : \mathtt{S}$. From T-NEW, $\mathtt{S} = \mathtt{D<FO,VI>}$ and

$$\mathit{mtype}(\bot, \mathtt{build}, \mathtt{D<FO,VI>}) = \overline{\mathtt{T}} \to \mathtt{U} \qquad \emptyset, \Gamma \vdash \overline{\mathtt{e}} : \overline{\mathtt{T}'} \qquad \emptyset, \Gamma \vdash \overline{\mathtt{T}'} \leq \overline{\mathtt{T}} \tag{5}$$

By induction on $\mathtt{e}_i$, we have that

$$\begin{aligned} &K, \Gamma_H \vdash [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}, \overline{\mathtt{v}}/\overline{\mathtt{x}}, \mathtt{l/this}]\mathtt{e}_i : \mathtt{V}_i \\ &K, \Gamma_H \vdash \mathtt{V}_i \leq [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{T}_i' \end{aligned} \tag{6}$$

From Lem. 2.6 and (5), we have that

$$K, \Gamma_H \vdash [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{T}_i' \leq [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{T}_i \tag{7}$$

From transitivity, (6), and (7), we have

$$K, \Gamma_H \vdash \mathtt{V}_i \leq [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{T}_i \tag{8}$$

From definition of *mtype* we have

$$\mathit{mtype}(\bot, \mathtt{build}, [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{D<FO,VI>}) = [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}](\overline{\mathtt{T}} \to \mathtt{U}) \tag{9}$$

From T-NEW, (8), and (9),

$$K, \Gamma_H \vdash [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}, \overline{\mathtt{v}}/\overline{\mathtt{x}}, \mathtt{l/this}]\mathtt{new\ D<FO,VI>}(\overline{\mathtt{e}}) : [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{D<FO,VI>}$$

$\mathbf{e}" = (\mathbf{e.f})$ From T-FIELD-ACCESS,

$$\emptyset, \Gamma \vdash \mathtt{e} : \mathtt{D<MO,IP>}$$
$$\emptyset, \Gamma \vdash \mathtt{e.f} : \mathtt{S}$$
$$\mathtt{S} = \mathit{ftype}(\mathtt{e}, \mathtt{f}, \mathtt{D<MO,IP>})$$

Recall that $\mathtt{S} = \mathit{ftype}(\mathtt{e}, \mathtt{f}, \mathtt{D<MO,IP>}) = \mathit{substitute}(\mathtt{e}, \mathtt{D<MO,IP>}, \mathit{ftype}(\mathtt{f}, \mathtt{D}))$. Note that $\mathtt{e}$ is not a location, and thus if $\mathtt{f}$ is $\mathtt{this}$-owned, then $\mathtt{e} = \mathtt{this}$. Consider first the case that $\mathtt{f}$ is $\mathtt{this}$-owned, thus $\mathtt{e} = \mathtt{this}$. Let $\mathit{ftype}(\mathtt{f}, \mathtt{D}) = \mathtt{FT}$, and $O(\mathtt{FT}) = \mathtt{This}$. We assumed that $\emptyset, \Gamma \vdash \mathtt{this}.f : \mathtt{S}$. Then, $\mathtt{S} = \mathtt{FT}$. We need to show that

$$K, \Gamma_H \vdash [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}, \overline{\mathtt{v}}/\overline{\mathtt{x}}, \mathtt{l/this}]\mathtt{this.f} : \mathtt{S}'$$
$$K, \Gamma_H \vdash \mathtt{S}' \leq [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{FT}$$

From T-FIELD-ACCESS

$$K, \Gamma_H \vdash \mathtt{l.f} : [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{FT}$$

, i.e., $\mathtt{S}' = [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{FT}$.

Now suppose that $\mathtt{f}$ is not $\mathtt{this}$-owned. Therefore, $\mathtt{S} = [\mathtt{IP/I}, \mathtt{MO/O}]\mathtt{FT}$. We need to show that

$$\begin{aligned} &K, \Gamma_H \vdash [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}, \overline{\mathtt{v}}/\overline{\mathtt{x}}, \mathtt{l/this}]\mathtt{e.f} : \mathtt{S}' \\ &K, \Gamma_H \vdash \mathtt{S}' \leq [\mathtt{l/This}, \mathtt{NI/I}, \mathtt{NO/O}]\mathtt{S} \qquad \mathtt{S} = [\mathtt{IP/I}, \mathtt{MO/O}]\mathtt{FT} \end{aligned} \tag{10}$$

From the induction on e, we have that

$$\begin{array}{l} K,\Gamma_H \vdash [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O},\overline{\mathtt{v}}/\overline{\mathtt{x}},\mathtt{l}/\mathtt{this}]\mathtt{e} : \mathtt{S}'' \\ K,\Gamma_H \vdash \mathtt{S}'' \leq [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O}]\mathtt{D<MO,IP>} \end{array} \tag{11}$$

From (11), $O(\mathtt{S}'') = [\mathtt{l}/\mathtt{This},\mathtt{NO}/\mathtt{O}]\mathtt{MO}$. From (10) and (11), and Lem. 2.2, we have that

$$\mathtt{S}' = \mathit{ftype}(\bot,\mathtt{f},\mathtt{S}'') = [I(\mathtt{S}'')/\mathtt{I},O(\mathtt{S}'')/\mathtt{O}]\mathtt{FT} \leq [([\mathtt{NI}/\mathtt{I}]\mathtt{IP})/\mathtt{I},([\mathtt{l}/\mathtt{This},\mathtt{NO}/\mathtt{O}]\mathtt{MO})/\mathtt{O}]\mathtt{FT} =$$
$$= [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O}]([\mathtt{IP}/\mathtt{I},\mathtt{MO}/\mathtt{O}]\mathtt{FT})$$

e$'' = (\mathtt{e.f{=}e}')$ The challenge in field assignment is that (by induction) the type of the substitution of e changed covariantly (i.e., it is a subtype of the substitution of the type), and $\mathtt{e}'$ also changed covariantly. However, we will prove that because $I(\mathtt{e})$ is Raw, and e is either this or this-owned, then e is invariant.

We assumed that $\emptyset,\Gamma \vdash \mathtt{e.f{=}e}' : \mathtt{S}$. From T-Field-Assignment, we know that

$$\begin{array}{l} \emptyset,\Gamma \vdash \mathtt{e.f} : \mathtt{T} \quad \emptyset,\Gamma \vdash \mathtt{e}' : \mathtt{S} \quad \emptyset,\Gamma \vdash \mathtt{S} \leq \mathtt{T} \quad \emptyset,\Gamma \vdash \mathtt{e} : \mathtt{D<MO,IP>} \\ \emptyset,\Gamma \vdash \mathtt{IP} \leq \mathtt{Raw} \quad \mathit{isTransitive}(\mathtt{e},\Gamma,\mathtt{D<MO,IP>}) \quad \mathtt{MO} \neq \mathtt{?} \end{array} \tag{12}$$

We wish to prove all the assumptions in (12) after substituting $[\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O},\overline{\mathtt{v}}/\overline{\mathtt{x}},\mathtt{l}/\mathtt{this}]$.

By induction on $\mathtt{e}'$ we have that

$$\begin{array}{l} K,\Gamma_H \vdash [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O},\overline{\mathtt{v}}/\overline{\mathtt{x}},\mathtt{l}/\mathtt{this}]\mathtt{e}' : \mathtt{S}' \\ K,\Gamma_H \vdash \mathtt{S}' \leq [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O}]\mathtt{S} \end{array} \tag{13}$$

By induction on e.f we have that

$$\begin{array}{l} K,\Gamma_H \vdash [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O},\overline{\mathtt{v}}/\overline{\mathtt{x}},\mathtt{l}/\mathtt{this}]\mathtt{e.f} : \mathtt{T}' \\ K,\Gamma_H \vdash \mathtt{T}' \leq [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O}]\mathtt{T} \end{array} \tag{14}$$

From the proof of field access above, we see that the class of $\mathtt{T}'$ and T is the same. By induction on e we have that

$$\begin{array}{l} K,\Gamma_H \vdash [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O},\overline{\mathtt{v}}/\overline{\mathtt{x}},\mathtt{l}/\mathtt{this}]\mathtt{e} : \mathtt{D'<MO',IP'>} \\ K,\Gamma_H \vdash \mathtt{D'<MO',IP'>} \leq [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O}]\mathtt{D<MO,IP>} \end{array} \tag{15}$$

Because $\mathtt{MO} \neq \mathtt{?}$, from Lem. 2.1 part (i), we have that $\mathtt{MO}' = \mathtt{MO}$.

We now prove that the following holds:

$$\begin{array}{l} K,\Gamma_H \vdash [\mathtt{NI}/\mathtt{I}]\mathtt{IP} \leq \mathtt{Raw} \\ \mathit{isTransitive}([\mathtt{l}/\mathtt{this}]\mathtt{e},\Gamma,[\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O}]\mathtt{D<MO,IP>}) \end{array} \tag{16}$$

From (12), $\emptyset,\Gamma \vdash \mathtt{IP} \leq \mathtt{Raw}$. From our syntax, and because Γ does not contain locations:

$$\mathtt{IP} = \mathtt{ReadOnly} \mid \mathtt{Immut} \mid \mathtt{Mutable} \mid \mathtt{I}$$

If $\mathtt{IP} = \mathtt{Mutable}$ then we proved (16). Therefore, it must be that $\mathtt{IP} = \mathtt{I}$ (thus $[\mathtt{NI}/\mathtt{I}]\mathtt{IP} = \mathtt{NI}$) and $\Gamma(\mathtt{I}) = \mathtt{IG} \leq \mathtt{Raw}$. From (4) ($K,\Gamma_H \vdash \mathtt{NI} \leq \mathtt{IG}$), we proved the first part of (16) that $K,\Gamma_H \vdash [\mathtt{NI}/\mathtt{I}]\mathtt{IP} \leq \mathtt{Raw}$. If $\mathtt{IG} = \mathtt{Mutable}$ then, $\mathtt{NI} = \mathtt{Mutable}$, which proved (16). Therefore $\mathtt{IG} = \mathtt{Raw}$, and from the definition of *isTransitive* we have that e=this or $\mathtt{MO} = \mathtt{This}$. If $\mathtt{e} = \mathtt{this}$ then *isTransitive*$(\mathtt{l},\ldots)$, which proved (16). Thus $\mathtt{MO} = \mathtt{This}$, and

$$\begin{array}{r} \mathtt{D<MO,IP>} = \mathtt{D<This,I>} \\ [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O}]\mathtt{D<MO,IP>} = \mathtt{D<l,NI>} \end{array} \tag{17}$$

From (4), $H[\mathtt{l}] = \mathtt{C<NO,NI>}$. If $\mathtt{NI} = \mathtt{Mutable}$ then we proved (16). Otherwise $\mathtt{NI} = \mathtt{Immut}_{\mathtt{l}'}$. Because H is well-typed for K, $\mathtt{l}' \preceq^T \mathtt{l}$, thus $\mathtt{l}' \not\prec_\theta \mathtt{l}$, proving (16) (because if $\mathtt{a} \prec_\theta \mathtt{b}$ then $\mathtt{b} \prec^T \mathtt{a}$).

(i) If $\mathtt{e} = \mathtt{this}$. Let $\mathit{ftype}(\mathtt{f},\mathtt{D}) = \mathtt{FT}$. We assumed in (12) that $\emptyset,\Gamma \vdash \mathtt{this}.f : \mathtt{T}$. Then, $\mathtt{T} = \mathtt{FT}$. From T-Field-Access

$$K,\Gamma_H \vdash \mathtt{l.f} : [\mathtt{l}/\mathtt{This},\mathtt{NI}/\mathtt{I},\mathtt{NO}/\mathtt{O}]\mathtt{FT}$$

From definition of *isTransitive*, we have that *isTransitive*(l,...) holds. From (12) ($\emptyset, \Gamma \vdash \texttt{S} \leq \texttt{T}$) and Lem. 2.6, we have

$$K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{S} \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{T} \tag{18}$$

From (13), (18), and transitivity, we have

$$K, \Gamma_H \vdash \texttt{S}' \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{FT} \tag{19}$$

To summarize, from (19), (16) ($K, \Gamma_H \vdash \texttt{NI} \leq \texttt{Raw}$), we have that

$$\begin{array}{l} K, \Gamma_H \vdash \texttt{l.f} : [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{FT} \quad K, \Gamma_H \vdash \texttt{e}' : \texttt{S}' \quad K, \Gamma_H \vdash \texttt{S}' \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{FT} \\ K, \Gamma_H \vdash \texttt{l} : \texttt{C<NO,NI>} \quad K, \Gamma_H \vdash \texttt{NI} \leq \texttt{Raw} \quad \textit{isTransitive}(\texttt{l}, \ldots) \end{array} \tag{20}$$

Therefore, from (20), and T-Field-Assignment, we proved that

$$\begin{array}{l} K, \Gamma_H \vdash \texttt{l.f=e}' : \texttt{S}' \\ K, \Gamma_H \vdash \texttt{S}' \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{FT} \end{array}$$

(ii) If $\texttt{e} \neq \texttt{this}$, then from (16), we have

$$\textit{isTransitive}(\bot, \Gamma, [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{D<MO,IP>}) \tag{21}$$

From (15) and (21) and Lem. 2.3, we have that $\texttt{IP}' = [\texttt{NI}/\texttt{I}]\texttt{IP}$. To summarize, from (13), (14), (18), (21),

$$\begin{array}{l} K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}, \overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}]\texttt{e.f} : [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{T} \\ K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}, \overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}]\texttt{e}' : \texttt{S}' \\ K, \Gamma_H \vdash \texttt{S}' \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{S} \\ K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{S} \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{T} \\ K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}, \overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}]\texttt{e} : \texttt{D'<MO,IP'>} \\ K, \Gamma_H \vdash \texttt{IP}' \leq \texttt{Raw} \\ \quad \textit{isTransitive}(\bot, \Gamma, \texttt{D'<MO,IP'>}) \quad \texttt{MO} \neq \texttt{?} \end{array} \tag{22}$$

Thus, from (22), and T-Field-Assignment, we proved that

$$\begin{array}{l} K, \Gamma_H \vdash \texttt{e.f=e}' : \texttt{S}' \\ K, \Gamma_H \vdash \texttt{S}' \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{S} \end{array}$$

e" $= (\texttt{e}_0\texttt{.m}(\overline{\texttt{e}}))$ The proof is similar in spirit to field assignment: the challenge is that both $\texttt{e}_0$ and $\texttt{e}_i$ change covariantly. Let $\texttt{IG}'$ be the guard of m. If $\texttt{IG}' = \texttt{ReadOnly}$ then the parameters of m cannot include I. If $\texttt{IG}' = \texttt{Mutable} \mid \texttt{Immut}$, then I remains with the same bound. The challenge is when $\texttt{IG}' = \texttt{Raw}$, then we use either the fact the $\texttt{e}_0$ is either this or this-owned, to prove that $\texttt{e}_0$ is invariant (like in field assignment).

With respect to wildcards, if the receiver $\texttt{e}_0$ has a wildcard, then after the covariant change it might no longer be the case. Therefore we require that the owner of method parameters in this case must be World. (it cannot be O nor This).

From T-Invoke,

$$\frac{\begin{array}{l} \emptyset, \Gamma \vdash \texttt{e}_0 : \texttt{D<MO,IP>} \quad \textit{mtype}(\texttt{e}_0, \texttt{m}, \texttt{D<MO,IP>}) = \overline{\texttt{T}} \to \texttt{W} \quad \emptyset, \Gamma \vdash \overline{\texttt{e}} : \overline{\texttt{T}'} \quad \emptyset, \Gamma \vdash \overline{\texttt{T}'} \leq \overline{\texttt{T}} \quad \textit{mguard}(\texttt{m}, \texttt{D}) = \texttt{IG}' \\ \emptyset, \Gamma \vdash \texttt{IP} \leq \texttt{IG}' \quad \texttt{IG}' = \texttt{Raw} \Rightarrow \textit{isTransitive}(\texttt{e}_0, \Gamma, \texttt{D<MO,IP>}) \quad \textit{mtype}(\texttt{m}, \texttt{D}) = \overline{\texttt{U}} \to \texttt{V} \quad O(\overline{\texttt{T}}) = \texttt{?} \Rightarrow O(\overline{\texttt{U}}) = \texttt{?} \end{array}}{\emptyset, \Gamma \vdash \texttt{e}_0\texttt{.m}(\overline{\texttt{e}}) : \texttt{W}} \tag{23}$$

By induction on $\texttt{e}_0$, we have that

$$\begin{array}{l} K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}, \overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}]\texttt{e}_0 : \texttt{D'<MO',IP'>} \\ K, \Gamma_H \vdash \texttt{D'<MO',IP'>} \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{D<MO,IP>} \end{array} \tag{24}$$

Note that, in contrast with field assignment, here we might have $\texttt{MO} = \texttt{?}$, and then $\texttt{MO}' \neq \texttt{MO}$. However, from Lem. 2.1 part (i),

$$\texttt{MO} \neq \texttt{?} \Rightarrow \texttt{MO}' = [\texttt{l}/\texttt{This}, \texttt{NO}/\texttt{O}]\texttt{MO} \tag{25}$$

By induction on $\texttt{e}_i$, we have that

$$\begin{array}{l} K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}, \overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}]\texttt{e}_i : \texttt{S}'_i \\ K, \Gamma_H \vdash \texttt{S}'_i \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{T}'_i \end{array} \tag{26}$$

From (26) and (23) and Lem. 2.6, we have

$$K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{T}'_i \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{T}_i \tag{27}$$

From transitivity, (26), and (27),

$$K, \Gamma_H \vdash \texttt{S}'_i \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{T}_i \tag{28}$$

Because method overriding maintains the same signature, we have that

$$\mathit{mtype}(\texttt{m}, \texttt{D}') = \mathit{mtype}(\texttt{m}, \texttt{D}) = \overline{\texttt{U}} \to \texttt{V} \tag{29}$$

From definition of *mtype*, (29), and because $\texttt{e}_0$ does not contain locations, we have that

$$\begin{array}{r} \mathit{mtype}(\texttt{e}_0, \texttt{m}, \texttt{D<MO, IP>}) = \overline{\texttt{T}} \to \texttt{W} = [\texttt{MO}/\texttt{O}, \texttt{IP}/\texttt{I}](\overline{\texttt{U}} \to \texttt{V}) \\ \mathit{mtype}([\texttt{l}/\texttt{this}]\texttt{e}_0, \texttt{m}, \texttt{D}'\texttt{<MO}', \texttt{IP}'\texttt{>}) = [\texttt{l}/\texttt{This}, \texttt{MO}'/\texttt{O}, \texttt{IP}'/\texttt{I}](\overline{\texttt{U}} \to \texttt{V}) \end{array} \tag{30}$$

We will always prove that the parameters are invariant, i.e.,

$$\mathit{mtype}([\texttt{l}/\texttt{this}]\texttt{e}_0, \texttt{m}, \texttt{D}'\texttt{<MO}', \texttt{IP}'\texttt{>}) = \overline{[\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{T}_i} \to \texttt{W}' \tag{31}$$

We will also prove that

$$K, \Gamma_H \vdash \texttt{IP}' \leq \texttt{IG}' \qquad \texttt{IG}' = \texttt{Raw} \Rightarrow \mathit{isTransitive}([\texttt{l}/\texttt{this}]\texttt{e}_0, \Gamma, \texttt{D}'\texttt{<MO}', \texttt{IP}'\texttt{>}) \tag{32}$$

Because there are no wildcards after substitution, from (31) and (32), we will have that

$$\begin{array}{l} K, \Gamma_H \vdash [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}, \overline{\texttt{v}}/\overline{\texttt{x}}, \texttt{l}/\texttt{this}]\texttt{e}_0\texttt{.m}(\overline{\texttt{e}}) : \texttt{W}' \\ K, \Gamma_H \vdash \texttt{W}' \leq [\texttt{l}/\texttt{This}, \texttt{NI}/\texttt{I}, \texttt{NO}/\texttt{O}]\texttt{W} \end{array}$$

Next we prove (31) just for the owner parameter, i.e., we want to show that (from (30) and (31))

$$O([\texttt{l}/\texttt{This}, \texttt{MO}'/\texttt{O}]\texttt{U}_i) = O([\texttt{l}/\texttt{This}, \texttt{NO}/\texttt{O}]\texttt{T}_i) \tag{33}$$

From (30),

$$O(\texttt{T}_i) = O([\texttt{MO}/\texttt{O}]\texttt{U}_i) \tag{34}$$

If $O(\texttt{U}_i) = \texttt{This}$, then both sides of (33) are $\texttt{l}$. If $O(\texttt{U}_i) \neq \texttt{O}$, then both sides of (33) are $O(\texttt{U}_i)$. The last case is that $O(\texttt{U}_i) = \texttt{O}$. From (23), we have

$$O(\texttt{T}_i) = \texttt{?} \Rightarrow O(\texttt{U}_i) = \texttt{?} \tag{35}$$

On the one hand, if $\texttt{MO} = \texttt{?}$, then $O(\texttt{T}_i) = \texttt{?}$, thus $O(\texttt{U}_i) = \texttt{?}$, and both sides of (33) are $\texttt{?}$. On the other hand, if $\texttt{MO} \neq \texttt{?}$, then from (25)

$$\texttt{MO}' = [\texttt{l}/\texttt{This}, \texttt{NO}/\texttt{O}]\texttt{MO}$$

which proves (33).

From (33), in order to prove (31), we just need to show it for the immutability parameter, i.e., (from (30))

$$I([\texttt{IP}'/\texttt{I}]\texttt{U}_i) = I([\texttt{NI}/\texttt{I}]([\texttt{IP}/\texttt{I}]\texttt{U}_i)) \tag{36}$$

Recall the following: From (24), we know that $K,\Gamma_H \vdash \texttt{D'<MO',IP'>} \leq [\texttt{l/This,NI/I,NO/O}]\texttt{D<MO,IP>}$. From (23), $\emptyset,\Gamma \vdash \texttt{IP} \leq \texttt{IG'}$. From (4), $K,\Gamma_H \vdash \texttt{NI} \leq \texttt{IG}$ and $\texttt{I} : \texttt{IG} \in \Gamma$.

We will split the proof by the four possible values of $\texttt{IG'} = \texttt{ReadOnly} \mid \texttt{Mutable} \mid \texttt{Immut} \mid \texttt{Raw}$. For each case we need to prove (32) and (36).

(i) $\texttt{IG'} = \textbf{ReadOnly}$ Because any immutability is a subtype of ReadOnly, we proved (32). Furthermore, when $\texttt{IG'} = \texttt{ReadOnly}$ the signature of parameters cannot contain I, i.e., $I(\texttt{U}_i) \neq \texttt{I}$, which proved (36).

(ii) $\texttt{IG'} = \textbf{Mutable}$ From (23), $\emptyset,\Gamma \vdash \texttt{IP} \leq \texttt{Mutable}$, thus either $\texttt{IP} = \texttt{Mutable}$ or ($\texttt{IG} = \texttt{Mutable}$ and $\texttt{IP} = \texttt{I}$). If $\texttt{IP} = \texttt{I}$, then from (4), $K,\Gamma_H \vdash \texttt{NI} \leq \texttt{Mutable}$, thus $\texttt{NI} = \texttt{Mutable}$. Thus, $K,\Gamma_H \vdash [\texttt{NI/I}]\texttt{IP} \leq \texttt{Mutable}$. Therefore, from (24) and Lem. 2.4 part (i), we have that $K,\Gamma_H \vdash \texttt{IP'} \leq \texttt{Mutable}$, which proved (32).

We showed that if $\texttt{IP} = \texttt{I}$, then $\texttt{NI} = \texttt{Mutable}$ and $\texttt{IP'} = \texttt{Mutable}$, proving (36). We also showed that if $\texttt{IP} = \texttt{Mutable}$ then $\texttt{IP'} = \texttt{Mutable}$, proving (36).

(iii) $\texttt{IG'} = \textbf{Immut}$ Exactly like part (ii), but we use Lem. 2.4 part (ii) instead of part (i), and ($\texttt{Immut}_{\texttt{l'}}$ where $\texttt{l'} \notin K$) instead of Mutable.

(iv) $\texttt{IG'} = \textbf{Raw}$ Exactly like in field assignment, we prove that:

$$\begin{array}{l} K,\Gamma_H \vdash [\texttt{NI/I}]\texttt{IP} \leq \texttt{Raw} \\ \mathit{isTransitive}([\texttt{l/this}]\texttt{e},\Gamma,[\texttt{l/This,NI/I,NO/O}]\texttt{D<MO,IP>}) \end{array} \tag{37}$$

If $\texttt{e} = \texttt{this}$, then $\texttt{D} = \texttt{C}$, $\texttt{IP} = \texttt{I}$ (because $\emptyset,\Gamma \vdash \texttt{this} : \texttt{D<O,I>}$) and $\texttt{IP'} = \texttt{NI}$ (because $K,\Gamma_H \vdash \texttt{l} : \texttt{D<NO,NI>}$), therefore,

$$I([\texttt{NI/I}]\texttt{U}_i) = I([\texttt{NI/I}]([\texttt{I/I}]\texttt{U}_i))$$

which proved (36). Furthermore, because $\emptyset,\Gamma \vdash \texttt{IP} \leq \texttt{Raw}$ (and $\texttt{IP} = \texttt{I}$), we know that $\texttt{IG} \leq \texttt{Raw}$. Thus from (4), $K,\Gamma_H \vdash \texttt{NI} \leq \texttt{Raw}$. Finally because this was replaced with l, and $\mathit{isTransitive}(\texttt{l},\ldots)$ always holds, then we proved (32).

If $\texttt{e} \neq \texttt{this}$, then from (37), we have that $\mathit{isTransitive}(\bot,\Gamma,[\texttt{l/This,NI/I,NO/O}]\texttt{D<MO,IP>})$, thus from definition of *isTransitive* we have that $\texttt{MO} \neq \texttt{?}$. From Lem. 2.4 part (iii) and (24), we know that $\texttt{IP'} = [\texttt{NI/I}]\texttt{IP}$, which proved (36). Combined with (37), we have that $K,\Gamma_H \vdash \texttt{IP'} \leq \texttt{Raw}$. From (25), we have $\texttt{MO'} = [\texttt{l/This,NO/O}]\texttt{MO}$. Therefore, $\texttt{D'<MO',IP'>} = [\texttt{l/This,NI/I,NO/O}]\texttt{D<MO,IP>}$, which proved (32).

□

Lemma 5.6. ***(Subtype preservation)*** *For every closed expression* $\texttt{e"} \neq \texttt{v}$*, H, and K, if* $K,\Gamma_H \vdash \texttt{e"} : \texttt{T"}$ *and* $K \vdash H,\texttt{e"} \rightarrow H',\grave{\texttt{e}}$ *and H is well-typed for* $K \cup K(\texttt{e"})$*, then* $K,\Gamma_{H'} \vdash \grave{\texttt{e}} : \grave{\texttt{T}}$ *and* $K,\Gamma_{H'} \vdash \grave{\texttt{T}} \leq \texttt{T"}$.

Proof. We prove by examining all possible reduction rules.

Congruence for field access Consider the congruence rule for field access

$$\frac{K \vdash H,\texttt{e} \rightarrow H',\texttt{e}'}{K \vdash H,\texttt{e.f} \rightarrow H',\texttt{e}'\texttt{.f}}$$

We assumed that $K,\Gamma_H \vdash \texttt{e.f} : \texttt{T"}$ and by induction $K,\Gamma_H \vdash \texttt{e} : \texttt{T}$, $K,\Gamma_{H'} \vdash \texttt{e}' : \texttt{T}'$ and $K,\Gamma_{H'} \vdash \texttt{T}' \leq \texttt{T}$. Let $\texttt{T} = \texttt{C<MO,IP>}$. Because $\texttt{e} \neq \texttt{this}$ (cause e is closed) and $\texttt{e} \neq \texttt{l}$ (cause a location cannot be reduced further), then field f is not this-owned, and $\texttt{T"} = \mathit{ftype}(\bot,\texttt{f},\texttt{T})$.

Because $K,\Gamma_{H'} \vdash \texttt{T}' \leq \texttt{T}$, from Lem. 2.1 part (iii), then $\texttt{C'}$ is a subtype of C. Therefore $\mathit{fields}(\texttt{C'})$ must contain the same field f which is not this-owned. Thus, $K,\Gamma_{H'} \vdash \grave{\texttt{e}} : \grave{\texttt{T}}$, where $\grave{\texttt{T}} = \mathit{ftype}(\bot,\texttt{f},\texttt{T}')$. The last thing we need to prove is that $K,\Gamma_{H'} \vdash \grave{\texttt{T}} \leq \texttt{T"}$, which follows from Lem. 2.7.

Congruence for method receiver Consider the congruence rule for method receiver

$$\frac{K \vdash H,\texttt{e}_0 \rightarrow H',\texttt{e}_0'}{K \vdash H,\texttt{e}_0\texttt{.m}(\overline{\texttt{e}}) \rightarrow H',\texttt{e}_0'\texttt{.m}(\overline{\texttt{e}})}$$

Similarly to field access, because $\texttt{e}_0$ is not a location, then none of the parameters or return type of method m is this-owned. Proving that $K,\Gamma_{H'} \vdash \grave{\texttt{T}} \leq \texttt{T"}$ (i.e., the return type is preserved) is done similarly to field

access, by noting that method overriding maintains the same return type. (The return type could also change covariantly and the proof would still hold.) However, proving that $K, \Gamma_{H'} \vdash \grave{e} : \grave{T}$ is more challenging because: (i) we need to show that e_0' satisfies the guard, and (ii) the type of method parameters after substitution can change covariantly (as opposed to FGJ, which is invariant).

We will first prove that e_0' satisfies the guard. We assumed that $K, \Gamma_H \vdash e_0.m(\overline{e}) : T$" and by induction $K, \Gamma_H \vdash e_0 : T$, $K, \Gamma_{H'} \vdash e_0' : T'$ and $K, \Gamma_{H'} \vdash T' \leq T$. From T-INVOKE, we know that

$$K, \Gamma_H \vdash e_0 : C\texttt{<}MO, IP\texttt{>} \qquad \mathit{mguard}(m, C) = IG \qquad K, \Gamma_H \vdash IP \leq IG \qquad IG = Raw \Rightarrow \mathit{isTransitive}(e_0, \Gamma, C\texttt{<}MO, IP\texttt{>})$$

Because e_0 is neither this nor l (because it was reduced), then $IG = Raw \Rightarrow \mathit{isTransitive}(\bot, \Gamma, C\texttt{<}MO, IP\texttt{>})$. Let $T' = C'\texttt{<}MO, IP'\texttt{>}$ and $\mathit{mguard}(m, C') = IG'$. Because $K, \Gamma_{H'} \vdash T' \leq T$, from Lem. 2.1 part (iii), then C' is a subtype of C. From the restriction on method overriding with guards, $IG \leq IG'$. From Lem. 2.4 part (iv), we have that $IP' \leq IG$. From transitivity, $IP' \leq IG'$.

Next we show that $IG = Raw \Rightarrow \mathit{isTransitive}(e_0', \Gamma, C'\texttt{<}MO, IP'\texttt{>})$. If $IG = Raw$ then we showed that $\mathit{isTransitive}(\bot, \Gamma, C\texttt{<}MO, IP\texttt{>})$. From Lem. 2.3 we have that $IP' = IP$, thus $IG = Raw \Rightarrow \mathit{isTransitive}(e_0', \Gamma, C'\texttt{<}MO, IP'\texttt{>})$.

Let

$$\begin{aligned} \mathit{mtype}(e_0, m, C\texttt{<}MO, IP\texttt{>}) &= \overline{U} \to V \\ \mathit{mtype}(e_0', m, C'\texttt{<}MO, IP'\texttt{>}) &= \overline{U'} \to V' \\ K, \Gamma_H &\vdash \overline{e} : \overline{U''} \end{aligned}$$

Because $K, \Gamma_H \vdash \overline{U''} \leq \overline{U}$, from Lem. 2.8, we have that $K, \Gamma_H \vdash \overline{U''} \leq \overline{U'}$.

Therefore, all the assumptions in T-INVOKE are fulfilled (the requirement for wildcards is fulfilled because all types are closed), and we proved that $K, \Gamma_{H'} \vdash \grave{e} : \grave{T}$.

Congruence for method argument Trivial.

Congruence for new instance Trivial.

Congruence for the rvalue of field assignment Trivial.

Congruence for the receiver of field assignment Consider the congruence rule for the receiver of field assignment

$$\frac{K \vdash H, e \to H', e'}{K \vdash H, e.f{=}e'' \to H', e'.f{=}e''}$$

We assumed that $K, \Gamma_H \vdash e.f{=}e'' : T$" and by induction $K, \Gamma_H \vdash e : T$, $K, \Gamma_{H'} \vdash e' : T'$ and $K, \Gamma_{H'} \vdash T' \leq T$. We will show that $K, \Gamma_H \vdash e.f{=}e'' : T$". From T-FIELD-ASSIGNMENT:

$$K, \Gamma_H \vdash e.f : F \quad K, \Gamma_H \vdash e'' : T'' \quad K, \Gamma_H \vdash T'' \leq F \quad K, \Gamma_H \vdash e : C\texttt{<}MO, IP\texttt{>} \quad K, \Gamma_H \vdash IP \leq Raw \quad \mathit{isTransitive}(e, \Gamma_H, C\texttt{<}MO, IP\texttt{>})$$

We need to show that:

$$K, \Gamma_H \vdash e'.f : F' \quad K, \Gamma_H \vdash T'' \leq F' \quad K, \Gamma_H \vdash e' : C'\texttt{<}MO, IP'\texttt{>} \quad K, \Gamma_H \vdash IP' \leq Raw \quad \mathit{isTransitive}(e, \Gamma_H, C'\texttt{<}MO, IP'\texttt{>})$$

Because e was reduced, we know it is not a location, so $\mathit{isTransitive}(\bot, \Gamma_H, C\texttt{<}MO, IP\texttt{>})$. From Lem. 2.3, we have that $IP = IP'$. Therefore $F' = F$ (because $\mathit{ftype}(f, C) = \mathit{ftype}(f, C')$), and $\mathit{isTransitive}(e, \Gamma_H, C'\texttt{<}MO, IP'\texttt{>})$.

Congruence for return R-C1 Consider the congruence rule for $e;\texttt{return l}$

$$\frac{K \cup \{l\} \vdash H, e \to H', e'}{K \vdash H, e;\texttt{return l} \to H', e';\texttt{return l}}$$

We assumed that

$$\begin{aligned} K, \Gamma_H \vdash e;\texttt{return l} : T" \quad e" = e;\texttt{return l} \quad \grave{e} = e';\texttt{return l} \\ K \cup \{l\}, \Gamma_H \vdash e : T \quad K \cup \{l\}, \Gamma_{H'} \vdash e' : T' \quad K \cup \{l\}, \Gamma_{H'} \vdash T' \leq T \end{aligned}$$

We need to prove that $K, \Gamma_{H'} \vdash \grave{e} : \grave{T}$ and $K, \Gamma_{H'} \vdash \grave{T} \leq T$".

According to T-RETURN:

$$\frac{K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{e}' : \mathtt{T}'}{K, \Gamma_{H'} \vdash \mathtt{e}'\mathtt{;return\ l} : \Gamma_{H'}(\mathtt{l})}$$

Because $K \cup \{\mathtt{l}\}, \Gamma_{\grave{H}} \vdash \mathtt{e}' : \mathtt{T}'$, we proved that $K, \Gamma_{H'} \vdash \mathtt{e}'\mathtt{;return\ l} : \Gamma_{H'}(\mathtt{l})$, i.e., $K, \Gamma_{H'} \vdash \grave{\mathtt{e}} : \grave{\mathtt{T}}$.

We still need to prove that $K, \Gamma_{H'} \vdash \grave{\mathtt{T}} \leq \mathtt{T}$". Because $\grave{\mathtt{T}} = \Gamma_{H'}(\mathtt{l})$ and $\mathtt{T}" = \Gamma_H(\mathtt{l})$ then $\grave{\mathtt{T}} = \mathtt{T}"$, and from reflexivity (S2) we have $K, \Gamma_{H'} \vdash \grave{\mathtt{T}} \leq \mathtt{T}"$.

Rule R-RETURN Trivial

Rule R-NEW According to R-NEW

$$\frac{\mathtt{l} \notin \mathrm{dom}(H) \qquad \mathtt{VI}' = \begin{cases} \mathtt{Immut}_{\mathtt{l}} & \text{if } \mathtt{VI} = \mathtt{Immut} \text{ or } (\mathtt{VI} = \mathtt{Immut}_{\mathtt{C}} \text{ and } \mathtt{c} \notin K) \\ \mathtt{VI} & \text{otherwise} \end{cases} \qquad H' = H[\mathtt{l} \mapsto \mathtt{C<NO,VI'>}(\overline{\mathtt{null}})]}{K \vdash H, \mathtt{new\ C<NO,VI>}(\overline{\mathtt{v}}) \rightarrow H', \mathtt{l.build}(\overline{\mathtt{v}})\mathtt{;return\ l}}$$

We assumed that

$$K, \Gamma_H \vdash \mathtt{e}" : \mathtt{T}" \qquad \mathtt{e}" = \mathtt{new\ C<NO,VI>}(\overline{\mathtt{v}}) \qquad \grave{\mathtt{e}} = \mathtt{l.build}(\overline{\mathtt{v}})\mathtt{;return\ l}$$

We need to prove that $K, \Gamma_{H'} \vdash \grave{\mathtt{e}} : \grave{\mathtt{T}}$ and $K, \Gamma_{H'} \vdash \grave{\mathtt{T}} \leq \mathtt{T}"$.

From T-NEW

$$\frac{\mathit{mtype}(\bot, \mathtt{build}, \mathtt{C<NO,VI>}) = \overline{\mathtt{U}} \rightarrow \mathtt{Z} \qquad K, \Gamma_H \vdash \overline{\mathtt{v}} : \overline{\mathtt{V}} \qquad K, \Gamma_H \vdash \overline{\mathtt{V}} \leq \overline{\mathtt{U}}}{K, \Gamma_H \vdash \mathtt{new\ C<NO,VI>}(\overline{\mathtt{v}}) : \mathtt{C<NO,VI>}} \tag{38}$$

Thus, $\mathtt{T}" = \mathtt{C<NO,VI>}$. From T-RETURN, $\grave{\mathtt{T}} = \mathtt{C<NO,VI'>}$. Because $\mathtt{l} \notin K$, then $K, \Gamma_{H'} \vdash \mathtt{Immut}_{\mathtt{l}} \leq \mathtt{Immut}$, thus $K, \Gamma_{H'} \vdash \grave{\mathtt{T}} \leq \mathtt{T}"$.

We still need to prove that $K, \Gamma_{H'} \vdash \grave{\mathtt{e}} : \grave{\mathtt{T}}$, and from T-RETURN we need to prove that

$$K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{l.build}(\overline{\mathtt{v}}) : \mathtt{Z}$$

Because l is a new location, all the equations in (38) are still true if we replace Γ_H with $\Gamma_{H'}$. From T-INVOKE, and because the guard of build is Raw:

$$\frac{\begin{array}{ccc} K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{l} : \grave{\mathtt{T}} & \mathit{mtype}(\mathtt{l}, \mathtt{build}, \grave{\mathtt{T}}) = \overline{\mathtt{W}} \rightarrow \mathtt{Z}' & K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \overline{\mathtt{v}} : \overline{\mathtt{V}} \qquad K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \overline{\mathtt{V}} \leq \overline{\mathtt{W}} \\ K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{VI}' \leq \mathtt{Raw} & \mathit{isTransitive}(\mathtt{l}, \Gamma, \grave{\mathtt{T}}) & \end{array}}{K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{l.build}(\overline{\mathtt{v}}) : \mathtt{Z}'}$$

Assumption $\mathit{isTransitive}(\mathtt{l}, \Gamma, \grave{\mathtt{T}})$ holds because l is a location. Because $\mathtt{l} \in K \cup \{\mathtt{l}\}$ and VI is either Mutable or Immut or $\mathtt{Immut}_{\mathtt{l}'}$, then this assumption holds $K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{VI}' \leq \mathtt{Raw}$. The only assumption left to prove is that $K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \overline{\mathtt{V}} \leq \overline{\mathtt{W}}$. From (38) we have that $K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \overline{\mathtt{V}} \leq \overline{\mathtt{U}}$. If $\mathtt{VI} = \mathtt{Mutable}$ then $\overline{\mathtt{U}} = \overline{\mathtt{W}}$. Otherwise $\mathtt{VI} = \mathtt{Immut}$, $\grave{\mathtt{T}} = \mathtt{C<NO,Immut\ }_{\mathtt{l}}\mathtt{VI>}$, and we have that

$$\begin{array}{l} \mathit{mtype}(\mathtt{build}, \mathtt{C}) = \overline{\mathtt{FT}} \rightarrow \mathtt{Z"} \\ \mathit{mtype}(\mathtt{l}, \mathtt{build}, \mathtt{C<NO,Immut\ }_{\mathtt{l}}\mathtt{VI>}) = \overline{\mathtt{W}} \rightarrow \mathtt{Z}' \\ \mathit{mtype}(\bot, \mathtt{build}, \mathtt{C<NO,Immut\ >}) = \overline{\mathtt{U}} \rightarrow \mathtt{Z} \\ \mathtt{W}_i = [\mathtt{NO}/\mathtt{O}, \mathtt{Immut}_{\mathtt{l}}/\mathtt{I}]\mathtt{FT}_i \\ \mathtt{U}_i = [\mathtt{NO}/\mathtt{O}, \mathtt{Immut}/\mathtt{I}]\mathtt{FT}_i \\ K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{V}_i \leq \mathtt{U}_i \end{array}$$

We want to prove that $K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{V}_i \leq \mathtt{W}_i$. If $I(\mathtt{FT}_i) \neq \mathtt{I}$ then $\mathtt{W}_i = \mathtt{U}_i$. Because $O(\mathtt{FT}_i) \neq \mathtt{This}$, then it is either O or World, thus, $\theta(\mathtt{W}_i)$ is either NO or World. Finally note that $\mathtt{l} \prec_\theta \mathtt{NO}$ (because $\mathtt{l} \mapsto \mathtt{C<NO,} \ldots \mathtt{>}$), and we always have that $\mathtt{NO} \preceq_\theta \mathtt{World}$, thus $\mathtt{l} \prec_\theta \mathtt{World}$. According to subtyping rule S13, we have that $K \cup \{\mathtt{l}\}, \Gamma_{H'} \vdash \mathtt{U}_i \leq \mathtt{W}_i$, and from transitivity $\mathtt{V}_i \leq \mathtt{U}_i \leq \mathtt{W}_i$.

Rule R-FIELD-ACCESS According to R-FIELD-ACCESS

$$\frac{H[\mathtt{l}] = \mathtt{C<NO,NI>}(\overline{\mathtt{v}}) \qquad \mathit{fields}(\mathtt{C}) = \overline{\mathtt{f}}}{K \vdash H, \mathtt{l.f}_i \rightarrow H, \mathtt{v}_i}$$

We assumed that $K, \Gamma_H \vdash \mathtt{l.f}_i : \mathtt{T}"$, and we need to prove that $K, \Gamma_H \vdash \mathtt{v}_i : \grave{\mathtt{T}}$ and $K, \Gamma_H \vdash \grave{\mathtt{T}} \leq \mathtt{T}"$. If $\mathtt{v}_i = \mathtt{null}$ then we can choose $\grave{\mathtt{T}} = \mathtt{T}"$, otherwise $\mathtt{v}_i \neq \mathtt{null}$, and because the heap is well-typed for K, then $K, \Gamma_H \vdash \grave{\mathtt{T}} \leq \mathtt{T}"$.

Rule R-FIELD-ASSIGNMENT According to R-FIELD-ASSIGNMENT

$$\frac{\ldots}{K \vdash H, \mathtt{l.f}_i = \mathtt{v}' \rightarrow H', \mathtt{v}'}$$

We assumed that $K, \Gamma_H \vdash \mathtt{l.f}_i = \mathtt{v}' : \mathtt{T}$", and we need to prove that $K, \Gamma_{H'} \vdash \mathtt{v}' : \grave{\mathtt{T}}$ and $K, \Gamma_{H'} \vdash \grave{\mathtt{T}} \leq \mathtt{T}$". Because we did not add any new locations, we have $\Gamma_H = \Gamma_{H'}$. From T-FIELD-ASSIGNMENT, we have that $K, \Gamma_H \vdash \mathtt{v}' : \mathtt{T}$", i.e., $\grave{\mathtt{T}} = \mathtt{T}$".

Rule R-INVOKE According to R-INVOKE

$$\frac{H[\mathtt{l}] = \mathtt{C<NO,NI>(\ldots)} \qquad \mathit{mbody}(\mathtt{m}, \mathtt{C}) = \overline{\mathtt{x}}.\mathtt{e}'}{K \vdash H, \mathtt{l.m}(\overline{\mathtt{v}}) \rightarrow H, [\overline{\mathtt{v}}/\overline{\mathtt{x}}, \mathtt{l}/\mathtt{this}, \mathtt{l}/\mathtt{This}, \mathtt{NO}/\mathtt{O}, \mathtt{NI}/\mathtt{I}]\mathtt{e}'}$$

We assumed that

$$K, \Gamma_H \vdash \mathtt{e}" : \mathtt{T}" \qquad \mathtt{e}" = \mathtt{l.m}(\overline{\mathtt{v}}) \qquad \grave{\mathtt{e}} = [\overline{\mathtt{v}}/\overline{\mathtt{x}}, \mathtt{l}/\mathtt{this}, \mathtt{l}/\mathtt{This}, \mathtt{NO}/\mathtt{O}, \mathtt{NI}/\mathtt{I}]\mathtt{e}'$$

From T-INVOKE we have that

$$\mathit{mtype}(\mathtt{l}, \mathtt{m}, \mathtt{C<NO,NI>}) = \overline{\mathtt{T}} \rightarrow \mathtt{T}" \qquad K, \Gamma_H \vdash \overline{\mathtt{v}} : \overline{\mathtt{T}'} \qquad K, \Gamma_H \vdash \overline{\mathtt{T}'} \leq \overline{\mathtt{T}} \qquad \mathit{mguard}(\mathtt{m}, \mathtt{C}) = \mathtt{IG} \qquad K, \Gamma_H \vdash \mathtt{NI} \leq \mathtt{IG}$$

We know the method `m` was typed-checked in `C`, i.e.,

$$\begin{aligned} \Gamma &= \{\mathtt{I} : \mathtt{IG}, \overline{\mathtt{x}} : \overline{\mathtt{U}}, \mathtt{this} : \mathtt{C<O,I>}\} \\ \mathit{mtype}(\mathtt{m}, \mathtt{C}) &= \overline{\mathtt{U}} \rightarrow \mathtt{FT} \\ \emptyset, \Gamma &\vdash \mathtt{e}' : \mathtt{S} \\ \emptyset, \Gamma &\vdash \mathtt{S} \leq \mathtt{FT} \end{aligned}$$

From the definition of *mtype*:

$$\begin{aligned} \mathit{mtype}(\mathtt{l}, \mathtt{m}, \mathtt{C<NO,NI>}) &= \mathit{substitute}(\mathtt{l}, \mathtt{C<NO,NI>}, \mathit{mtype}(\mathtt{m}, \mathtt{C})) \\ \mathtt{T}" &= [\mathtt{NO}/\mathtt{O}, \mathtt{NI}/\mathtt{I}, \mathtt{l}/\mathtt{This}]\mathtt{FT} \\ \mathtt{T}_i &= [\mathtt{NO}/\mathtt{O}, \mathtt{NI}/\mathtt{I}, \mathtt{l}/\mathtt{This}]\mathtt{U}_i \end{aligned}$$

We need to prove that $K, \Gamma_H \vdash \grave{\mathtt{e}} : \grave{\mathtt{T}}$ and $K, \Gamma_H \vdash \grave{\mathtt{T}} \leq \mathtt{T}$", which follows immediately from Lem. 5.5.

□

Lemma 5.7. ***(Well-typed heap preservation)*** *For every closed expression* $\mathtt{e}" \neq \mathtt{v}$*, K, and H, if* $K, \Gamma_H \vdash \mathtt{e}" : \mathtt{T}"$ *and* $K \vdash H, \mathtt{e}" \rightarrow H', \grave{\mathtt{e}}$ *and H is well-typed for* $K \cup K(\mathtt{e}")$*, then* H' *is well-typed for* $K \cup K(\grave{\mathtt{e}})$*.*

Proof. Recall that a well-typed heap H satisfies: (i) there is a linear order $\preceq^T$ over $\mathrm{dom}(H)$ such that for every location `l`, $\theta(\mathtt{l}) = \mathtt{World}$ or $\theta(\mathtt{l}) \prec^T \mathtt{l}$, and $I(\mathtt{l}) = \mathtt{Mutable}$ or $\kappa(\mathtt{l}) \preceq^T \mathtt{l}$, and (ii) each non-null field location is a subtype of the declared field type. Recall also that from the definition of a heap H, every location `l` in H has the form: (iii) $\mathtt{l} \mapsto \mathtt{C<\underline{NO,NI}>}(\overline{\mathtt{v}})$.

Consider the congruence rules, such as

$$\frac{K \vdash H, \mathtt{e} \rightarrow H', \mathtt{e}'}{K \vdash H, \mathtt{e}.f \rightarrow H', \mathtt{e}'.f}$$

By the induction hypothesis H' is well-typed $K \cup K(\grave{\mathtt{e}})$.

The only rule that changes K is R-C1:

$$\frac{K \cup \{\mathtt{l}\} \vdash H', \mathtt{e} \rightarrow H", \mathtt{e}'}{K \vdash H, \mathtt{e;return\ l} \rightarrow H", \mathtt{e}'\mathtt{;return\ l}}$$

By induction $H"$ is well-typed for $(K \cup \{\mathtt{l}\}) \cup K(\mathtt{e}')$. We need to prove that $H"$ is well-typed for $K \cup K(\mathtt{e}'\mathtt{;return\ l})$. By definition of $K(\ldots)$, we have that $K \cup K(\mathtt{e}'\mathtt{;return\ l}) = K \cup \{\mathtt{l}\} \cup K(\mathtt{e}')$.

Rules R-FIELD-ACCESS and R-INVOKE do not change the heap.

Rule R-RETURN does not change the heap nor K, however $K(\grave{e}) = K(e'') \setminus \{l\}$, According to Lem. 4.1, the resulting heap H' is well-typed for $K \cup K(\grave{e})$.

Rule R-NEW creates a new object with `null` fields:

$$VI' = \begin{cases} \texttt{Immut}_l & \text{if } VI = \texttt{Immut} \text{ or } (VI = \texttt{Immut}_c \text{ and } c \notin K) \\ VI & \text{otherwise} \end{cases} \qquad l \notin \text{dom}(H) \qquad H' = H[l \mapsto \texttt{C<NO,}VI'\texttt{>(}\overline{\texttt{null}}\texttt{)}]$$

The fields of the new object are all `null`, thus fulfilling demand (ii).

We extend the linear order $\preceq^T$ by adding the new location `l` at the end. Its owner `NO` is either `World` or an existing object l', and either $VI = \texttt{Mutable}$ or $VI = \texttt{Immut}_{l'}$ (where l' is either an existing location or `l`), thus fulfilling demand (i).

Note that $e'' = \texttt{new C<}\underline{\texttt{NO,VI}}\texttt{>(}\overline{v}\texttt{)}$, and because e'' is closed, then we have that `NO` and `VI` do not contain `0`, `I`, nor `This`. And if $VI = \texttt{Immut}$ then it is substituted with $\texttt{Immut}_l$, thus fulfilling demand (iii). Finally, note that $\grave{e} = (\texttt{l.build(}\overline{v}\texttt{);return l})$, i.e., $K(\grave{e}) = K(e'') \cup \{l\}$. However, because `l` is a *new* location, it does not change existing subtype relations (it does not affect existing objects that do not refer to `l`). Therefore, H' is well-typed for $K \cup K(\grave{e})$.

Finally, in rule R-FIELD-ASSIGNMENT, $e'' = \texttt{l.f}_i = v'$, $\grave{e} = v'$, and $H' = H[l \mapsto \texttt{C<NO,NI>(}[v'/v_i]\overline{v}\texttt{)}]$. Note that $K(e'') = K(\grave{e}) = \{\}$, thus a if H' is well typed then it is well-typed for $K \cup K(\grave{e})$. Because the typing rule T-FIELD-ASSIGNMENT require that:

$$K, \Gamma_H \vdash \texttt{l.f} : T \qquad K, \Gamma_H \vdash v' : T' \qquad K, \Gamma_H \vdash T' \leq T$$

then the heap H' is well-typed for $K \cup K(\grave{e})$. □

What follows is our camera-ready SPLASH2010 (formerly known as OOPSLA) main technical track paper.

Ownership and Immutability in Generic Java

Yoav Zibin Alex Potanin Paley Li
Victoria University of Wellington
Wellington, New Zealand
yoav|alex|lipale@ecs.vuw.ac.nz

Mahmood Ali
Massachusetts Institute of Technology
Cambridge, MA, USA
mali@csail.mit.edu

Michael D. Ernst
University of Washington
Seattle, WA, USA
mernst@cs.washington.edu

Abstract

The Java language lacks the important notions of *ownership* (an object owns its representation to prevent unwanted aliasing) and *immutability* (the division into mutable, immutable, and readonly data and references). Programmers are prone to design errors, such as representation exposure or violation of immutability contracts. This paper presents *Ownership Immutability Generic Java* (OIGJ), a backward-compatible purely-static language extension supporting ownership and immutability. We formally defined a core calculus for OIGJ, based on Featherweight Java, and proved it sound. We also implemented OIGJ and performed case studies on 33,000 lines of code.

Creation of immutable cyclic structures requires a "*cooking phase*" in which the structure is mutated but the outside world cannot observe this mutation. OIGJ uses *ownership* information to facilitate creation of *immutable* cyclic structures, by safely prolonging the cooking phase even after the constructor finishes.

OIGJ is easy for a programmer to use, and it is easy to implement (flow-insensitive, adding only 14 rules to those of Java). Yet, OIGJ is more expressive than previous ownership languages, in the sense that it can type-check more good code. OIGJ can express the factory and visitor patterns, and OIGJ can type-check Sun's `java.util` collections (except for the `clone` method) without refactoring and with only a small number of annotations. Previous work required major refactoring of existing code in order to fit its ownership restrictions. Forcing refactoring of well-designed code is undesirable because it costs programmer effort, degrades the design, and hinders adoption in the mainstream community.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features; D.1.5 [*Programming Techniques*]: Object-oriented Programming

General Terms Experimentation, Languages, Theory

OOPSLA/SPLASH'10, October 17–21, 2010, Reno/Tahoe, Nevada, USA.

1. Introduction

This paper presents *Ownership Immutability Generic Java* (OIGJ), a simple and practical language extension that expresses both ownership and immutability information. OIGJ is purely static, without any run-time representation. This enables executing the resulting code on any JVM without run-time penalty. Our ideas, though demonstrated using Java, are applicable to any statically typed language with generics, such as C++, C#, Scala, and Eiffel.

OIGJ follows the *owner-as-dominator* discipline [1, 11, 32] where an object cannot leak beyond its owner: outside objects cannot access it. If an object owns its representation, then there are no aliases to its internal state. For example, a `LinkedList` should own all its `Entry` objects (but not its elements); entries should not be exposed to clients, and entries from different lists must not be mixed.

The keyword `private` does not offer such strong protection as ownership, because a careless programmer might write a public method that exposes a private object (a.k.a. representation exposure). Phrased differently, the name-based protection used in Java hides the variable but not the object, as opposed to ownership that ensures proper encapsulation. The key idea in ownership is that representation objects are nested and encapsulated inside the objects to which they belong. Because this nesting is transitive, this kind of ownership is also called *deep* ownership [12].

OIGJ is based on our previous type systems for ownership (OGJ [32]) and immutability (IGJ [39]). Although ownership and immutability may seem like two unrelated concepts, a design involving both enhances the expressiveness of each individual concept. On the one hand, immutability enhances ownership by relaxing owner-as-dominator to owner-as-modifier [23], i.e., it constrains modification instead of aliasing. On the other hand, the benefits of adding ownership on top of immutability have not been investigated before. One such benefit is easier creation of immutable cyclic data-structures by using ownership information.

Constructing an immutable object must be done with care. An object begins in the *raw* (not fully initialized) state and transitions to the *cooked* state [6] when initialization is complete. For an immutable object, field assignment is allowed only when the object is *raw*, i.e., the object cannot

be modified after it is *cooked*. An immutable object should not be visible to the outside world in its raw state because it would seem to be mutating. The challenge in building an immutable cyclic data-structure is that many objects must be raw simultaneously to create the cyclic structure. Previous work restricted cooking an object to the constructor, i.e., an object becomes cooked when its constructor finishes.

Our key observation is that *an object can be cooked when its owner's constructor finishes*. More precisely, in OIGJ, a programmer can choose between cooking an object until its constructor finishes, or until its owner becomes cooked. Because the object is encapsulated within its owner, the outside world will not see this cooking phase. By adding ownership information, we can prolong the cooking time to make it easier to create complex data-structures.

Consider building an immutable `LinkedList` (Sun's implementation is similar):

```
LinkedList(Collection<E> c) {
  this();
  Entry<E> succ = this.header, pred = succ.prev;
  for (E e : c)
  { Entry<E> entry = new Entry<E>(e,succ,pred);
    // It is illegal to change an entry after it is cooked.
    pred.next = entry;  pred = entry; }
  succ.prev = pred;
}
```

An immutable list contains immutable entries, i.e., the fields `next` and `prev` cannot be changed after an entry is cooked. In IGJ and previous work on immutability, an object becomes cooked after its constructor finishes. Because `next` and `prev` are changed after that time, this code is illegal. In contrast, in OIGJ, this code type-checks if we specify that the list (the `this` object) owns all its entries (the entries are the list's representation). The entries will become cooked when their owner's (the list's) constructor finishes, thus permitting the underlined assignments during the list's construction. Therefore, there was no need to refactor the constructor of `LinkedList` for the benefit of OIGJ type-checking.

Informally, OIGJ provides the following ownership and immutability guarantees. Let $\theta(o)$ denote the owner of o and let $\preceq_\theta$ denote the ownership tree, i.e., the transitive, reflexive closure of $o \preceq_\theta \theta(o)$. Phrased differently, $\theta(o)$ is the parent of o in the tree, and the top (biggest) node in the tree is the root `World`. We say that o_1 is *inside* o_2 iff $o_1 \preceq_\theta o_2$, and it is *strictly inside* if $o_1 \neq o_2$.

Ownership guarantee: An object o' can point to object o iff $o' \preceq_\theta \theta(o)$, i.e., o is owned by o' or by one of its owners.

Immutability guarantee: An immutable object cannot be changed after it is cooked.

Contributions The main contributions of this paper are:

Simplify ownership concepts OIGJ expresses ownership concepts without introducing new mechanisms, by using Java's underlying generic mechanisms. Specifically, owner-polymorphic methods and scoped regions [36] are implemented using *generic methods*, and existential owners [37] are implemented using *generic wildcards*. By contrast, previous work used new mechanisms that are orthogonal to generics.

No refactoring of existing code Java's collection classes (`java.util`) are properly encapsulated. We have implemented OIGJ, and verified the encapsulation by running the OIGJ type-checker without changing the source code (except for the `clone` method). Verifying Sun's `LinkedList` requires only 3 ownership annotations (see Sec. 4). Previous approaches to ownership or immutability required major refactoring of this codebase.

Flexibility As illustrated by our case study, OIGJ is more flexible and practical than previous type systems. For example, OIGJ can type-check the factory and visitor design patterns (see Sec. 2), but other *ownership* languages cannot [25]. Another advantage of OIGJ is that it uses ownership information to facilitate creating immutable cyclic structures by prolonging their cooking phase.

Formalization We define a Featherweight OIGJ (FOIGJ) calculus to formalize the concepts of raw/cooked objects and wildcards as owner parameters. We prove FOIGJ is sound and show our ownership and immutability guarantees.

Outline. Sec. 2 presents the OIGJ language. Sec. 3 defines the OIGJ formalization. Sec. 4 discusses the OIGJ implementation and the collections case study. Sec. 5 compares OIGJ to related work, and Sec. 6 concludes.

2. OIGJ Language

This section presents the OIGJ language extension that expresses both ownership and immutability information. We first describe the OIGJ syntax (Sec. 2.1). We then proceed with a `LinkedList` class example (Sec. 2.2), followed by the OIGJ typing rules (Sec. 2.3). We conclude with the factory (Sec. 2.4) and visitor (Sec. 2.5) design patterns in OIGJ.

2.1 OIGJ syntax

OIGJ introduces two new type parameters to each type, called the *owner parameter* and the *immutability parameter*. For simplicity of presentation, in the rest of this paper we assume that the special type parameters are at the beginning of the list of type parameters. We stress that generics in Java are erased during compilation to bytecode and do not exist at run time, therefore OIGJ does not incur any run-time overhead (nor does it support run-time casts).

In OIGJ, all classes are subtypes of the parameterized root type `Object<O,I>` that declares an owner and an immutability parameter. In OIGJ, the first parameter is the owner (O), and the second is the immutability (I). All subclasses must invariantly preserve their owner and immutability parameter. The owner and immutability parameters form two separate

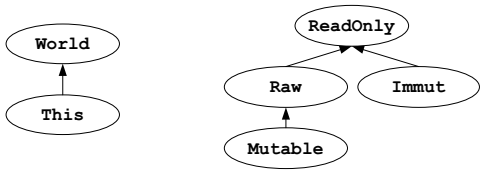

Figure 1. The type hierarchy of (a) ownership and (b) immutability parameters. `World` means the entire world can access the object, whereas `This` means that `this` owns the object and no one else can access it. The meaning of `Mutable`/`Immut` is obvious. A `ReadOnly` reference points to a mutable or immutable object, and therefore cannot be used to mutate the object. `Raw` represents an object under construction whose fields can be assigned.

hierarchies, which are shown in Fig. 1. These parameters cannot be extended, and they have no subtype relation with any other types. The subtyping relation is denoted by $\leq$, e.g., `Mutable` $\leq$ `ReadOnly`. Subtyping is invariant in the owner parameter and covariant in the immutability parameter. (See also paragraph ***Subtype relation*** in Sec. 2.3.)

Note that the *owner parameter* `O` is a *type*, whereas the *owner* of an object is an *object*. For example, if the owner parameter is `This`, then the owner is the object `this`. Therefore, the owner parameter (which is a type) at compile time corresponds to an owner (which is an object) at run time. (See also paragraph ***Owner vs. Owner-parameter*** below.)

OIGJ syntax borrows from *conditional Java* (cJ) [19], where a programmer can write method *guards*. A guard of the form `<X extends Y>? METHOD_DECLARATION` has a dual meaning: (i) the method is applicable only if the type argument that substitutes `X` extends `Y`, and (ii) the bound of `X` inside `METHOD_DECLARATION` changes to `Y`. The guards are used to express the immutability of `this`: a method receiver or a constructor result. For example, a method guarded with `<I extends Mutable>?` means that (i) the method is applicable only if the receiver is mutable and therefore (ii) `this` can be mutated inside the method.

Class definition example Fig. 2 shows an example of OIGJ syntax. A class definition declares the owner and immutability parameters (line 1); by convention we always denote them by `O` and `I` and they always extend `World` and `ReadOnly`. If the `extends` clause is missing from a class declaration, then we assume it extends `Object<O,I>`.

Immutability example Lines 2–4 show different kinds of immutability in OIGJ: immutable, mutable, and readonly. A readonly and an immutable reference may seem similar because neither can be used to mutate the referent. However, line 4 shows the difference between the two: a readonly reference may point to a mutable object. Phrased differently, a readonly reference may not mutate its referent, though the referent may be changed via an aliasing mutable reference.

```
1:class Foo<O extends World,I extends ReadOnly> {
2:  // An immutable reference to an immutable date.
    Date<O,Immut> imD = new Date<O,Immut>();
3:  // A mutable reference to a mutable date.
    Date<O,Mutable> mutD = new Date<O,Mutable>();
4:  // A readonly reference to any date. Both roD and imD cannot
    mutate their referent, however the referent of roD might be
    mutated by an alias, whereas the referent of imD is immutable.
    Date<O,ReadOnly> roD = ... ? imD : mutD;
5:  // A date with the same owner and immutability as this.
    Date<O,I> sameD;
6:  // A date owned by this; it cannot leak.
    Date<This,I> ownedD;
7:  // Anyone can access this date.
    Date<World,I> publicD;
8:  // Can be called on any receiver; cannot mutate this. The
    method guard "<...>?" is part of cJ's syntax [19].
    <I extends ReadOnly>? int readonlyMethod() {...}
9:  // Can be called only on mutable receivers; can mutate this.
    <I extends Mutable>? void mutatingMethod() {...}
10: // Constructor that can create (im)mutable objects.
    <I extends Raw>? Foo(Date<O,I> d) {
11:   this.sameD = d;
12:   this.ownedD = new Date<This,I>();
13:   // Illegal, because sameD came from the outside.
      // this.sameD.setTime(...);
14:   // OK, because Raw is transitive for owned fields.
      this.ownedD.setTime(...);
15:}}
```

Figure 2. An example of OIGJ syntax.

Java's type arguments are invariant (neither covariant nor contravariant), to avoid a type loophole [20], so line 4 is illegal in Java. Line 4 is legal in OIGJ, because OIGJ safely allows covariant changes in the immutability parameter (but not in the owner parameter). OIGJ *restricts* Java by having additional typing rules, while at the same time OIGJ also *relaxes* Java's subtyping relation. Therefore, neither OIGJ nor Java subsumes the other, i.e., a legal OIGJ program may be illegal in Java (and vice versa). However, because generics are erased during compilation, the resulting bytecode can be executed on any JVM.

The immutability of `sameD` (line 5) depends on the immutability of `this`, i.e., `sameD` is (im)mutable in an (im)mutable `Foo` object. Similarly, the owner of `sameD` is the same as the owner of `this`.

Ownership example Lines 5–7 show three different owner parameters: `O`, `This`, and `World`. The owner parameter is invariant, i.e., the subtype relation preserves the owner parameter. For instance, the types on lines 5–7 have no subtype relation with each other because they have different owner parameters.

Reference `ownedD` cannot leak outside of `this`, whereas references `sameD` and `publicD` can potentially be accessed by

anyone with access to this. Although sameD and publicD can be accessed by the same objects, they cannot be stored in the same places: publicD can be stored anywhere on the heap (even in a static public variable) whereas sameD can only be stored inside its owner.

We use $O(\ldots)$ to denote the function that takes a type or a reference, and returns its owner parameter; e.g., $O(\texttt{ownedD}) = \texttt{This}$. Similarly, function $I(\ldots)$ returns the immutability parameter; e.g., $I(\texttt{ownedD}) = \texttt{I}$. We say that an object o is *this-owned* (i.e., owned by this) if $O(\texttt{o}) = \texttt{This}$; e.g., ownedD is this-owned, but sameD is not. OIGJ prevents leaking this-owned objects by requiring that this-owned fields (and methods with this-owned arguments or return-type) can only be used via this. For example, this.ownedD is legal, but foo.ownedD is illegal.

Owner vs. owner-parameter Now we explain the connection between the *owner parameter* $O(o)$, which is a generic type parameter at *compile time*, and the *owner* $\theta(o)$, which is an object at *run time*. This is an owner parameter that represents an owner that is the current this object, and World represents the root of the ownership tree (we treat World both as a type parameter and as an object that is the root of the ownership tree). Formally, if $O(o) = \texttt{This}$ then $\theta(o) = \texttt{this}$, if $O(o) = \texttt{O}$ then $\theta(o) = \theta(\texttt{this})$, and if $O(o) = \texttt{World}$ then $\theta(o) = \texttt{World}$. Two references (in the same class) with the same owner parameter (at compile time) will point to objects with the same owner (at run time), i.e., $O(o_1) = O(o_2)$ implies $\theta(o_1) = \theta(o_2)$.

Finally, recall the **Ownership guarantee**: o' can point to o iff $o' \preceq_\theta \theta(o)$. By definition of $\preceq_\theta$, we have that for all o: (i) $o \preceq_\theta o$, (ii) $o \preceq_\theta \theta(o)$, and (iii) $o \preceq_\theta \texttt{World}$. By part (iii), if $\theta(o) = \texttt{World}$ then anyone can point to o. On lines 5–7, we see that this can point to ownedD, sameD, publicD, whose owner parameters are This, O, World, and whose owners are this, θ(this), World. This conforms with the ownership guarantee according to parts (i), (ii), and (iii), respectively. More complicated pointing patterns can occur by using multiple owner parameters, e.g., an entry in a list can point to an element owned by the list's owner, such as in List<This,I,Date<O,I>>.

There is a similar connection between the immutability type parameter (at compile time) and the object's immutability (at run time). Immutability parameter Mutable or Immut implies the object is mutable or immutable (respectively), ReadOnly implies the referenced object may be either mutable or immutable and thus the object cannot be mutated through the read-only reference. Raw implies the object is still raw and thus can still be mutated, but it might become immutable after it is cooked.

Method guard example Lines 8 and 9 show a readonly and a mutating method. These methods are *guarded* with <...>?. Conditional Java (cJ) [19] extends Java with such guards (a.k.a. conditional type expressions). Note that cJ changed Java's syntax by using the question mark in the guard <...>?. The exposition in this paper uses cJ for convenience. However, our implementation of OIGJ (Sec. 4) uses type annotations [15] without changing Java's syntax, for conciseness and compatibility with existing tools and code bases.

A guard such as <T extends U>? METHOD_DECLARATION has a dual purpose: (i) the method is included only if T extends U, and (ii) the bound of T is U inside the method. In our example, the guard on line 9 means that (i) this method can only be called on a Mutable receiver, and (ii) inside the method the bound of I changes to Mutable. For instance, (i) only a mutable Foo object can be a receiver of mutatingMethod, and (ii) field sameD is mutable in mutatingMethod. cJ also ensures that the condition of an overriding method is equivalent or weaker than the condition of the overridden method.

IGJ [39] used *declaration annotations* to denote the immutability of this. In this paper, OIGJ uses cJ to reduce the number of typing rules and handle inner classes more flexibly.[1] OIGJ does not use the full power of cJ: it only uses guards with immutability parameters. Moreover, we modified cJ to treat guards over constructors in a special way described in the **Object creation rule** of Fig. 4.

To summarize, on lines 8–10 we see three guards that change the bound of I to ReadOnly, Mutable, and Raw, respectively. Because the bound of I is already declared on line 1 as ReadOnly, the guard on line 8 can be removed.

Constructor example The constructor on line 10 is guarded with Raw, and therefore can create both mutable and immutable objects, because all objects start their life cycle as raw. This constructor illustrates the interplay between *ownership* and *immutability*, which makes OIGJ more expressive than previous work on immutability. OIGJ uses ownership information to prolong the *cooking phase* for owned objects: the cooking phase of this-owned fields (ownedD) is longer than that of non-owned fields (sameD). This property is critical to type-check the collection classes, as Sec. 2.2 will show.

Consider the following code:

```
class Bar<O extends World,I extends ReadOnly>
  { Date<O,Immut> d = new Date<O,Immut>();
    Foo<O,Immut> foo = new Foo<O,Immut>(d); }
```

Recall our **Immutability guarantee**: an immutable object cannot be changed after it is *cooked*. A This-owned object is cooked when its owner is cooked (e.g., foo.ownedD). Any other object is cooked when its constructor finishes (e.g., d and foo). The intuition is that ownedD cannot leak and so the outside world cannot observe this longer cooking phase, whereas d is visible to the world after its constructor finishes and must not be mutated further. The constructor on lines 10–15 shows this difference between the assignments to sameD

[1] Our implementation uses *type annotations* to denote immutability of this. A type annotation @Mutable on the receiver is similar to a cJ <I extends Mutable>? construct, but it separates the distinct roles of the receiver and the result in inner class constructors.

(line 11) and to ownedD (line 12): sameD can come from the outside world, whereas ownedD must be created inside this. Thus, sameD cannot be further mutated (line 13) whereas ownedD can be mutated (line 14) until its owner is cooked.

An object in a raw method, whose immutability parameter is I, is still considered raw (thus the modified body can still assign to its fields or call other raw methods) iff the object is this or this-owned. Informally, we say that Raw is *transitive* only for this or this-owned objects. For example, the receiver of the method call sameD.setTime(...) is not this nor this-owned, and therefore the call on line 13 is illegal; however, the receiver of ownedD.setTime(...) is this-owned, and therefore the call on line 14 is legal.

2.2 `LinkedList` example

Fig. 3 shows an implementation of LinkedList in OIGJ that is similar in spirit to Sun's implementation. We explain this example in three stages: (i) we first explain the data-structure, i.e., the fields of a list and its entries (lines 1–6), (ii) then we discuss the Raw constructors that enable creation of immutable lists (lines 7–24), and (iii) finally we dive into the complexities of inner classes and iterators (lines 27–53).

LinkedList data-structure A linked list has a header field (line 6) pointing to the first entry. Each entry has an element and pointers to the next and prev entries (line 3). We explain first the immutability and then the ownership of each field.

Recall that we implicitly assume that O extends World and that *I extends ReadOnly on lines 1, 5, 35 and 49.

An (im)mutable list contains (im)mutable entries, i.e., the entire data-structure is either mutable or immutable as a whole. Hence, all the fields have the same immutability I. The underlying generic type system propagates the immutability information without the need for special typing rules.

Next consider the ownership of the fields of LinkedList and Entry. This on line 6 expresses that the reference header points to an Entry owned by this, i.e., the entry is encapsulated and cannot be aliased outside of this. O on line 3 expresses that the owner of next is the same as the owner of the entry, i.e., a linked-list owns *all* its entries. Note how the generics mechanism propagates the owner parameter, e.g., the type of this.header.next.next is Entry<This,I,E>. Thus, the owner of all entries is the this object, i.e., the list.

Finally, note that the field element has no immutability nor owner parameters, because they will be specified by the client that instantiates the list type, e.g.,
LinkedList<This,Mutable,Date<World,ReadOnly>>

Immutable object creation A constructor that is making an immutable object must be able to set the fields of the object. It is not acceptable to mark such constructors as Mutable, which would permit arbitrary side effects, possibly including making mutable aliases to this. OIGJ uses a fourth kind of immutability, Raw, to permit constructors to perform limited side effects without permitting modification of immutable

```
 1: class Entry<O,I,E> {
 2:   E element;
 3:   Entry<O,I,E> next, prev;
 4: }
 5: class LinkedList<O,I,E> {
 6:   Entry<This,I,E> header;
 7:   <I extends Raw>? LinkedList() {
 8:     this.header = new Entry<This,I,E>();
 9:     header.next = header.prev = header;
10:   }
11:   <I extends Raw>? LinkedList(
12:                 Collection<?,ReadOnly,E> c) {
13:     this();  this.addAll(c);
14:   }
15:   <I extends Raw>? void addAll(
16:                 Collection<?,ReadOnly,E> c) {
17:     Entry<This,I,E> succ = this.header,
18:                     pred = succ.prev;
19:     for (E e : c) {
20:       Entry<This,I,E> en=new Entry<This,I,E>();
21:       en.element=e; en.next=succ; en.prev=pred;
22:       pred.next = en;  pred = en;  }
23:     succ.prev = pred;
24:   }
25:   int size() {...}
26:   // iterator is a generic method; this is not a cJ guard:
27:   <ItrI extends ReadOnly> Iterator<O,ItrI,I,E>
28:                 iterator() {
29:     return this.new ListItr<ItrI>();
30:   }
31:   void remove(Entry<This,Mutable,E> e) {
32:     e.prev.next = e.next;
33:     e.next.prev = e.prev;
34:   }
35:   class ListItr<ItrI> implements
36:         Iterator<O,ItrI,I,E> {
37:     Entry<This,I,E> current;
38:     <ItrI extends Raw>? ListItr() {
39:       this.current = LinkedList.this.header;
40:     }
41:     <ItrI extends Mutable>? E next() {
42:       this.current = this.current.next;
43:       return this.current.element;
44:     }
45:     <I extends Mutable>? void remove() {
46:       LinkedList.this.remove(this.current);
47:     }
48: } }
49: interface Iterator<O,ItrI,CollectionI,E> {
50:   boolean hasNext();
51:   <ItrI extends Mutable>? E next();
52:   <CollectionI extends Mutable>? void remove();
53: }
```

Figure 3. LinkedList<O,I,E> in OIGJ.

objects. Raw represents a partially-initialized *raw* object that can still be arbitrarily mutated, but after it is cooked (fully initialized), then the object might become immutable. The constructors on lines 7 and 11 are guarded with Raw, and therefore can create both mutable and immutable lists.

Objects must not be captured in their raw state to prevent further mutation after the object is cooked. If a programmer could declare a field, such as Date<O,Raw>, then a raw date could be stored there, and later it could be used to mutate a cooked immutable date. Therefore, a programmer can write the Raw type only after the extends keyword, but *not* in any other way. As a consequence, in a Raw constructor, this can only escape as ReadOnly.

Recall that an object becomes *cooked* either when its constructor finishes or when its owner is *cooked*. The entries of the list (line 6) are this-owned. Indeed, the entries are mutated after their constructor finished, but before the list is cooked, on lines 9, 22, and 23. This shows the power of combining immutability and ownership: we are able to create immutable lists *only* by using the fact that the list owns its entries. If those entries were *not* owned by the list, then this mutation of entries might be visible to the outside world, thus breaking the guarantee that an immutable object never changes. By enforcing ownership, OIGJ ensures that such illegal mutations cannot occur.

OIGJ requires that all access and assignment to a this-owned field must be done via this. For example, see header, on lines 8, 9, 17, and 39. In contrast, fields next and prev (which are not this-owned) do not have such a restriction, as can be seen on lines 32–33.

Iterator implementation and inner classes An *iterator* has an underlying *collection*, and the immutability of these two objects might be different. For example, you can have

- a mutable iterator over a mutable collection (the iterator supports both remove() and next()),
- a mutable iterator over a readonly/immutable collection (the iterator supports next() but not remove()), or
- a readonly iterator over a mutable collection (the iterator supports remove() but not next(), which can be useful if you want to pass an iterator to a method that may not advance the iterator but may remove the current element).

Consider the Iterator<O,ItrI,CollectionI,E> interface defined on lines 49–53, and used on lines 27 and 36. ItrI is the iterator's immutability, whereas CollectionI is intended to be the underlying collection's immutability (see on line 36 how the collection's immutability I is used in the place of CollectionI). Line 51 requires a mutable ItrI to call next(), and line 52 requires a mutable CollectionI to call remove().

Inner class ListItr (lines 35–48) is the implementation of Iterator for list. Its full name is LinkedList<O,I,E>.ListItr<ItrI>, and on line 35 it extends Iterator<O,ItrI,I,E>. It reuses the owner parameter O from LinkedList, but declares a new immutability parameter ItrI. An inner class, such as ListItr<ItrI>, only declares an immutability parameter because it inherits the owner parameter from its outer class. ListItr and LinkedList have the same owner O, but different immutability parameters (ItrI for ListItr, and I for LinkedList). ListItr must inherit LinkedList's owner because it directly accesses the (this-owned) representation of LinkedList (line 39), which would be illegal if their owner was different. For example, consider the types of this and LinkedList.this on line 39:

```
Iterator<O,ItrI,...> thisIterator = this;
LinkedList<O,I,...> thisList = LinkedList.this;
```

Because line 38 sets the bound of ItrI to be Raw, this can be mutated. By contrast, the bound of I is ReadOnly, so LinkedList.this cannot.

An inner class must have a distinct immutability parameter, but it must reuse the owner parameter of its outer class. We could have several This types, e.g., LinkedList.This vs. ListItr.This, but this would complicate the typing rules.

Finally, consider the creation of a new *inner* object on line 29 using this.new ListItr<ItrI>(). This expression is type-checked both as a method call (whose receiver is this) and as a constructor call. Observe that the bound of ItrI is ReadOnly (line 27) and the guard on the constructor is Raw (line 38), which is legal because a Raw constructor can create both mutable and immutable objects.

2.3 OIGJ typing rules

Fig. 4 contains all the OIGJ typing rules. We now discuss each rule. Sec. 3 presents a formal type system based on a simplified version of these rules. Some of the rules are identical to those found in OGJ [32] and IGJ [39] (see Sec. 5 for a comparison with OIGJ).

Ownership nesting Consider the following example:

```
List<This,I,Date<World,I>> l1; // Legal nesting
List<World,I,Date<This,I>> l2; // Illegal!
```

Definition of l2 has illegal ownership nesting because owned dates might leak, e.g., we can store l2 in this variable:

```
public static Object<World,ReadOnly> publicAliasToL2;
```

On the one hand, types in OIGJ may have multiple owner parameters, e.g., the type of l1 has two owner parameters (This and World). On the other hand, an object may only have a single owner at run time. For example, the type of l1 will correspond at run time to a list that is owned by this while its elements are owned by World, and observe that this is always inside World.

Recall that an owner o_1 is *inside* o_2 iff o_1 is a descendant in the ownership tree of o_2, i.e., $o_1 \preceq_\theta o_2$. We extend this definition from owners to owner parameters as follows: given two owner parameters O_1 and O_2 in the same type, then O_1 is *inside* O_2 iff in any possible execution, these owner parameters correspond to some owners o_1 and o_2 (respectively) where $o_1 \preceq_\theta o_2$. For example, This is inside O, and any owner parameter is inside World.

Ownership nesting The first owner parameter of type `T` must be inside any other owner parameter in `T`.

Field access Field access `o.f` is legal iff $\underline{O(\texttt{f}) = \texttt{This} \Rightarrow \texttt{o} = \texttt{this}}$.

Field assignment Field assignment `o.f=...` is legal iff (i) $I(\texttt{o}) \leq \texttt{Raw}$, and (ii) $\big(I(\texttt{o}) = \texttt{Raw} \Rightarrow (\texttt{o} = \texttt{this}$ or $O(\texttt{o}) = \texttt{This})\big)$, and (iii) field access `o.f` is legal.

Method invocation Consider method $\texttt{T}_0\ \texttt{m(T}_1,\ldots,\texttt{T}_n\texttt{)}$. The invocation `o.m(...)` is legal iff (i) $\underline{O(\texttt{T}_i) = \texttt{This} \Rightarrow \texttt{o} = \texttt{this}}$ for $i = 0,\ldots,n$, and (ii) $I(\texttt{m}) = \texttt{Raw}$ implies field assignment part (ii).

cJ's method guard (i) An invocation `o.m(...)` is legal if the type of `o` satisfies the guard of `m`. (ii) When typing method `m`, the bound of type variables that appear in the guard changes to their bound in the guard. (iii) The guard of an overriding method is equivalent or weaker than that of the overridden method.

Inner classes An inner class reuses the owner parameter of the outer class. However, it has a distinct immutability parameter.

Invariant The programmer marks each type parameter as invariant or covariant. An immutability parameter is always covariant, whereas an owner parameter is always invariant.

A type parameter must be invariant if it is used in a superclass that contains `Mutable`, a field that contains `Mutable` but is not `this`-owned, or in the position of another invariant type parameter.

Same-class subtype relation Let $\texttt{C<X}_1,\ldots,\texttt{X}_n\texttt{>}$ be a class. Type $\texttt{S} = \texttt{C<S}_1,\ldots,\texttt{S}_n\texttt{>}$ is *a subtype of* $\texttt{T} = \texttt{C<T}_1,\ldots,\texttt{T}_n\texttt{>}$, written as $\texttt{S} \leq \texttt{T}$, iff $(\texttt{S} = \texttt{T})$ or $\Big($(all immutability parameters $\texttt{T}_j$ are either `ReadOnly` or `Immut`), and for $i = 1,\ldots,n$, $\big(\texttt{S}_i = \texttt{T}_i$ or $(\texttt{S}_i \leq \texttt{T}_i$ and $\texttt{X}_i$ is covariant in `C`$)\big)\Big)$.

Erased signature If method $\texttt{m}'$ overrides a readonly/immutable method `m`, then the erased signatures of $\texttt{m}'$ and `m`, excluding invariant type parameters, must be identical. (The *erased signature* of a method is obtained by replacing type parameters with their bounds.)

Object creation A constructor cannot have any `this`-owned arguments. Furthermore, `new SomeClass<X,...>(...)` is legal iff the constructor's guard `<I extends Y>?` satisfies: `Y` = `Mutable` and `X` = `Mutable`, or `Y` = `Raw`.

Generic Wildcards OIGJ prohibits using a generic wildcard (`?`) in the position of the immutability parameter. For the owner parameter, OIGJ prohibits using a wildcard in a field or in a method return type, but permits it for stack variables and method parameters.

`Raw` parameter `Raw` can only be used after the `extends` keyword. It cannot be used in the position of a generic parameter.

Fresh owners A *fresh owner* is a method owner parameter that is not used in the method signature. It is a descendant in the ownership tree of all other owners in scope.

Static context `This` cannot be used in a static context, i.e., in static methods or fields.

Figure 4. All the OIGJ typing rules (beyond those of Java), in English. Also see Sec. 3 for a formalization. Underlined sentences show similarities among the rules.

OIGJ requires that owner parameters are properly nested, i.e., that the first owner parameter of type `T` is inside any other owner parameter in `T`. To enforce this rule, OIGJ maintains ordering constraints among owner parameters in the same way as described in OGJ [32].

Field access This rule enforces ownership: `this`-owned fields can be assigned only via `this`. In Fig. 3, note that all accesses and assignments to `header` are done via `this`.

Field assignment Assigning to a field should respect both immutability and ownership constraints. Part (i) of the rule enforces immutability constraints: a field can be assigned only by a `Mutable` or `Raw` reference. Part (ii) ensures `Raw` is transitive only for `this` or `this`-owned objects. Part (iii) enforces ownership constraints as in field access.

For example, consider the assignments on lines 8 and 9 of Fig. 3. Note that the bound of `I` is `Raw`, thus the assignments satisfy part (i). Part (ii) holds, i.e., `Raw` is transitive in the first assignment because the target object is `this` and in the second assignment because it is `this`-owned (the type of `this.header` is `Entry<`$\underline{\texttt{This}}$`,I,E>`). Finally, part (iii) holds in the first assignment because `header` was assigned via `this`, and in the second assignment because field `next` (`Entry<`$\underline{\texttt{O}}$`,I>`) is not `this`-owned.

Method invocation Method invocation is handled in the same way as field access/assignment: parts (i) and (ii) are similar to field access and field assignment part (ii). For example, consider the following method: `R m(A a) { ... }` Then, the method call `o.m(e)` is handled as if there is an assignment to a field of type `A`, and the return value is typed as if there was an access to a field of type `R`. Note that regarding the transitivity of `Raw`, we check both the immutability of the receiver object ($I(\texttt{o})$) and that of the method, i.e., its guard ($I(\texttt{m})$). If both are `Raw`, then we require that `o` is either `this` or `this`-owned.

Inner classes An *inner class* is a non-static nested class, e.g., iterators in `java.util` are implemented using inner classes. An inner class reuses the owner parameter of the outer class, i.e., the inner object is seen as an extension of the outer object. However, it has a distinct immutability parameter. Therefore, both `this` and `OuterClass.this` are treated identically by the typing rules that involve ownership.

Nested classes that are *static* can be treated the same as normal classes.

Invariant A user can annotate a type parameter `X` in class `C` with `@InVariant` to prevent covariant changes, in which case we say that `X` is invariant. Otherwise we say that `X` is covariant. An immutability parameter must be covariant, or else a mutable reference could not be a receiver when calling a readonly method. An owner parameter must be invariant, because the owner of an object cannot change.

A type parameter must be invariant if it is used in a field/superclass that contains `Mutable`, or if the erased signature differs. For example, if a class has a field of type `Foo<O,Mutable,X>`, then `X` must be invariant (the owner parameter `O` is always invariant).

Subtype relation Java is *invariant* in generic arguments, i.e., it prohibits *covariant* (or contravariant) changes. `Vector<Integer>` is not a subtype of `Vector<Object>`. If it were, then mutating a `Vector<Integer>` by inserting, e.g., a `String`, breaks type-safety.

OIGJ permits covariant changes for non-mutable references because the object cannot be mutated in a way that is not type-safe. OIGJ's subtyping rules includes Java's subtyping rules, therefore OIGJ's subtype relation is a superset of Java's subtype relation. If mutation is disallowed, OIGJ's subtyping rule allows covariant changes in other type parameters, within the *same class*. For example, `List<O,ReadOnly,Integer>` is a subtype of `List<O,ReadOnly,Number>`. Note that covariance is allowed iff *all* immutability parameters of the supertype are `ReadOnly` or `Immut`, e.g., `Iterator<O,ReadOnly,Mutable,Integer>` is *not* a subtype of `Iterator<O,ReadOnly,Mutable,Number>`, but it is a subtype of `Iterator<O,ReadOnly,ReadOnly,Number>`.

Erased signature When the erased signature of an overriding method differs from the overridden method, the normal `javac` compiler inserts a *bridge method* to cast the arguments to the correct type [7]. Such bridge methods work correctly only under the assumptions that subtyping is invariant. For example, consider an integer comparator `intComp` that implements `Comparable<Integer>`. If `Comparable<Integer>` were a subtype of `Comparable<Object>`, then we could pass a `String` to `intComp`'s implementation of `compareTo(Integer)`: `((Comparable<Object>)intComp).compareTo("a")`

OIGJ requires that the *erased signature* of an overriding method remains the same (excluding invariant parameters) if the overridden method is either readonly or immutable. For example, the erased signature of `compareTo` in `intComp` differs from the one in the interface `Comparable<O,I,X>`. Therefore, this rule requires that the type parameter `X` must be invariant:

```
interface Comparable<O,I, @InVariant X> {
  int compareTo(X o); }
```

Object creation A constructor should not have any `this`-owned parameters, because `this`-owned objects can only be created inside `this`.

Recall that the immutability of a constructor (or any method in general) is defined to be the bound of the immutability parameter in that constructor, e.g., a mutable constructor has the guard `<I extends Mutable>?`. Recall that cJ prohibits calling a `Raw` constructor to create an `Immut` object because the guard is not satisfied: `Immut` is not a subtype of `Raw`. OIGJ changed cJ and treats constructor calls using this object creation rule: a `Raw` constructor can create any object (mutable and immutable). A `Mutable` constructor can only create `Mutable` objects. A constructor cannot be `Immut` or `ReadOnly`, so that it is able to assign to the fields of `this`.

Generic wildcards OIGJ uses Java's existing generic wildcard syntax (`?`) to express existential owners [8, 29, 37]. A programmer can use existential owners when the exact owner of an object is unknown. One motivation for existential owners is the downcast performed in the `equals` method [37].

Consider the following two casts in normal Java:

```
boolean equals(Object o)
{  List<?> l = (List<?>)o; // OK
   List<Object> l = (List<Object>)o; } // Warning!
```

The second cast is a warning since erasure makes it impossible to check at run time that the generic parameter is `Object`.

OIGJ prohibits wildcards on the owner parameter of *fields*, e.g., `Date<?,ReadOnly> field`, because one can declare a static field of that type and store a `this`-owned date, thus breaking owner-as-dominator. Wildcards on a method return type are also prohibited because they can be used to leak `this`-owned fields. However, wildcards on *stack variables* (method parameters or local variables) are allowed.

Note that the immutability parameter is covariant, and therefore there is no need to use a wildcard for immutability. For example, consider the `DateList` class, which is parameterized by its owner parameter (`O`) and the dates' owner parameter (`DO`):

```
class DateList<O,I,DO extends World> {
 boolean equals(Object<?,ReadOnly> o)
 { DateList<?,ReadOnly,?> l =
     // No need to check ownership or immutability at run time.
     (DateList<?,ReadOnly,?>) o;
   return listEquals(l); }
 <O2 extends World,DO2 extends World> boolean
    listEquals(DateList<O2,ReadOnly,DO2> l) {...}
}
```

Method `listEquals` shows that it is possible to name the existential owner—the unknown list's owner parameter is `O2` and the unknown dates' owner parameter is `DO2`. Phrased differently, the two wildcards in `DateList<?,ReadOnly,?>` are now named `DateList<O2,ReadOnly,DO2>`.

Recall that Java's generics can be bypassed by using reflection or raw types such as `List`. Similarly, one can bypass OIGJ when using these features.

`Raw` *parameter* `Raw` can only be used after the `extends` keyword. For example, it is prohibited to write `Date<O,Raw>`. If this was possible, then such a date could leak from a `Raw` constructor that is building an immutable object resulting in an alias that could mutate such immutable object.

Fresh owner A *fresh owner* is a method owner parameter that is not used in the method signature. In OIGJ, a fresh owner expresses *temporary ownership* within the method. This allows a method to create stack-local objects with access to any object visible at the point of creation, but with a guarantee that stack-locals will not leak. Hence, stack-local objects can be garbage-collected when the method returns. For example, consider a method that deserializes a `ByteStream` by creating a temporary `ObjectStream` wrapper:

```
<O,TmpO> void deserialize(ByteStream<O> bs) {
  ObjectStream<TmpO,ByteStream<O>> os = ... }
```

Note that `TmpO` is a fresh owner, whereas `O` is not. Because `TmpO` is strictly inside other owner parameters such as `O`, there

cannot be any aliases from bs to os. In fact, os can only be referenced from other stack-local objects, and therefore, when the method returns, os can be garbage-collected.

Technically, *a fresh owner is strictly inside all other non-fresh owners in scope*, to make sure it cannot exist after the method returns. (Multiple fresh owners are incomparable with each other.) Because a fresh owner is inside several other owners that might be incomparable in the ownership tree, the ownership structure is a DAG rather than a tree.

To type-check temporary ownership and DAG ownership structures, OIGJ adopts Wrigstad's *scoped confinement* [36] ownership model, in which the fresh owners are owned by the current *stack-entry*. Briefly stated, each method invocation pushes a new stack-entry (the first stack-entry corresponds to the static main method), which is the root of a new ownership tree. Objects in this new tree may point to objects in previous trees, but not vice versa.

Static context This represents that an object is owned by this, and so OIGJ prohibits using it in a static context, such as static fields or methods. Static fields can use the owner parameter World, and static methods can also use generic method parameters extending World. For example, method:

```
static <LO extends World,E> void sort(
                    List<LO,Mutable,E> l) { ... }
```

is parameterized by the list's owner LO.

2.4 Factory method design pattern

The *factory method pattern* [17] is a creational design pattern for creating objects without specifying the exact class of the object that will be created. The solution is to define an interface with a method for creating an object. Implementers can override the method to create objects of a derived type.

The challenge of the factory method pattern with respect to *ownership* [25] is that the point of *creation* and *usage* are in different classes, and the created object must be owned by its user. Previous work makes a newly-created object be owned by its creator, and then changes the ownership after the fact via sophisticated ownership transfer mechanisms [24] using capture and release.

In OIGJ's approach, an object has its final owner from its moment of creation. When requesting creation of a new object, the client of the factory also specifies the owner. The type-checker ensures that the created object cannot be *captured* (stored in a location that would require a different owner) in the process. Specifically, a *generic factory method* can abstract over the owner (and immutability) parameter of the constructed object. The underlying generics mechanism finds the correct generic method arguments.

We will show how to use the factory method pattern in the context of *synchronized lists*. Consider this client code:

```
b = new LinkedList<T>();
l = Collections.synchronizedList(b);
```

The documentation of Collections.synchronizedList states: "*In order to guarantee serial access, it is* critical *that all access to the backing list is accomplished through the returned list.*" That means that there might be concurrency problems if one accidentally uses the backing list b instead of l.

```
 1: class SafeSyncList<O,I,E> implements List<O,I,E> {
 2:   List<This,I,E> l;
 3:   <I extends Raw>? SafeSyncList(
 4:                       Factory<?,ReadOnly,E> f)
 5:   { List<This,I,E> b = f.create();
 6:     l = Collections.synchronizedList(b); }
 7:   ... // delegate methods to l
 8: }
 9: class Collections<O,I> {
10:   // Sun's original implementation, augmented only by O2 and I2
11:   static <O2,I2,E> List<O2,I2,E>
12:    synchronizedList(List<O2,I2,E> list) { ... }
13: }
14: interface Factory<O,I,E>
15: {  <O2,I2> List<O2,I2,E> create();  }
16: class LinkedListFactory<O,I,E> implements
17:                 Factory<O,I,E> {
18:   <O2,I2> List<O2,I2,E> create() {
19:     return new LinkedList<O2,I2,E>();
20: } }
```

Figure 5. Factory method design pattern in OIGJ. OIGJ guarantees that the backing list b (line 5) is never accessed directly, e.g., it cannot be captured on line 19.

So, you want to own a list l, which is backed by another list b. The challenge is that b should be owned by l (and not by you), in order to guarantee that you do not accidentally access b directly and comprise thread-safety.

Fig. 5 shows how owner-as-dominator can ensure that the backing list b has no outside aliases. This solution avoids refactoring of existing Java code by delegating calls to the synchronized list l. Specifically, class SafeSyncList (lines 1–8) owns both the list l (line 2) and the backing list b (line 5). A factory method is used on lines 3–6.

The Factory interface is defined on lines 14–15. The owner and immutability of the Factory is irrelevant because it only has a readonly method. However, the newly created list has a generic owner and immutability, which are statically unknown at the creation point (line 15). The generics mechanism fills in the correct generic arguments from the usage point (line 6) to the actual creation point (line 19). Note that the factory implementation cannot capture an alias to the newly created list on line 19, because its owner parameter O2 is a generic method parameter that cannot be used in fields.

To conclude, one can use SafeSyncList instead of using Sun's unsafe synchronizedList, and be certain no one else can access the backing list. All this was achieved using *generic factory methods* on lines 15 and 18.

```
1: interface Visitor<O,I,NodeO,NodeI> {
2:  <I extends Mutable>? void
3:    visit(Node<NodeO,NodeI> n);
4: }
5: class Node<O,I> {
6:  void accept(Visitor<?,Mutable,O,I> v)
7:  {  v.visit(this)  }
8: }
9: // Visiting a readonly node hierarchy.
10: Node<This,ReadOnly> readonlyNode = ...;
11: readonlyNode.accept( new
12:  Visitor<World,Mutable,This,ReadOnly>() {
13:   <I extends Mutable>? void
14:     visit(Node<This,ReadOnly> n)
15:     { ... // Can mutate the visitor, but not the nodes.  }
16:  });
17: // Visiting a mutable node hierarchy.
18: Node<This,Mutable> mutableNode = ...;
19: mutableNode.accept( new
20:  Visitor<World,Mutable,This,Mutable>() {
21:   <I extends Mutable>? void
22:     visit(Node<This,Mutable> n)
23:     { ... // Can mutate the visitor and the nodes.  }
24:  });
```

Figure 6. Visitor pattern in OIGJ. The `Node`'s ownership and immutability are underlined. We omit the `extends` clause for generic parameters, e.g., we assume that `NodeO extends World`. A single visitor interface can be used both for `Mutable` and `ReadOnly` nodes.

2.5 Visitor pattern

The *visitor design pattern* [17] is a way of separating an algorithm from a node hierarchy upon which it operates. Instead of distributing the node processing code among all the node implementations, the algorithm is written in a single *visitor* class that has a visit method for every node in the hierarchy. This is desirable when the algorithm changes frequently or when new algorithms are frequently created. The standard implementation (that does not use reflection) defines a tiny `accept` method that is overridden in all the nodes, that calls the appropriate visit method for that node.

Nägeli [25] discusses ownership in design patterns, and shows that previous *ownership* work was not flexible enough to express the visitor pattern. A visitor is always mutable because it may accumulate information during the traversal of the nodes hierarchy. However, some visitors only need readonly access to the nodes, and some need to modify the nodes. In the former case, the owner of the visitor and nodes may be different, and in the latter case, it must be the same owner. The challenge is to use the same `visit` and `accept` methods, and to avoid duplicating the traversal code.

OIGJ can express the visitor pattern by relying on owner-polymorphic methods: the owner of an object `o`, can pass it to an *owner-polymorphic method*, which cannot capture `o`.

Fig. 6 shows the visitor pattern in OIGJ. As mentioned before, the *owner* of the visitor and nodes may be different, and some visitors may or may not *modify* the nodes. Therefore, the visitor is parameterized on line 1 by the owner (`NodeO`) and immutability (`NodeI`) of the nodes. The `visit` method on line 2 is mutable because it changes the visitor that accumulates information during the traversal. Different visitor implementations may have different immutability for the nodes, e.g., readonly on line 14 or mutable on line 22.

Finally, note how the type arguments `This,ReadOnly` of the node on line 10 match the last two arguments of the visitor on line 12, and on line 18 the type arguments `This,Mutable` match those on line 20. This shows that the same `accept` method (without duplicating the nodes' hierarchy traversal code) can be used both for readonly and mutable hierarchies.

3. Formalization and Type Soundness

Proving soundness is essential in the face of complexities such as wildcards and raw/cooked objects. This section gives the typing rules and operational semantics of a simplified version of OIGJ and sketches the proofs of our immutability and ownership guarantees. For lack of space, the full proofs are included in our technical report [38].

Our type system, called Featherweight OIGJ (FOIGJ), is based on Featherweight Java (FJ) [20]. FOIGJ models the essence of OIGJ: the fact that every object has an ownership and immutability, and the cooking phase when creating immutable objects. FOIGJ adds imperative constructs to FJ, such as `null` values, field assignment, locations/objects, and a heap. FOIGJ also adds a constructor body (to model the cooking phase), owner and immutability parameters to classes, guards as in cJ [19], wildcard owners, and the runtime notion of raw/cooked objects.

FOIGJ poses two main challenges: (i) modeling *wildcards* in the typing rules, and (ii) the representation for *raw objects*. We use the following example (similar to Fig. 2) to demonstrate these two challenges:

```
class Foo<O,I> {
 Date<O,I> sameD;
 Date<This,I> ownedD;
 Date<This,Immut> immutD;
 <I extends Raw>? void Foo(){
   this.ownedD = new Date<This,I>();
   this.immutD = new Date<This,Immut>();
   ... } }
```

Wildcards pose a difficulty due to a process in Java called *wildcard capture* in which a wildcard is replaced with a fresh type variable. For example, the two underlined wildcards below might represent two distinct owners:

```
Foo<?,I> f = ...;
Date<?,I> d = ...;
f.sameD = d; // Illegal assignment! Different owners!
```

A Java compiler rejects the assignment due to incompatible types, because the wildcards were captured by dif-

ferent type variables. Formalizing the full power of wildcards (with upper and lower bounds) was only recently achieved [9]. FOIGJ does *not* model wildcard capture. Instead, it is enough to augment the field assignment rule with the following check: assigning to o is illegal if $O(o) = \texttt{?}$ (similarly for method invocation). This extra check is needed only in FOIGJ, and not in OIGJ, because OIGJ is built on top of Java, which supports wildcard capture.

The second challenge is modeling raw objects in the *non-erased* operational semantics. Recall that generics are erased in Java and are not present at run time. FOIGJ's *erased* operational semantics is identical to that of normal Java: ownership and immutability information is not kept. In contrast, the *non-erased* version stores with each object its owner and immutability, and it checks at run time the ownership and immutability guarantees (i.e., that field assignment respects owner-as-dominator and is done only on mutable or raw objects). The non-erased version is used only in the formalism. Storing the owner and immutability of every object at run time would be a huge overhead, and is not required for correctness if the program satisfies OIGJ's type rules.

The non-erased semantics of Featherweight Generic Java [20] (FGJ) performs variable substitution for method calls, however FGJ's way of doing substitution does not work in FOIGJ. For example, consider the following reduction as done in imperative FGJ:

```
new Foo<World,Immut>() →
                l.ownedD = new Date<l,Immut>();
                l.immutD = new Date<l,Immut>();...
```

The variable `I` in the constructor was substituted with `Immut`, and the variables `this` and `This` were substituted with a new location `l` that was created on the heap, i.e., the heap H now contains a new object in location `l` whose fields are all `null`: $H = \{\texttt{l} \mapsto \texttt{Foo<World,Immut>}(\overline{\texttt{null}})\}$. (Locations are pointers to objects; we treat locations and objects identically because they have a one-to-one mapping, e.g., the owner of a location is defined to be the owner of its object.) Note how owner parameters ($O(o)$) at compile time are replaced with owners ($\theta(o)$) at run time, e.g., `This` was replaced with location `l`.

There are two reasons why substituting `I` with `Immut` does not work in FOIGJ: (i) the reduction does not type-check because we mutate an immutable object (`l.ownedD = ...`), and (ii) we lost information about the two new `Date` objects, namely that `ownedD` can still be mutated after its constructor finishes (because it is `this`-owned) whereas `immutD` cannot.

FOIGJ solves these two issues by introducing an auxiliary type $\texttt{Immut}_\texttt{l}$. An object o of immutability $I(o) = \texttt{Immut}_\texttt{l}$ becomes cooked when the constructor of `l` finishes, therefore we call `l` its *cooker*, denoted by $\kappa(o) = \texttt{l}$. Phrased differently, an object is cooked when its cooker is *constructed* (i.e., the cooker's constructor finishes). Note that the cooker `l` can be o itself, its owner, or even some other incomparable object.

The connection between the cooker and the owner will be shown in the subtyping and typing rules below. Intuitively, for a reference of type $\texttt{C<o,Immut}_\texttt{l}\texttt{>}$, if the cooker `l` is not inside the owner `o`, then that reference must point to an object whose cooker is `l`. Otherwise (if `l` is inside `o`), then that reference might point to any cooked immutable object (even one with a cooker that is not `l`).

In our example, the location `l` that was created with `new Foo<World,Immut>()` becomes cooked when it is constructed, i.e., $\kappa(\texttt{l}) = \underline{\texttt{l}}$, and $H = \{\texttt{l} \mapsto \texttt{Foo<World,Immut}_\underline{\texttt{l}}\texttt{>}(\overline{\texttt{null}})\}$. Now FGJ's way of doing the substitution works for FOIGJ, because `I` is replaced with $\texttt{Immut}_\texttt{l}$, i.e.,

```
l.ownedD = new Date<l,Immut_l>();
l.immutD = new Date<l,Immut>();
```

Note how the cooker of `ownedD` is `l`, whereas the cooker of `immutD` is `immutD` itself. Therefore, `ownedD` has a longer cooking phase than `immutD`.

FOIGJ also maintains the set of currently executing constructors K, where $K \subseteq \text{dom}(H)$. We maintain the invariant that a location `l` is *raw* iff $\kappa(\texttt{l}) \in K$, and require that *only mutable or raw objects can be mutated*. Specifically, $\texttt{Immut}_\texttt{l}$ is a subtype of `Raw` when $\texttt{l} \in K$, and it is a subtype of `Immut` when $\texttt{l} \notin K$.

Type $\texttt{Immut}_\texttt{l}$ also helps understand the **Object creation rule** better. Recall that an `Immut` object can be created from a `Raw` constructor, even though `Immut` is not a subtype of `Raw`, which seems to contradict **cJ's method guard rule**. However, type $\texttt{Immut}_\texttt{l}$ is in fact a subtype of `Raw` when the object is created, because in our typing rules we have that $\texttt{Immut}_\texttt{l} \leq \texttt{Raw}$ iff $\texttt{l} \in K$, and when an object is created, its cooker must be in K. Phrased differently, the type of an immutable object never changes (always $\texttt{Immut}_\texttt{l}$), but during the program execution the set K changes, and therefore the subtyping relation changes: initially $\texttt{Immut}_\texttt{l}$ is a subtype of `Raw`, but later the object becomes cooked, and then $\texttt{Immut}_\texttt{l}$ is no longer a subtype of `Raw`, but instead it becomes a subtype of `Immut`.

Faced with such major challenges, we removed from FOIGJ anything that was not needed to prove our run-time guarantees. Specifically, FOIGJ does *not* model: generics (except for the owner and immutability parameters), owner polymorphic methods, casting, inner classes, fresh owners, or multiple immutability/owner parameters. On the one hand, the interaction between generics and *immutability* (which enables covariant subtyping) was previously proven sound in Featherweight IGJ (FIGJ) [39]. On the other hand, the interaction between generics and *ownership* (as found in the ownership nesting rule) was previously proven sound in Featherweight OGJ (FOGJ) [32]. Because covariant subtyping (as found in IGJ) and ownership nesting (as found in OGJ) was not changed in OIGJ, we decided not to model generics in FOIGJ. We note that FOIGJ does model `Raw`, which was not modeled previously in FIGJ.

FT ::= C<FO, IP>	Field (and method return) Type.
T ::= C<MO, IP>	Type.
N ::= C<NO, NI>	Non-variable type (for objects).
NO ::= World | l	Non-variable Owner parameter (for objects).
FO ::= NO | This | O	Field Owner parameter.
MO ::= FO | ?	Method Owner parameter (including generic wildcard).
NI ::= Mutable | $\texttt{Immut}_{\texttt{l}}$	Non-variable Immutability parameter (for objects).
VI ::= NI | Immut | I	Variable Immutability for new.
IP ::= ReadOnly | VI	Immutability Parameter.
IG ::= ReadOnly | Immut | Mutable | Raw	Immutability method Guard.
M ::= <I extends IG>? FT m($\overline{\texttt{T}}$ $\overline{\texttt{x}}$) { return e; }	Method declaration.
L ::= class C<O, I> extends C′<O, I>{ $\overline{\texttt{FT}}$ $\overline{\texttt{f}}$; $\overline{\texttt{M}}$ }	cLass declaration.
v ::= null | l	Values: either null or a location l.
e ::= v | x | e.f | e.f = e | e.m($\overline{\texttt{e}}$) | new C<FO, VI>($\overline{\texttt{e}}$) | e ;return l	Expressions.

Figure 7. FOIGJ Syntax. The terminals are null, owner parameters (O, This, World), and immutability parameters (I, ReadOnly, Mutable, Raw, Immut). Given a location l, $\texttt{Immut}_{\texttt{l}}$ represents an immutable object with cooker l. The program source code cannot contain the grayed elements (locations are only created during execution/reduction in R-New of Fig. 10).

Consider the typing rules in Fig. 4. Classes in FOIGJ have a single Raw constructor, therefore **Object creation rule** is always satisfied and can be ignored. Furthermore, because FOIGJ does not model generics, static, or inner classes, then the following rules can also be ignored: **Ownership nesting**, **Inner classes**, **Inheritance**, **Invariant**, **Erased signature**, and **Fresh owners**. Covariant subtyping and erased signatures were described in FIGJ, and ownership nesting and (limited) owner-polymorphic methods in FOGJ. We stress that FOIGJ *does* model wildcard for the owner parameter (?), which is used in owner-polymorphic methods such as sort or equals. In our view, extending the formalism with fresh owners or inner classes increases the complexity of the calculus without providing new insights.

The following rules are enforced by the syntax of FOIGJ (Fig. 7): **Generic Wildcards** and **Raw parameter**. The remaining rules are: **Field assignment**, **Field access**, **Method invocation**, **cJ's [19] method guard**, and a **Subtype relation** (without generics). These rules are formalized in FOIGJ in the subtyping rules of Fig. 8 ($K, \Gamma \vdash \texttt{T} \leq \texttt{T}'$) and the typing rules of Fig. 9 ($K, \Gamma \vdash \texttt{e} : \texttt{T}$). Finally, the reduction rules are described in Fig. 10 ($K \vdash H, \texttt{e} \rightarrow H', \texttt{e}'$).

Sec. 3.1 describes the syntax of FOIGJ, Sec. 3.2 the subtyping rules, Sec. 3.3 the typing rules, Sec. 3.4 the reduction rules, and Sec. 3.5 proves preservation, progress, and our run-time immutability and ownership guarantees.

3.1 Syntax of FOIGJ

FOIGJ adds imperative extensions to FJ such as assignment to fields, object locations, null, and a heap [31]. A constructor initializes all the fields to null, and then calls a build method that constructs the object. Having null values is important because this-owned fields must be initialized with null since they cannot be assigned from the outside, i.e., they must be created within this. For example, a list constructor cannot receive its entries as constructor arguments; instead it must create the entries within the build method.

Fig. 7 presents the syntax of FOIGJ. Expressions in FOIGJ include the four expressions in FJ (method parameter, field access, method invocation, and new instance creation; *without* casting), as well as the imperative extensions (field update, e;return l, and values). Expression e;return l is created when reducing a constructor call, e.g., $K \vdash$ new N(...) $\rightarrow$ l.build(...);return l, then we proceed to reduce l.build(...) and finally return l. Note that O and I are terminals, i.e., the owner and immutability parameters are always named O and I.

An evaluation of an expression (e) is either infinite, or is stuck on null-pointer exception, or terminates with a value (v), which is either null or a location l.

Note how the syntax limits the usage of wildcards and Raw: wildcards (?) can be used only as the owner of method arguments (FOIGJ does not have local variables), and Raw only as a method guard (IG).

We represent sequences using an over-line notation, similarly to FJ, i.e., comma denotes concatenation of sequences, and $\overline{\texttt{FT}}$ $\overline{\texttt{f}}$; represents the sequence $\texttt{FT}_1$ $\texttt{f}_1$; ... $\texttt{FT}_n$ $\texttt{f}_n$;

A class in FOIGJ has a single constructor that can create both mutable and immutable objects, i.e., it is a Raw constructor. The constructor is not shown in the syntax because it can be inferred from the class declaration: it always assigns null to the fields of the newly created object, and then invokes this special method (ignoring the return value):

<I extends Raw>? T′ build($\overline{\texttt{T}}$ $\overline{\texttt{e}}$) { return e; }

We require that each class has such a method, and that its parameters are not this-owned nor have wildcards. The reduction rules call that method after the fields are set to null.

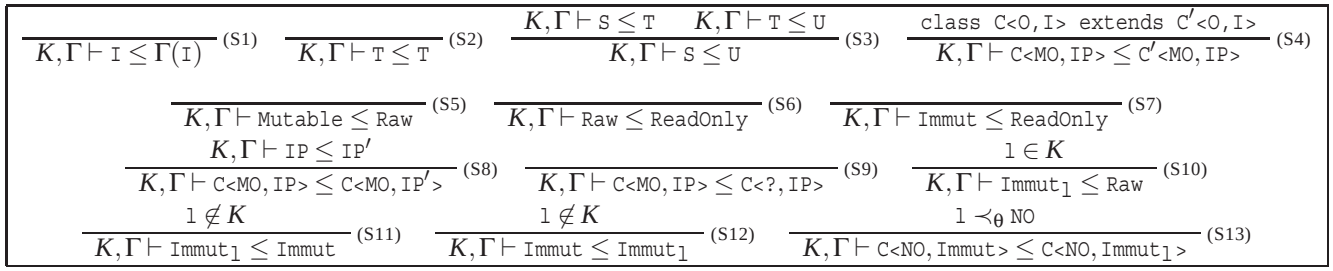

Figure 8. FOIGJ Subtyping Rules ($K, \Gamma \vdash \texttt{T} \leq \texttt{T}'$). Rule S13 shows the connection between cooker `l` and owner `NO`.

3.2 Subtyping in FOIGJ

An environment Γ is a finite mapping from variables `x` and locations `l` to types `T`, e.g., $\texttt{x} : \texttt{T} \in \Gamma$. The location types define the ownership tree $\preceq_\theta$ (or without reflexivity $\prec_\theta$). The set of currently executing constructors is denoted K. In addition, Γ maps the immutability parameter `I` to its bound according to the current method's guard (`IG` in Fig. 7). For example, $\texttt{I} : \texttt{Raw} \in \Gamma$ when typing the expression `e` in method `build` above.

Fig. 8 shows FOIGJ subtyping rules. Rules S1–S4 are the same as FGJ rules: S1 means that a generic variable is a subtype of its bound, S2 is reflexivity, S3 is transitivity, and S4 is that subclassing defines subtyping. Rules S5–S7 show subtyping among non-variable immutability parameters as shown in Fig. 1b. Rule S8 defines covariant subtyping for the immutability parameter. Rule S9 formalizes subtyping with a wildcard owner.

The last four rules S10–S13 are concerned with *cookers* such as $\texttt{Immut}_\texttt{l}$. Recall that an object is cooked when its cooker `l` is constructed, i.e., the constructor of `l` is no longer executing: $\texttt{l} \notin K$. Rule S10 views the type as `Raw`, while rules S11–S12 shows the equivalence to `Immut`. Note that subtyping is no longer antisymmetric, i.e., there are non-equal types $\texttt{T}_1$ and $\texttt{T}_2$ for which $\texttt{T}_1 \leq \texttt{T}_2 \leq \texttt{T}_1$. For example, $\texttt{T}_1 = \texttt{C<O,Immut}_\texttt{l}\texttt{>}$ and $\texttt{T}_2 = \texttt{C<O,Immut>}$, when $\texttt{l} \notin K$. In fact, this is not surprising because these types both represent immutable object, and after the cooker is cooked, the identity of the cooker is irrelevant.

Cooker vs. owner Rule S13 assumes that the cooker is inside the owner ($\texttt{l} \prec_\theta \texttt{NO}$), which means the object might came from the outside. This rule addresses the difference between the cooker of (i) a location `l` or (ii) that of an expression such as field access `l.f`: (i) location `l` will be cooked *exactly* when the constructor of $\kappa(\texttt{l})$ is finished, however, (ii) the cooker of `l.f` is an *over-approximation*, i.e., the object stored in that field might have been cooked earlier. Rule S13 allows an over-approximation only when the cooker is inside the owner.

Consider this example:

```
class Foo<O,I> { Date<O,I> same;
 <I extends Raw>? Foo(Date<O,I> d) { same=d; } }
```

Field `same` is assigned from the outside, but it might still be `this`-owned. We will show the reduction of two expressions: one where `same` is assigned a cooked (outside) date, and one with a raw date. The expressions are inside the constructor of some object `b` whose cooker is `b` itself.

The reduction of the first expression:

```
new Foo<This,Immut>(new Date<This,Immut>())
```

results in the heap: $H = \{\texttt{d1} \mapsto \texttt{Date<b,Immut}_\texttt{d1}\texttt{>()}, \texttt{f1} \mapsto \texttt{Foo<b,Immut}_\texttt{f1}\texttt{>(d1)}\}$. Note that the type of `d1` must be a subtype of `f1.same` in a well-typed heap (formally defined later). The type of `d1` is a subtype of `Date<b,Immut>` (because $\texttt{d1} \notin K$ in rule S12), which is a subtype of $\texttt{Date<b,Immut}_\texttt{f1}\texttt{>}$ (because $\texttt{f1} \prec_\theta \texttt{b}$ in rule S13). Phrased differently, the cooker of `f1.same` is `f1`, but it may point to an object that was cooked before, and indeed it points to an object whose cooker is `d1` (so it is an over-approximation).

The reduction of the second expression:

```
new Foo<This,I>(new Date<This,I>())
```

results in the heap (because $\texttt{I} = \texttt{Immut}_\texttt{b}$): $H = \{\texttt{d2} \mapsto \texttt{Date<b,Immut}_\texttt{b}\texttt{>()}, \texttt{f2} \mapsto \texttt{Foo<b,Immut}_\texttt{b}\texttt{>(d2)}\}$. Note that in this case, both `d2` and `f2` have the same cooker `b`. The type of `f2.same` is $\texttt{Date<b,Immut}_\texttt{b}\texttt{>}$, and because $\texttt{b} \not\prec_\theta \texttt{b}$ (see rule S13), then we know that this is not an over-approximation, i.e., that field points to an object whose cooker must be `b`.

To summarize, consider a type $\texttt{Foo<o,Immut}_\texttt{c}\texttt{>}$. If the cooker `c` is inside the owner `o` ($\texttt{c} \prec_\theta \texttt{o}$), or the cooker is cooked ($\texttt{c} \notin K$), then the type is an over-approximation, i.e., it can point to any `Immut` object (that is, to any object with cooker $\texttt{c}' \notin K$). Otherwise, it points to an object whose cooker is exactly `c`. Formally,

LEMMA 3.1. *If* $K, \Gamma \vdash \texttt{C<MO,IP>} \leq \texttt{C}'\texttt{<NO,Immut}_\texttt{l}\texttt{>}$, $\texttt{l} \not\prec_\theta \texttt{NO}$, *and* $\texttt{l} \in K$, *then* $\texttt{IP} = \texttt{Immut}_\texttt{l}$.

We also prove in the technical report that:

LEMMA 3.2. *If* $K, \Gamma \vdash \texttt{C<MO,IP>} \leq \texttt{C}'\texttt{<MO}', \texttt{IP}'\texttt{>}$, *then (i)* $\texttt{MO}' \neq \texttt{?} \Rightarrow \texttt{MO} = \texttt{MO}'$, *(ii)* $(\texttt{IP}' \neq \texttt{Immut}_\texttt{l}$ *or* $\texttt{l} \not\prec_\theta \texttt{MO}') \Rightarrow K, \Gamma \vdash \texttt{IP} \leq \texttt{IP}'$, *and (iii)* `C` *is a subclass of* $\texttt{C}'$, *(iv)* $K, \Gamma \vdash \texttt{D<l,IP>} \leq \texttt{D<l,IP}'\texttt{>}$ *for any class* `D` *and location* `l` *where* $\texttt{MO}' \preceq_\theta \texttt{l}$.

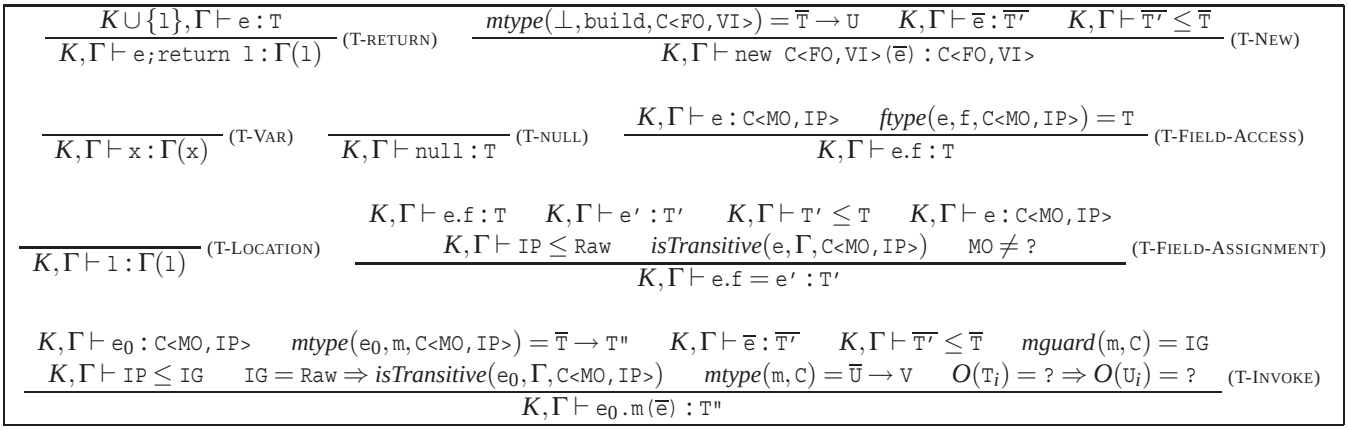

Figure 9. FOIGJ Expression Typing Rules ($K, \Gamma \vdash \texttt{e} : \texttt{T}$).

3.3 Typing rules of FOIGJ

Auxiliary functions We use the following auxiliary functions: *fields*(C) returns all the field names (including inherited fields) of class C, *ftype*(f,C) returns the type of field f in class C, *mtype*(m,C) returns the type of method m in class C, and *mbody*(m,C) returns its body. Their definitions are based on their counterparts in FJ, and thus omitted from this paper. In addition, function *mguard*(m,C) returns the method's guard (IG in Fig. 7).

We overload the auxiliary functions above to work also for types (T = C<MO,IP>) and not just classes (C) by substituting [MO/O,IP/I]. However, we also need to carefully substitute This when the receiver is this or locations. Function *ftype*(e,f,T) returns the type of the field access e.f where T is the type of e, or error if the access is illegal. For example,

$$\begin{aligned}
\textit{ftype}(\texttt{ownedD},\texttt{Foo}) &= \texttt{Date<This,I>}\\
\textit{ftype}(\texttt{this},\texttt{ownedD},\texttt{Foo<O,Immut>}) &= \texttt{Date<This,Immut>}\\
\textit{ftype}(\texttt{this.f},\texttt{ownedD},\texttt{Foo<O,Immut>}) &= \texttt{error}\\
\textit{ftype}(\texttt{l},\texttt{ownedD},\texttt{Foo<o,Immut}_{\texttt{C}}\texttt{>}) &= \texttt{Date<l,Immut}_{\texttt{C}}\texttt{>}
\end{aligned}$$

Formally, *ftype*(e,f,C<MO,IP>)=[MO/O,IP/I,z/This]*ftype*(f,C), where (i) z = l if e = l, (ii) z = This if e = this, (iii) otherwise z = error (and if a type contains error then it means the call to *ftype* failed). When we know the field is not this-owned, then the expression e is not used, and we write *ftype*(⊥,f,C). However, if the field is this-owned, then e must be this or a location l.

Recall that **Field access rule** in Fig. 4 required that this-owned fields can be accessed only via this. At run time, this is substituted with a location l. Therefore, there is a *duality* between this and a location l in the definition of *ftype*. For example, the field access this.ownedD of type Date<This,I> is legal because we accessed a this-owned field via this. At run time, this is substituted with some location l, and the access l.ownedD of type Date<l,...> is now still legal because we accessed a this-owned field via a location l. Note that if the access is not done via a location, e.g., bar.l.ownedD, then we cannot type-check the resulting expression (because we do not know what should be the substitute for This).

There is a similar duality in **Field assignment rule** part (ii), that checks that Raw is transitive for this or this-owned objects. The dual of this is a location l, and the dual of a this-owned object (C<This,I>) is an object whose cooker is not inside its owner (C<o,Immut$_l$> where $\texttt{l} \not\prec_\theta \texttt{o}$). The second duality holds because, for type C<This,I>, the cooker (I) is never inside the owner (This). At run time, the owner will be the location of this, and the cooker is either this or some other object that was created before this, i.e., the cooker is never inside the owner (but they might be equal).

Function *isTransitive* checks whether Raw is transitive. The underlined part shows the *dual* version of **Field assignment rule** part (ii): (i) this vs. l′, and (ii) MO = This vs. $\texttt{l} \not\prec_\theta$ MO.

$$\begin{aligned}
&\textit{isTransitive}(\texttt{e},\Gamma,\texttt{C<MO,IP>}) =\\
&\quad\big(\texttt{IP} = \texttt{I} \text{ and } \Gamma(\texttt{I}) = \texttt{Raw} \Rightarrow (\texttt{e=this or } \texttt{MO} = \texttt{This})\big) \text{ or}\\
&\qquad\underline{\big(\texttt{IP} = \texttt{Immut}_{\texttt{l}} \Rightarrow (\texttt{e=l'} \text{ or } \texttt{l} \not\prec_\theta \texttt{MO})\big)}
\end{aligned}$$

Typing class declarations FOIGJ program consists of class declarations followed by the program's expression. Next, we describe in words the rules for typing the class declarations, and the rules for typing an expression are given formally in Fig. 9. When typing an expression, we assumed the class declarations are well-formed.

To check that class declarations are well-formed, FOIGJ first performs all the checks done in FJ, e.g., that there are no cycles in the inheritance relation, that field and method names are unique in a class, that this is not a legal method parameter name, that an overriding method maintains the same signature, etc. FOIGJ performs additional checks related to method guards when typing method declarations, i.e., we modify rule T-Method in FJ as follows: (i) An overriding method can only make the guard weaker, i.e., if a method with guard IG overrides one with guard

IG′ then IG′ ≤ IG. (ii) In class C, when typing a method: <I extends IG>? FT m(T̄ x̄) { return e; }
we use an environment Γ in which the bound of I is IG, i.e., $\Gamma = \{\texttt{I} : \texttt{IG}, \overline{\texttt{x}} : \overline{\texttt{T}}, \texttt{this} : \texttt{C<O,I>}\}$, and we must prove that $\emptyset, \Gamma \vdash \texttt{e} : \texttt{S}$ and $\emptyset, \Gamma \vdash \texttt{S} \leq \texttt{FT}$. Finally, we require that if IG = ReadOnly then $I(\texttt{T}_i) \neq \texttt{I}$. This last requirement is not really a limitation, because a programmer can replace I with ReadOnly for parameters in readonly methods, and previously legal programs would remain legal. (This requirement is needed to prove preservation for the congruence rule of method receiver, see our technical report for details.)

Typing expressions Fig. 9 shows the typing rules for expressions in FOIGJ. Most of these rules are a direct translation from Fig. 4. The main challenge was handling *wildcards* correctly without resulting to wildcard-capture. Rule T-Field-assignment requires that MO ≠ ?, i.e., one cannot assign to an object with unknown owner. Typing method parameters is similar to typing field-assignment, however, method parameters can have a wildcard owner whereas fields cannot (see the difference between T and FT in Fig. 7). Therefore, rule T-Invoke requires that $O(\overline{\texttt{T}}) = \texttt{?} \Rightarrow O(\overline{\texttt{U}}) = \texttt{?}$, i.e., if $\texttt{T}_i = \texttt{C<?,IP>}$ then $\texttt{U}_i = \texttt{C<?,IP'>}$. Phrased differently, if the owner of $\texttt{e}_0$ is unknown, then the owner of the method parameters cannot be O.

Rule T-New performs less checks compared to a method call (e.g., no need to check the guard, *isTransitive*, nor wildcards) because build has several restrictions: its guard is Raw and it does not contain wildcards nor this-owned parameters. Because This does not appear in the signature of build, we know that *mtype* will not use ⊥ in the call: $\mathit{mtype}(\bot, \texttt{build}, \texttt{C<FO,VI>})$.

An expression/type is called *closed* if it does not contain any free variables (such as wildcards, this, I, O, or This), but it may contain World, ReadOnly, Mutable, Immut, $\texttt{Immut}_\texttt{l}$ or locations. Note that the type of a closed expression is also closed.

LEMMA 3.3. *If* $K, \Gamma \vdash \texttt{e}' : \texttt{T}'$ *and* $\texttt{e}'$ *is closed and* $\texttt{e}' \neq \texttt{null}$, *then* $\texttt{T}'$ *is closed.*

The delicate part of the proof is showing that T′ does not contain This. Note that *ftype* returns a type with This only if e = this (which cannot happen since e is closed).

3.4 Reduction rules of FOIGJ

The initial expression to be reduced is *closed*, and we guarantee that a closed expression is always reduced to another closed expression:

LEMMA 3.4. *If* e *is closed and* $K \vdash H, \texttt{e} \rightarrow H', \texttt{e}'$, *then* $\texttt{e}'$ *is closed.*

The heap (or store) H maps each location l to an object $\texttt{C<NO,NI>}(\overline{\texttt{v}})$, where $\theta_H(\texttt{l}) = \texttt{NO}$ is its owner, and $I_H(\texttt{l}) = \texttt{NI}$ is its immutability, and $\overline{\texttt{v}}$ are the values of its fields. If $\texttt{NI} = \texttt{Immut}_{\texttt{l}'}$ then we say that its cooker is $\kappa_H(\texttt{l}) = \texttt{l}'$. (We added the subscript H to the functions that return the owner, immutability and cooker, in order to explicitly show the dependence on the heap.) We define a *heap-typing* $\Gamma_H : \texttt{l} \mapsto \texttt{T}$ that gives a type to each location in the obvious way (simply removing the list of fields $(\overline{\texttt{v}})$).

The set of currently executing constructors is K. A heap H is *well-typed for K* if it satisfies two conditions: (i) Each field location is a subtype (using K, Γ_H) of the declared field type, i.e., for every location l, where $H[\texttt{l}] = \texttt{C<NO,NI>}(\overline{\texttt{v}})$ and $\mathit{fields}(\texttt{C}) = \overline{\texttt{f}}$, and for every field $\texttt{f}_i$, we have that either $\texttt{v}_i = \texttt{null}$ or $K, \Gamma_H \vdash \Gamma_H(\texttt{v}_i) \leq \mathit{ftype}(\texttt{l}, \texttt{f}_i, \texttt{C<NO,NI>})$. (ii) There is a linear order $\preceq^T$ over $\mathrm{dom}(H)$ such that for every location l, $\theta_H(\texttt{l}) = \texttt{World}$ or $\theta_H(\texttt{l}) \prec^T \texttt{l}$, and $I_H(\texttt{l}) = \texttt{Mutable}$ or $\kappa_H(\texttt{l}) \preceq^T \texttt{l}$. The linear order $\preceq^T$ can order the objects according to their *creation time*, because $\theta_H(\texttt{l})$ is always created before l, and $\kappa_H(\texttt{l})$ is either l or created before l.

In our technical report we prove that if H is well-typed for K then (a) owner-as-dominator holds (Lem. 3.5), and (b) H is well-typed for any subset of K (Lem. 3.6). Part (b) is not trivial, because the subtyping relation for a subset of K is different because raw objects become immutable. Intuitively, during execution objects become cooked (when their cooker is removed from K), and therefore Lem. 3.6 guarantees that the heap remains well-typed when K decreases.

LEMMA 3.5. *If heap H is well-typed for K, then for every location* $\texttt{l} \in \mathrm{dom}(H)$, $\texttt{l} \mapsto \texttt{C<NO,NI>}(\overline{\texttt{v}})$, *then either* $\texttt{v}_i = \texttt{null}$ *or* $\texttt{l} \preceq_\theta \theta_H(\texttt{v}_i)$.

LEMMA 3.6. *Given a heap H that is well-typed for K, then for any S ⊂ K, the heap H is well-typed for S.*

Fig. 10 presents the reduction rules in a small-step notation, excluding all congruence rules except R-C1.

Rule R-Return ignores the return value of build and returns l. Rule R-Field-Access is trivial. Rule R-Field-assignment enforces our immutability guarantee (only mutable or raw objects can be mutated) and our ownership guarantee (owner-as-dominator, i.e., l can point to v′ iff $\texttt{l} \preceq_\theta \theta_H(\texttt{v}')$). Rule R-Invoke finds the method body according to the receiver, and substitutes all the free variables in the method body.

Rule R-C1 is the congruence rule for e;return l. Note that this rule is the only place the set K is modified, i.e., when reducing e, the set of ongoing constructors is $K \cup \{\texttt{l}\}$. It is easy to prove that if $K \vdash H, \texttt{e} \rightarrow H', \texttt{e}'$ then $\Gamma_H \subseteq \Gamma_{H'}$. The other congruence rules are not shown because they are trivial, e.g., in order to reduce a method call $\texttt{e}_0.\texttt{m}(\overline{\texttt{e}})$, we first reduce $\texttt{e}_0$ to a location, then reduce the first argument to a value, etc.

Rule R-New creates a new location l, sets the fields to null, sets the cooker of l (VI′) and finally calls build. In order to build the newly created object l, then it must be raw, i.e., its cooker VI′ must be in K. (Note that l will be in K according to R-C1.) Therefore, if $\texttt{VI} = \texttt{Immut}_\texttt{c}$ and $\texttt{c} \notin K$, then we must set the cooker to l. This can happen if there is a method

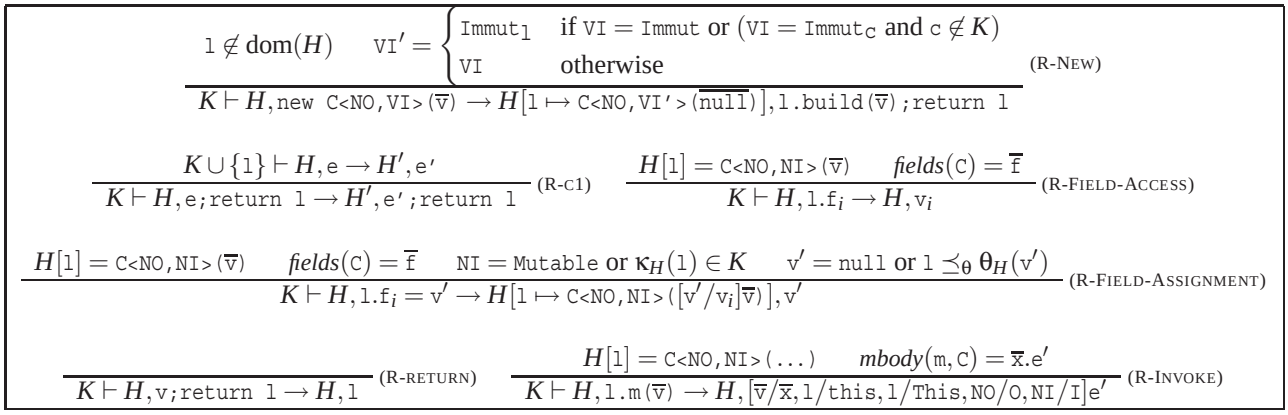

Figure 10. FOIGJ Reduction Rules ($K \vdash H, \mathtt{e} \to H', \mathtt{e}'$), excluding all congruence rules except R-C1.

that returns `new C<O,I>(...)` and the receiver is a *cooked* immutable object.

3.5 Guarantees of FOIGJ

We now turn to prove various properties of FOIGJ, including preservation theorem, ownership and immutability guarantees, and an erasure property. In the remainder of this section, we assume that reduction does not get stuck on *null-pointer exceptions*, i.e., the receiver/target of field access, assignment and method calls is never `null`. Under this assumption, then `e` can always be reduced to another expression `e'`.

Before stating the preservation theorem, we need to establish a connection between K and the reduced expression `e`, which may contain `return l`. Given an expression `e`, we define $K(\mathtt{e})$ to be the set of all ongoing constructors in `e`, i.e., all the locations in subexpressions `e';return l`. Formally, $K(\mathtt{e;return\ l}) = K(\mathtt{e}) \cup \{l\}$, and for any other expression we just recurse into all subexpressions, e.g., $K(\mathtt{e.f=e'}) = K(\mathtt{e}) \cup K(\mathtt{e'})$.

We will maintain the invariant that H is well-typed for $K \cup K(\mathtt{e})$. From Lem. 3.6, then H will also be well-typed for K. Initially, we start with a closed expression `e` without any locations (therefore $K(\mathtt{e}) = \emptyset$), an empty heap H, and an empty set of constructors K.

THEOREM 3.7. ***(Progress and Preservation)*** *For every closed expression* $\mathtt{e} \neq \mathtt{v}$*, K, and H, if* $K, \Gamma_H \vdash \mathtt{e} : \mathtt{T}$ *and H is well-typed for* $K \cup K(\mathtt{e})$*, then there exists* $H', \mathtt{e}', \mathtt{T}'$ *such that* $K \vdash H, \mathtt{e} \to H', \mathtt{e}'$*, H' is well-typed for* $K \cup K(\mathtt{e}')$*,* `T`*,* `T'`*, and* `e'` *are closed,* $K, \Gamma_{H'} \vdash \mathtt{e}' : \mathtt{T}'$*, and* $K, \Gamma_{H'} \vdash \mathtt{T}' \leq \mathtt{T}$.

Proved by showing there is always (exactly) one applicable reduction rule, which preserves subtyping. From Lem. 3.4, we know that `e'` is closed, and from Lem. 3.3, we know that `T` and `T'` are closed. Next we mention some highlights from the proof. In rule R-RETURN, we have that $K(\mathtt{e}') = K(\mathtt{e}) \setminus \{\mathtt{l}\}$, but even though we shrink K, we still have a well-typed heap from Lem. 3.6. In rule R-FIELD-ASSIGNMENT, Lem. 3.5 shows that the assumption $\mathtt{l} \preceq_\theta \theta_H(\mathtt{v}')$ holds, and the resulting heap is well-typed for $K \cup K(\mathtt{e}')$ because $K(\mathtt{e}') = \{\}$ and from T-FIELD-ASSIGNMENT. In rule R-NEW, we need to type the call $\mathtt{l.build}(\overline{\mathtt{v}})$, and for parameters with immutability `I`, we use the subtyping rule S13.

Our ownership and immutability guarantees follow directly from the reduction rules, because rule R-FIELD-ASSIGNMENT enforces them.

Thm. 3.8 shows that there is no need to maintain at run time K nor to store the owner and immutability parameter of each object. Formally, we define an erased heap structure $E(H)$ that maps location to objects without these parameters, i.e., $\mathtt{l} \mapsto \mathtt{C}(\overline{\mathtt{v}}) \in E(H)$. We define the erasure of an expression `e`, $E(\mathtt{e})$, by deleting all generic parameters, and define new reduction rules $\to_E$ in the obvious way.

THEOREM 3.8. ***(Erasure)*** *If* $K \vdash H, \mathtt{e} \to H', \mathtt{e}'$ *then* $K \vdash E(H), E(\mathtt{e}) \to_E E(H'), E(\mathtt{e}')$.

4. OIGJ Case Studies

This section describes our implementation of OIGJ: the language syntax (Sec. 4.1) and the type-checker implementation (Sec. 4.2). Sec. 4.3 presents our case study that involved annotating Sun's implementation of the `java.util` collections, and our conclusions about the design of the collection classes w.r.t. ownership and immutability.

The prototype OIGJ type-checker is implemented and distributed as part of the Checker Framework [30][2], which supports pluggable type systems using type annotations.

4.1 Syntax: from generics to annotations

Whereas this paper uses generics to express ownership and immutability (e.g., `Date<O,I>`), our OIGJ implementation uses Java 7's type annotations [15] (e.g., `@O @I Date`). Java 7's receiver annotations play the role of cJ's guards.

Using annotations has the advantage of compatibility with existing compilers and other tools. Another advantage is the ability to use a default value, such as `@Mutable`. Fur-

[2] `http://types.cs.washington.edu/checker-framework/`

thermore, it is possible to customize these defaults per class. Defaults are not possible in generics, because a programmer must supply arguments for all generic parameters.

Using annotations has the disadvantage that some notions are no longer explicit in the syntax, such as transitivity, wildcards, and generic methods. For example, compare the annotation and generic syntax:

```
class Foo      { @O @I Bar bar; }
class Foo<O,I> { Bar<O,I>  bar; }
```

Note that the type of `new Foo<World,Immut>().bar` is explicit in the generic syntax, whereas the annotation syntax (`new @World @Immut Foo().bar`) requires additional rules that mimic generics. Use of annotations also complicates the implementation (see below). For practical use, the compatibility benefits of using annotations outweigh their disadvantages.

OIGJ's annotations are the Cartesian product of owner parameters and immutability parameters. Our implementation does not yet support wildcards (though in practice the `@I` and `@O` annotations subsume most need for wildcards), nor classes with multiple owner or immutability parameters, such as `Iterator<O,ItrI,CollectionI,E>` in Fig. 3. (In our case study, we implemented iterators using a single immutability parameter by declaring `next()` as mutable.)

A class declaration can be annotated as `@Immut` to indicate class immutability, i.e., all instances are immutable and no mutable methods exist.

4.2 OIGJ implementation

Because a pluggable type checker augments, rather than replaces, the type system of the underlying language, the Checker Framework permits only language extensions that are stricter than ordinary Java. A pluggable type system cannot relax Java's rules, as the OIGJ subtyping rule does. For example,

```
@Immut    List<@Immut    Date>   a;
@ReadOnly List<@ReadOnly Date>   b=a; // OK
@Immut    List<@Immut    Object> c=a; // Illegal!
```

The assignment `c=a` is illegal in Java and therefore in the Checker Framework, though it is legal in OIGJ itself. Phrased differently, in our implementation, the covariance is limited to annotations.

The OIGJ type-checker incorporates, extends, and in some places overrides the IGJ checker, and adds OGJ features. It consists of about 700 source lines of code (of which 100 lines is Java boilerplate to define the annotations). Most of the code handles default and implicit types.

4.3 `java.util` collections case study

As a case study, we type-checked Sun's implementations of the `java.util` collections (77 classes, 33,246 lines of code). This required us to write 85 *ownership-related* annotations and 46 *immutability-related* annotations in 102 lines of code (the lines with `new` usually contain 2 annotations).

Sun's collections are not type-safe with respect to generics because Java does not support generic arrays. However, the OIGJ implementation uses type annotations, which can be placed on arrays as well, and therefore our annotated collections type-check without any errors with respect to ownership and immutability.

Class `LinkedList` in Fig. 3 is similar in essence to Sun's implementation. We annotated the constructors with `Raw`, thus allowing creation of immutable instances. Since all instances of `Entry` are `this`-owned, using `@This @I` as the default annotation for `Entry` meant that only three *ownership-related*[3] annotations were needed in `LinkedList`:

```
@Default({This.class, I.class})
static class Entry<E> {
  E element; @O Entry<E> next; @O Entry<E> prev;
  ... }
```

Similarly, in a `HashMap`, both the array and the entries are `this`-owned: `@This @I Entry[@This @I] table;`

The case study supports these conclusions: (i) the collections classes are properly encapsulated (they own their representation), (ii) it is possible to create immutable instances (all constructors are `Raw`), and (iii) methods `Map.get` and `clone` contain design mistakes (see below). We were not previously aware of these design mistakes. We believe that if the collections were designed with ownership and immutability in mind, such mistakes could be avoided.

***Immutability of method* `get`** Let's start with a quick riddle: is there a `Map` implementation in `java.util` that might throw an exception when running the following *single-threaded* code?

```
for (Object key : map.keySet()) { map.get(key); }
```

The answer is that for a map created with

```
new LinkedHashMap(100, 1, /*accessOrder=*/ true)
```

that contains more than one element, the above code throws `ConcurrentModificationException` after printing one element.

Most programmers assume that `Map.get` is readonly, but there is no such guarantee in Java's specification. The documentation of `LinkedHashMap` states: "*A special constructor is provided to create a linked hash map whose order of iteration is the order in which its entries were last accessed, from least-recently accessed to most-recently (*access-order*). Invoking the* `put` *or* `get` *method results in an access to the corresponding entry.*"

Because calling `get` modified the list, the above code threw `ConcurrentModificationException`. Phrased differently, method `LinkedHashMap.get` is mutable! Because an overriding method can only strengthen the specification of the overridden method, `HashMap.get` and `Map.get` must be mutable as well.

[3] The other annotations are immutability-related, e.g., receiver annotations.

	OIGJ	OGJ [32]	IGJ [39]	GUT [14]	UTT [24]	IOJ [18]	JOE_3 [29]
Owner-as-dominator	+	+				+	+
Owner-as-modifier				+	+		
Readonly references	+		+	+	+		+
Immutable objects	+		+			+	+
Uniqueness							+
Ownership transfer					+		
Factory method pattern	+	+	+	+	+		+
Visitor pattern	+		+				
Sun's `LinkedList`	+						
Case studies available	+		+				

Figure 11. Features supported by various language designs.

***Ownership and method* `clone`** Method `clone` violates owner-as-dominator because it leaks `this`-owned references by creating a shallow copy, i.e., only immediate fields are copied. Furthermore, Sun's implementation of `LinkedList` assigns to `result.header`, which is a `this`-owned field. This violates **Field assignment rule** of Sec. 2.3, which only permits assignment to `this.header`.

```
// The following code appears in LinkedList.clone().
// Calling super.clone() breaks owner-as-dominator because
// it leaked this.header to result.header.
LinkedList result = (LinkedList) super.clone();
result.header = new Entry(); // Illegal in OIGJ!
```

We sketch a solution that, instead of initializing the cloned `result` from `this`, uses the idea of *inversion of control*. The solution has two parts. (1) The programmer writes a method `constructFrom` that initializes `this` from a parameter. (This is similar to a copy-constructor in C++, and indeed this method should be given all the privileges of a constructor, such as assignment to `final` fields.) (2) The compiler automatically generates a `clone` method that first nullifies all the reference fields and then calls the user generated `constructFrom` method. This approach enforces the ownership and immutability guarantees.

5. Related Work

In this section we discuss related work on ownership and immutability. We first highlight the relationship between OIGJ and our previous work on ownership (OGJ) and immutability (IGJ). We also survey some of the most relevant related language designs and show how OIGJ compares to them.

5.1 Relationship with OGJ and IGJ

OIGJ can be thought of as the "cartesian product" of OGJ and IGJ: OIGJ uses two type parameters to express ownership and immutability. However, the delicate intricacies between ownership and immutability required changes to both OGJ and IGJ, making OIGJ more expressive than a naive combination.

Ownership Generic Java (OGJ) [32] demonstrated how ownership and generic types can be unified as a language feature. OGJ featured a single owner parameter for every class that was treated in the same way as normal generic type parameters, simplifying the language, the formalism, and the implementation.

OGJ completely prohibits wildcards as owner parameters, e.g., `Point<?>`, whereas OIGJ relaxes this rule and allows wildcards on stack variables, which enables writing the `equals` method (see **Generic Wildcards rule** in Sec. 2.3).

In OGJ, a method may have generic parameters that are owner parameters, e.g.,

```
class Foo<O extends World> {
  <O2 extends World> void bar(Object<O2> o) {...}
```

However, OGJ required that the parametric owners are *outside* the owner of the class, e.g., $\texttt{O} \preceq_\theta \texttt{O2}$. This rule is very restrictive, however it guarantees that the ownership structure is a tree. OIGJ removed this rule at the cost of complicating the ownership structure: it is a *directed acyclic graph* (DAG) instead of a tree.

Finally, OIGJ can express temporary ownership within a method by using a fresh owner parameter (see **Fresh owners** in Sec. 2.3).

Immutability Generic Java (IGJ) [39] showed how generic types can be used to provide support for readonly references and object immutability. OIGJ used ownership information to improve the expressiveness of IGJ. Specifically, certain restrictions in IGJ no longer apply in OIGJ for `this`-owned objects. For example, `Raw` is *not* transitive in IGJ, e.g., the assignment to `next` in Fig. 3 on lines 9 and 22 is illegal in IGJ, thus limiting creation of immutable objects. In contrast, `Raw` is transitive in OIGJ for `this`-owned fields (see **Field assignment rule** in Sec. 2.3), and therefore there was no need to refactor the collections' code.

IGJ includes an `@Assignable` annotation on fields that permits field assignment even in immutable objects. The `@Assignable` annotation indicates that a given field is not part of the object's abstract state. This is necessary to typecheck caches, lazily-initialized fields, and other programming idioms. OIGJ removed this annotation to simplify the formalism. This also guarantees representation immutability as well as immutability of the abstraction: the fields of a cooked immutable object never change. Our implementation supports the `@Assignable` annotation.

IGJ only permits a single immutability parameter, which simplifies the subtyping rule. In contrast, types in OIGJ can have multiple immutability parameters, for example, `Iterator<O,ItrI,CollectionI,E>`. Because IGJ uses a single immutability parameter, the immutability of an iterator and its underlying collection must be the same. Thus, in IGJ, method `next()` must be readonly (or you couldn't iterate over a readonly list), and therefore we had to use an `@Assignable` annotation on `ListItr.current` (line 37 in Fig. 3). In contrast, in OIGJ, we guard `next()` with a mu-

table `ItrI` (line 51), and guard `remove()` with a mutable `CollectionI` (line 52).

5.2 Relationship with other work

OIGJ uses method guards borrowed from *cJ* [19]. (The OIGJ implementation uses annotations syntax instead.)

In what follows, we have room to survey only closely related papers. Fig. 11 compares OIGJ to some of the previous work described below.

Mutability and encapsulation were first combined by Flexible Alias Protection (FLAP) [27]. FLAP inspired a number of proposals including ownership types [13] and confined types [35]. Capabilities for Sharing [5] describes the fundamentals underlying various encapsulation and mutability approaches by separating "mechanism" (the semantics of sharing and exclusion) from "policy" (the guarantees provided by the resulting system). Capabilities gives a lower-level semantics that can be enforced at compile or run time. A reference can possess any combination of these 7 access rights: read, write, identity (permitting address comparisons), exclusive read, exclusive write, exclusive identity, and ownership (giving the capability to assert rights). Immutability, for example, is represented by the lack of the write right and possession of the exclusive write right. Finally, *Fractional Permissions* [6] can give semantics to various annotations such as unique, readonly, method effects, and an ownership variant called *owner-as-effector* in which one cannot read or write owned state without declaring the appropriate effect for the owner.

Ownership types [2, 3, 11] impose a structure on the references between objects in a program's memory. OIGJ and other work [29, 32] enforce the owner-as-*dominator* disciplines. Generic Universe Types (**GUT**) [14, 23] enforce owner-as-*modifier* by using three type annotations: `rep`, `peer`, and `readonly`. `rep` denotes representation objects (similar to `This`), while `peer` denotes objects owned by the same owner (similar to `O`). **UTT** [24] is an extension of Universe Types that supports ownership transfer by utilizing a modular static analysis, which is useful for merging data-structures or complex object initialization.

MOJO [10] can express multiple ownership, i.e., objects can have more than one owner at run time. OIGJ supports only a single *owner* at run time. (OIGJ supports multiple *owner parameters*, but according to the **Ownership nesting rule**, all owner parameters are inside the first owner parameter.)

Jo$\exists$ [8] supports variant subtyping over the owner parameter by using existential types. OIGJ supports wildcards used as owners for stack variables, but those are less flexible than Jo$\exists$. For example, Jo$\exists$ can distinguish a list of students that may have different owners, from a list of student that share the same unknown owner.

Immutability and ownership. Similarly to OIGJ, Immutable Objects for a Java-like Language (**IOJ**) [18] associates with each type its mutability and owner. In contrast to OIGJ, IOJ does not have generics, nor readonly *references* (only readonly and immutable *objects*). Moreover, in IOJ, the constructor cannot leak a reference to `this`.

X10 [28] supports constrained types that can refer to properties and final local variables. X10 supports cyclic immutable structures by using `proto` annotations, which are similar to our immutability `I` and the notion of cookers. However, both X10 and IGJ cannot type-check Sun's `LinkedList` because an object becomes cooked when its constructor finishes. It is possible to refactor `LinkedList` to fit X10's typing-rules by using a recursive implementation, but then you risk a stack-overflow when creating large lists. Delayed types [16], which are similar to X10's `proto`, are used to verify non-null fields or other heap-monotonic properties.

JOE$_3$ [29] combines ownership (as dominators, not modifiers), uniqueness, and immutability. It also supports owner-polymorphic methods, but not existential owners.

Frozen Objects [22] show how ownership can help support immutability by allowing programmers to decide when the object should become immutable. This system takes a verification approach rather than a simple type checker such as OIGJ. Frozen Objects show how flexible the initialization stage can potentially be in the presence of ownership and immutability, while OIGJ shows how much flexibility can be achieved while staying at the type checking level.

Readonly references are found in C++ (using the `const` keyword), JAC [21], modes [33], Javari [34], etc. Previous work on readonly references lack ownership information. Boyland [4] observes that readonly does not address observational exposure, i.e., modifications on one side of an abstraction boundary that are observable on the other side. Immutable objects address such exposure because their state cannot change.

List iterators pose a challenge to ownership because they require a direct pointer to the list's privately owned entries, thus breaking the owner-as-dominator property. Both OIGJ and SafeJava [3] allow an inner instance to access the outer instance's privately owned objects. Clarke [11] suggested to use iterators only with stack variables, i.e., you cannot store an iterator in a field. It is also possible to redesign the code and implement iterators without violating ownership, e.g., by using internal iterators or magic-cookies [26].

6. Conclusion

OIGJ is a Java language extension that supports both ownership and immutability, while enhancing the expressiveness of each individual concept. By using Java's generic types, OIGJ simplifies previous type mechanisms, such as existential owners, scoped regions, and owner-polymorphic methods. OIGJ is easy to understand and implement, using only 14 (flow-insensitive) typing rules beyond those of Java. We have formalized a core calculus called FOIGJ and proved it sound. Our implementation is backward-compatible with Java, and it scales to realistic programs. OIGJ can type-check

Sun's `java.util` collections (without the `clone` method), using a small number of annotations. Finally, various design patterns, such as the factory and visitor patterns, can be expressed in OIGJ, making it ready for practical use. An implementation is publicly available at `http://types.cs.washington.edu/checker-framework/`.

Future work includes inferring ownership and immutability annotations, conducting a bigger case study including client and library code, and extending OIGJ with concepts such as *owner-as-modifier* [40], *uniqueness*, and *external-uniqueness* [12].

Acknowledgments

This work was supported by the New Zealand Royal Society Marsden Grant, ISAT Grant, and NSF grant CNS-0855252. Werner Dietl, James Noble, the ELVIS group, and the anonymous referees provided valuable feedback.

References

[1] Chandrasekhar Boyapati. *SafeJava: A Unified Type System for Safe Programming*. PhD thesis, MIT Dept. of EECS, Feb. 2004.

[2] Chandrasekhar Boyapati, Robert Lee, and Martin Rinard. Ownership types for safe programming: Preventing data races and deadlocks. In *OOPSLA*, pages 211–230, Oct. 2002.

[3] Chandrasekhar Boyapati, Barbara Liskov, and Liuba Shrira. Ownership types for object encapsulation. In *POPL*, pages 213–223, Jan. 2003.

[4] John Boyland. Why we should not add `readonly` to Java (yet). In *FTfJP*, July 2005.

[5] John Boyland, James Noble, and William Retert. Capabilities for sharing: A generalisation of uniqueness and read-only. In *ECOOP*, pages 2–27, June 2001.

[6] John Boyland, William Retert, and Yang Zhao. Comprehending annotations on object-oriented programs using fractional permissions. In *IWACO*, pages 1–11, July 2009.

[7] Gilad Bracha, Martin Odersky, David Stoutamire, and Philip Wadler. Making the future safe for the past: Adding genericity to the Java programming language. In *OOPSLA*, pages 183–200, Oct. 1998.

[8] Nicholas Cameron and Sophia Drossopoulou. Existential quantification for variant ownership. In *ESOP*, pages 128–142, Mar. 2009.

[9] Nicholas Cameron, Sophia Drossopoulou, and Erik Ernst. A model for Java with wildcards. In *ECOOP*, pages 2–26, July 2008.

[10] Nicholas R. Cameron, Sophia Drossopoulou, James Noble, and Matthew J. Smith. Multiple ownership. In *OOPSLA*, pages 441–460, Oct. 2007.

[11] Dave Clarke and Sophia Drossopoulou. Ownership, encapsulation and the disjointness of type and effect. In *OOPSLA*, pages 292–310, Oct. 2002.

[12] Dave Clarke and Tobias Wrigstad. External uniqueness is unique enough. In *ECOOP*, pages 176–200, July 2003.

[13] David G. Clarke, John M. Potter, and James Noble. Ownership types for flexible alias protection. In *OOPSLA*, pages 48–64, Oct. 1998.

[14] Werner Dietl, Sophia Drossopoulou, and Peter Müller. Generic Universe Types. In *ECOOP*, pages 28–53, Aug. 2007.

[15] Michael D. Ernst. Type Annotations specification (JSR 308). `http://types.cs.washington.edu/jsr308/`, Sep. 12, 2008.

[16] Manuel Fähndrich and Songtao Xia. Establishing object invariants with delayed types. In *OOPSLA*, pages 337–350, Oct. 2007.

[17] Erich Gamma, Richard Helm, Ralph E. Johnson, and John Vlissides. *Design Patterns*. Addison-Wesley, Reading, MA, 1995.

[18] Christian Haack, Erik Poll, Jan Schäfer, and Aleksy Schubert. Immutable objects for a Java-like language. In *ESOP*, pages 347–362, Mar. 2007.

[19] Shan Shan Huang, David Zook, and Yannis Smaragdakis. cJ: Enhancing Java with safe type conditions. In *AOSD*, pages 185–198, Mar. 2007.

[20] Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. Featherweight Java: a minimal core calculus for Java and GJ. *ACM TOPLAS*, 23(3):396–450, May 2001. ISSN 0164-0925.

[21] Günter Kniesel and Dirk Theisen. JAC — access right based encapsulation for Java. *Software: Practice and Experience*, 31(6):555–576, 2001.

[22] K. Rustan M. Leino, Peter Müller, and Angela Wallenburg. Flexible immutability with frozen objects. In *VSTTE*, pages 192–208, Oct. 2008.

[23] P. Müller and A. Poetzsch-Heffter. Universes: A type system for controlling representation exposure. In *Programming Languages and Fundamentals of Programming*, pages 131–140, 1999.

[24] Peter Müller and Arsenii Rudich. Ownership transfer in universe types. In *OOPSLA*, pages 461–478, Oct. 2007.

[25] Stefan Nägeli. Ownership in design patterns. Master's thesis, ETH Zürich, Zürich, Switzerland, Mar. 2006.

[26] James Noble. Iterators and encapsulation. In *TOOLS Pacific*, pages 431–442, 2000.

[27] James Noble, Jan Vitek, and John Potter. Flexible alias protection. In *ECOOP*, pages 158–185, July 1998.

[28] Nathaniel Nystrom, Vijay Saraswat, Jens Palsberg, and Christian Grothoff. Constrained types for object-oriented languages. In *OOPSLA*, pages 457–474, Oct. 2008.

[29] Johan Östlund, Tobias Wrigstad, and Dave Clarke. Ownership, uniqueness and immutability. In *Tools Europe*, pages 178–197, 2008.

[30] Matthew M. Papi, Mahmood Ali, Telmo Luis Correa Jr., Jeff H. Perkins, and Michael D. Ernst. Practical pluggable types for Java. In *ISSTA*, pages 201–212, July 2008.

[31] Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.

[32] Alex Potanin, James Noble, Dave Clarke, and Robert Biddle. Generic Ownership for Generic Java. In *OOPSLA*, pages 311–324, Oct. 2006.

[33] Mats Skoglund and Tobias Wrigstad. A mode system for read-only references in Java. In *FTfJP*, June 2001.

[34] Matthew S. Tschantz and Michael D. Ernst. Javari: Adding reference immutability to Java. In *OOPSLA*, pages 211–230, Oct. 2005.

[35] Jan Vitek and Boris Bokowski. Confined types. In *OOPSLA*, pages 82–96, Nov. 1999.

[36] Tobias Wrigstad. *Ownership-Based Alias Management*. PhD thesis, Royal Institute of Technology, Sweden, May 2006.

[37] Tobias Wrigstad and Dave Clarke. Existential owners for ownership types. *J. Object Tech.*, 6(4):141–159, May–June 2007.

[38] Yoav Zibin. Featherweight Ownership and Immutability Generic Java (FOIGJ). Technical Report 10-16, ECS, VUW, June 2010. `http://ecs.victoria.ac.nz/Main/TechnicalReportSeries`.

[39] Yoav Zibin, Alex Potanin, Mahmood Ali, Shay Artzi, Adam Kieżun, and Michael D. Ernst. Object and reference immutability using Java generics. In *ESEC/FSE*, pages 75–84, Sep. 2007.

[40] Yoav Zibin, Alex Potanin, Paley Li, Mahmood Ali, and Michael D. Ernst. OIGJ with owners as modifiers. Technical Report 10-15, ECS, VUW, January 2010.