

A Program Specialization Relation Based on Supercompilation and its Properties

Andrei V. Klimov*

Keldysh Institute of Applied Mathematics
Russian Academy of Sciences
4 Miusskaya sq., Moscow, 125047, Russia
klimov@keldysh.ru

Abstract. An input-output relation for a wide class of program specializers for a simple functional language in the form of Natural Semantics inference rules is presented. It covers *polygenetic* specialization, which includes deforestation and supercompilation, and generalizes the author's previous paper on specification of *monogenetic* specialization like partial evaluation and restricted supercompilation.

The *specialization relation* expresses the idea of *what* is to be a specialized program, avoiding as much as possible the details of *how* a specializer builds it. The relation specification follows the principles of Turchin's supercompilation and captures its main notions: *configuration*, *driving*, *generalization of a configuration*, *splitting a configuration*, as well as *collapsed-jungle driving*. It is virtually a formal definition of supercompilation abstracting away the most sophisticated parts of supercompilers—strategies of *configuration analysis*.

Main properties of the program specialization relation—idempotency, transitivity, soundness, completeness, correctness—are formulated and discussed.

Keywords: specialization, input-output relation, partial evaluation, supercompilation, correctness.

1 Introduction

Program specialization is an equivalence transformation. A specializer *spec* maps a source program p to a residual program q , which is equivalent to p on a given subset D of the domain of the program p : $q = \text{spec}(p, D)$. The equivalence of the source and residual programs is understood *extensionally*, that is, nonconstructively: $p \approx_D q$ if for all $d \in D$: $p(d) = q(d)$ or both $p(d)$ and $q(d)$ do not terminate. The correctness of specializers is usually proven by reducing it to the extensional equivalence [4,2,11]:

$$q = \text{spec}(p, D) \Rightarrow p \approx_D q.$$

* Supported by Russian Foundation for Basic Research projects No. 06-01-00574-a and No. 08-07-00280-a and Russian Federal Agency of Science and Innovation project No. 2007-4-1.4-18-02-064.

In this paper we define a constructive, *intensional* relation of specialization, that is, a relation of equivalence of source and residual programs, which many specialization methods satisfy, including partial evaluation [5], deforestation [18], supercompilation [16,17]. Let $p \simeq_D q$ denote such intensional equivalence of a source program p and a residual one q on a set D of input data values. To be correct, it must be a subset of the extensional relation, $(\simeq) \subset (\approx)$, that is, $p \simeq_D q \Rightarrow p \approx_D q$. The intensional relation provides a shorter way for proving the correctness of specializers than the extensional equivalence does:

$$q = \text{spec}(p, D) \Rightarrow p \simeq_D q \Rightarrow p \approx_D q.$$

The specialization relation is defined in this paper by inference rules in the style of Natural Semantics [6]. The rules serve as formal specification of a wide class of specializers. By the nature of Natural Semantics, the specification allows for automated derivation of specializers as well as checkers of the correctness of residual programs, which can help in debugging practical specializers. Some notions (e.g. driving) are defined precisely enough to unambiguously derive the corresponding algorithm by the well-known methods. Other notions (e.g. generalization and splitting configurations) are defined with certain degrees of freedom to allow for various decision-taking algorithms and strategies.

The specialization relation is based on the ideas of supercompilation, but agrees with partial evaluation and deforestation as well. The inference rules model at abstract level the operational behavior of supercompilers. All essential notions of supercompilation are captured: *configuration*, *driving*, *generalization of a configuration*, *splitting a configuration*, as well as *collapsed-jungle driving* [12,13], while abstracting from the problems of algorithmic decisions of when, what and how to generalize and when to terminate.

In paper [8] a similar specialization relation was presented for the simpler case of *monogenetic* specialization [10] where any program point in the residual program is produced from a single program point of the source program. In this paper the relation definition is developed further to *polygenetic* specialization [10] where a residual program point is produced from one or several source program points. Monogenetic specialization includes partial evaluation but excludes deforestation and supercompilation. Polygenetic specialization covers all of them. For completeness sake, the basic notions and the definition of driving from [8] are repeated in Section 2, Figs. 6, 7.

The main contributions of the paper are as follows:

- a complete formal definition of what supercompilation is, in form of an input-output specialization relation, is given;
- several interesting properties that the presented specialization relation obeys are formulated and related to each other: idempotency, transitivity, soundness, completeness, correctness.

The paper is organized as follows. A simple object language, which is both the source and target language of specializers, is presented in Section 2.1 together with semantic domains for interpretation and supercompilation. In Section 2.2

$k \in Atom$	atomic data	$k ::= True \mid False \mid Nil \mid \dots$
$x \in Data$	ground data	$x ::= k \mid Cons \ x \ x$
$z \in CData$	configuration data	$z ::= k \mid Cons \ z \ z \mid u \mid l$
$a \in Arg$	source arguments	$a ::= z \mid Cons \ a \ a \mid v$
$d \in Prim$	source primitives	$d ::= a \mid \mathbf{fst} \ v \mid \mathbf{snd} \ v$
$s \in Term$	source program terms	$\quad \mid \mathbf{cons}^? \ v \mid \mathbf{equ}^? \ v \ a$
$r \in Term$	residual program terms	$s ::= d$
$v \in Var$	source program variables	$\quad \mid \mathbf{if} \ v \ \mathbf{then} \ s_1 \ \mathbf{else} \ s_2$
$u \in Var$	residual program variables and configuration variables	$\quad \mid \mathbf{let} \ v = s_1 \ \mathbf{in} \ s_2$
		$\quad \mid \mathbf{call} \ f \ b$
$l \in LVar$	liaison variables	
$f \in FName$	function names	$Prog = FName \rightarrow Term$
$p \in Prog$	source programs	$Args = Var \rightarrow Arg$
$q \in Prog$	residual programs	$Contr = Var \rightarrow CData$
$b \in Args$	argument bindings	$Expl = LVar \rightarrow MConf$
$c \in Contr$	contractions	$MConf = Term$
$\Delta \in Expl$	explications	$PConf = CData \times Expl$
$d, s \in MConf$	monogenetic configurations	$CMap = FName \rightarrow PConf$
$z \mid \Delta \in PConf$	polygenetic configurations	
$m \in CMap$	mapping of residual function names to configurations	

Fig. 1. Object language syntax and semantic domains

the notion of configuration is introduced. In Section 2.3 the operation of substitution as it is used in this paper is defined. Section 2.4 discusses the supercompilation notion of contraction. Sections from 3 to 5 present the definition of the specialization relation: in Section 3 the specifics of our definition of the language semantics and specialization is explained; in Section 4 the semantics and driving of the language primitives and in Section 5 the semantics and specialization of control program terms are specified. In Section 6 the most interesting properties of the specialization relation are formulated and discussed, and in Section 7 we conclude.

2 Basic Notions

2.1 Object Language and Semantic Domains

Figure 1 contains the definition of the abstract syntax of the object language together with semantic domains for interpretation and specialization. It is a simple first-order functional language. It has conventional control constructs **if-then-else**, **let-in**, **call**, adjusted a bit to make the specialization inference rules simpler. Figure 2 shows an example of a program and an initial configuration for specialization.

$\text{rev } v_1 = \text{loop } v_1 []$ $\text{loop } [] v_2 = v_2$ $\text{loop } (v_4 : v_5) v_2 = \text{loop } v_5 (v_4 : v_2)$	<p>— a program in Haskell</p>
$p = \left\{ \begin{array}{l} \text{rev} \mapsto \text{call loop } \{v_1 \mapsto v_1, v_2 \mapsto \text{Nil}\}, \\ \text{loop} \mapsto \text{let } v_3 = \text{cons? } v_1 \text{ in} \\ \quad \text{if } v_3 \\ \quad \text{then let } v_4 = \text{fst } v_1 \text{ in} \\ \quad \quad \text{let } v_5 = \text{snd } v_1 \text{ in} \\ \quad \quad \text{call loop } \{v_1 \mapsto v_5, v_2 \mapsto \text{Cons } v_4 v_2\} \\ \quad \text{else } v_2 \end{array} \right\}$	<p>— the same program in the object language</p>
$s_0 = \text{call rev } \{v_1 \mapsto \text{Cons A (Cons } u_1 (\text{Cons B } u_2))\}$	<p>— an initial term</p>
$z \mid \Delta = l_0 \mid \{l_0 \mapsto s_0\}$	<p>— an initial configuration</p>

Fig. 2. An example of a program p and an initial configuration

Data. A data domain *Data* is a *constructor-based* domain recursively defined from a set of atoms *Atom* by applying a binary constructor *Cons*. The set *Atom* contains at least *True*, *False*, and *Nil*.

Any constructor-based domain has the nice property that it can be easily extended to meta-data without the need for encoding. In particular, a constant in program code coincides with the value it represents. That is, $\text{Data} \subset \text{Term}$, where *Term* is the domain of program terms.

Configuration Data. Another extension of *Data* originates from the need to constructively represent sets of data values and sets of program states. The basic method to represent sets is to embed free variables into the representation of data. In the theory of supercompilation such variables are referred to as *configuration variables*. The general principle is that a configuration variable, $u \in \text{Var}$, can occur in any position where a ground value is allowed.

A characteristic feature of supercompilation, which is preserved in our specialization relation definition, is that configuration variables become residual program variables.

To specify polygenetic specialization, we use a representation of configurations in form of directed acyclic graphs [12,13], which we define in the next section and refer to as *polygenetic configurations* or *polyconfigurations* for short. It requires one more extension of the data domain by so called *liaison*¹ variables, $l \in \text{LVar}$, bound variables that link positions in terms to subterms.

¹ The term is due to V. Turchin, who suggested the use of such a representation of configurations in supercompilers in 1970s.

The data and configuration data domains, $Data$ and $CData$, and the $Term$ domain are embedded in each other: $Data \subset CData \subset Term$.

Primitives. Data values are analyzed by primitive predicates $\mathbf{equ?} v a$ (are two values equal?) and $\mathbf{cons?} v$ (is the value of a variable v a term of the form $\mathbf{Cons} z_1 z_2 ?$), which return atoms \mathbf{True} or \mathbf{False} , and selectors $\mathbf{fst} v$ and $\mathbf{snd} v$, which require the value of v to be a \mathbf{Cons} term and return its first and second argument respectively.

To avoid dealing with exceptions, we impose a context restriction on selectors $\mathbf{fst} v$ and $\mathbf{snd} v$: they can occur only on the positive branch of an \mathbf{if} -term with the conditional $\mathbf{cons?} v$.

Control. Control terms $\mathbf{if-then-else}$ and $\mathbf{let-in}$ are the usual conditional term and \mathbf{let} binding respectively. The following restriction is imposed for the simplicity of the specialization definition: the conditional must be a variable v that is bound to a conditional primitive, $\mathbf{equ?}$ or $\mathbf{cons?}$, by an enclosing \mathbf{let} term, e.g.

$$\mathbf{let} v = \mathbf{cons?} v_1 \mathbf{in} \dots \mathbf{if} v \mathbf{then} s_1 \mathbf{else} s_2 \dots$$

A program is a finite mapping of function names to program terms.

A function call, which usually looks like $f(a_1, \dots, a_n)$, is written in our language as

$$\mathbf{call} f \{v_1 \mapsto a_1, \dots, v_n \mapsto a_n\}$$

where v_1, \dots, v_n are the free variable names of the term that the name f is bound to in the program.

For simplicity, terms are in the so called *administrative normal form*, that is, the arguments of all terms except the $\mathbf{let-in}$ term and the \mathbf{then} and \mathbf{else} branches of the \mathbf{if} term, has trivial form: $v \in \mathbf{Var}$ or $a \in \mathbf{Arg}$.

Notation.

- $\mathbf{FVars}(t)$ denotes the set of free variables occurring in term t .
- $\mathbf{LVars}(t)$ denotes the set of liaison variables in a term or in a configuration.
- $\mathbf{Dom}(m)$ and $\mathbf{Rng}(m)$ denote the domain and range of mapping m respectively.

2.2 Configuration

While an interpreter runs a program on a ground data, a specializer runs a source program on a set of data. A representation of a program state in interpretation and that of a set of states in specialization is referred to as a *configuration*. We follow the general rule of construction of the notion of the configuration in a supercompiler from that of the program state in an interpreter that reads as follows: add configuration variables to the data domain, and allow the variables to occur anywhere where an ordinary ground value can occur. A configuration

represents the set that is obtained by replacing all configuration variables with all possible values.

There are two kinds of configurations, which we refer to as *monogenetic configurations* and *polygenetic configurations*, or *monoconfigurations* and *polyconfigurations* for short. In our previous work [8] monoconfigurations were actual configurations. In this work monoconfigurations are used as configurations in the rules for primitives, and polyconfigurations are configurations in the rules for control terms. Polyconfigurations comprise monoconfigurations.

Syntactically, a *monoconfiguration* is a source program term, in which program variables are replaced with their values.²

A *polyconfiguration* is a representation of a program state as directed acyclic graphs (as in [12,13]). A polyconfiguration can be thought of as obtained from a monoconfiguration in these steps:

- 1) decompose a monoconfiguration into a topmost term and some subterms;
- 2) bind the subterms to fresh liaison variables;
- 3) put the liaison variables into the topmost term instead of respective subterms;
- 4) decompose some subterms analogously.

A polyconfiguration is denoted by $z \mid \Delta$, where z is the topmost subterm, and Δ the binding of liaison variables to terms (monoconfigurations), referred to as an *explication*.³ Topmost terms are restricted to $z \in CData$, that is, primitives and control terms must be picked out (*explicated*) and put into the binding. As an example, see the initial configuration in Fig. 2:

$$z \mid \Delta = l_0 \mid \{l_0 \mapsto \mathbf{call} \text{ rev } \{v_1 \mapsto \mathbf{Cons} \ A \ (\dots)\}\}.$$

While initial configurations are usually trees, during specialization polyconfigurations form directed acyclic graphs in general. In the case of the applicative evaluation order, the polyconfiguration is a call stack. However, the inference rules do not fix the order of evaluation, and we consider the explication as an unordered set of bindings. Each time we write $\{l \mapsto s\} \Delta$, we imply $\Delta_1 \{l \mapsto s\} \Delta_2$ for some Δ_1 and Δ_2 such that $\Delta = \Delta_1 \Delta_2$.

When it is clear from the context what kind of configuration is meant, we say just a *configuration*. It is a monogenetic configuration when the rules of driving of primitives are considered, and a polygenetic configuration in other cases.

2.3 Substitution

To avoid the ambiguity of traditional postfix notation for substitution $t\theta$ when it is used in inference rules (either juxtaposition, or application of substitution), we lift up the substitution symbol θ and use a kind of power notation t^θ .

² An alternative is to keep program terms untouched and to represent the monoconfiguration as a pair consisting of a program term and an environment that binds program variables to their values. Although this representation is more common in implementations of interpreters and specializers, we prefer to substitute the environment into the term for conciseness of inference rules.

³ The term is due to V. Turchin.

$_{}^- : FName \times Prog \rightarrow Term$	
$f\{\dots, f \mapsto s, \dots\} = s$	
$_{}^- : Term \times (Contr \cup Args \cup Expl) \rightarrow Term$	
z^θ	$= z$ if $z \in Atom \cup Var \cup LVar$ and $z \notin Dom(\theta)$
$v\{\dots, v \mapsto z, \dots\}$	$= z$
$\overline{Cons\ a_1\ a_2}^\theta$	$= Cons\ a_1^\theta\ a_2^\theta$
$\overline{fst\ v}^\theta$	$= fst\ v^\theta$
$\overline{snd\ v}^\theta$	$= snd\ v^\theta$
$\overline{cons?\ v}^\theta$	$= cons?\ v^\theta$
$\overline{equ?\ v\ a}^\theta$	$= equ?\ v^\theta\ a^\theta$
$\overline{if\ v\ then\ s_1\ else\ s_2}^\theta$	$= if\ v^\theta\ then\ s_1^\theta\ else\ s_2^\theta$
$\overline{let\ v = s_1\ in\ s_2}^\theta$	$= let\ v = s_1^\theta\ in\ s_2^\theta$
$\overline{call\ f\ b}^\theta$	$= call\ f\ b^\theta$
$_{}^- : Args \times (Contr \cup Args \cup Expl) \rightarrow Args$	
$\{v_1 \mapsto a_1, \dots, v_n \mapsto a_n\}^\theta$	$= \{v_1 \mapsto a_1^\theta, \dots, v_n \mapsto a_n^\theta\}$
$_{}^- : Expl \times (Contr \cup Expl) \rightarrow Expl$	
$\{l_1 \mapsto s_1, \dots, l_n \mapsto s_n\}^\theta$	$= \{l_1 \mapsto s_1^\theta, \dots, l_n \mapsto s_n^\theta\}$
$_{}^- : PConf \times (Contr \cup Expl) \rightarrow PConf$	
$\overline{z \mid \Delta}^\theta$	$= z^\theta \mid \Delta^\theta$

Fig. 3. The definition of substitution t^θ for those domains which it is applied to in the specialization relation definition

Thus t^θ denotes the replacement of all occurrences of variables $v \in Dom(\theta)$ in t with their values from a binding θ . Notation $t^{\eta\theta}$ means sequential application of substitutions η and θ to t in that order. When the argument of a substitution is unclear, it is over-lined, e.g. $\overline{a\ b\ c\ d}$.

The bindings listed in Fig. 1 are used as substitutions as follows:

- f^p gets the term bound to a function name f in a program p ;
- f^{pb} builds a monoconfiguration from a program term f^b and the argument binding b from a monoconfiguration **call** $f\ b$:

- $\frac{}{z \mid \Delta} \{l \mapsto z'\}$ substitutes a value z' for a liaison variable l in a polyconfiguration $z \mid \Delta$;
- s^c and $\frac{}{z \mid \Delta}^c$ contracts a monoconfiguration s and a polyconfiguration $z \mid \Delta$ respectively by replacing a configuration variable u with the configuration value z' bound to u by a contraction $c = \{u \mapsto z'\}$.

See Fig. 3 for the formal definition of the substitution operation.

2.4 Contraction

After evaluation of a conditional, the current configuration divides into two sub-configurations, the initial configurations of the positive and negative branches. In our definition the subconfigurations represent the subsets precisely, that is, they are disjoint.

There are two Boolean primitives, **equ**? $v a$ and **cons**? v , in the object language. After substitution of configuration values into the arguments of the primitives, they ultimately reduce (by rules in Fig. 7 below) to the following checks on configuration variables that produce branching in residual code: **equ**? $u k$, **equ**? $u u'$, and **cons**? u , where k is an atom, u and u' configuration variables.

For each of the primitives, the set of the values of a configuration variable u that go to the positive branch can be represented by a substitution $\{u \mapsto k\}$, $\{u \mapsto u'\}$, or $\{u \mapsto \text{Cons } u_1 u_2\}$, where $k \in \text{Atom}$, u_1 and u_2 are new configuration variables. Such a substitution is referred to as a *contraction*. Being applied to a configuration, it produces a configuration representing a subset of the original one.

For uniformity’s sake, the opposite case of “negative” information—the set of the values that go to the negative branch—is represented by a substitution as well. To achieve this, we assume the representation of a configuration variable contains a *negative set*: a set of “negative entities” the variable must be *unequal to*. The negative entities are atoms, configuration variables, and the word **Cons**, which represents inequality to all terms of the form **Cons** $z_1 z_2$.

We denote the operation to add an entity n to the negative set of a configuration variable u by u^{-n} . Thus the following substitutions are *negative contractions*: $\{u \mapsto u^{-k}\}$, $\{u \mapsto u^{-u'}\}$, and $\{u \mapsto u^{-\text{Cons}}\}$.

2.5 Definition of Relations by Inference Rules

In the next sections the input-output interpretation and specialization relations are defined in the style of Natural Semantics.

The relations are formalized by judgments listed and commented in Fig. 4. A relation holds for some terms if the corresponding judgment is deducible.

The axioms, from which, and the inference rules, by which the judgments are deduced, are presented in Figs. 5–11.

When we say just “deducible”, it means “deducible from all axioms and by all inference rules presented in this paper”. When only part of axioms or rules is used (e.g. for defining interpretation), it is mentioned explicitly.

$d \rightsquigarrow z$	Interpretation and transient driving of a primitive d produces a value z . (The term d is a monoconfiguration.)
$d \prec \mathcal{T}(c_1, c_2)$	Driving with branching: Driving of a primitive d produces a branching represented by a residual conditional term $\mathcal{T}(_, _)$ with two free positions for positive and negative branches. In the right-hand side $\mathcal{T}(c_1, c_2)$ these positions are occupied by contractions c_1 and c_2 . The contractions being substituted to the configuration before the branching produce the initial configurations for the positive and negative branches respectively.
$p : s \Rightarrow q : r$	Specialization: A residual program q with an initial term r is a specialization of a source program p with an initial term s . (The term s is a monoconfiguration. The initial polyconfiguration is $l \mid \{l \mapsto s\}$.)
$p : z \mid \Delta \rightsquigarrow m : r$	Specialization to a term: A residual term r is a specialization of a source program p with an initial polyconfiguration $z \mid \Delta$ with respect to a mapping m of residual function names to polyconfigurations.
$p : z \mid \Delta \rightsquigarrow \emptyset : z'$	Interpretation as a subset of specialization: In the case where the residual program is empty, and hence $m = \emptyset$, and the residual term z' is a configuration value, $z' \in CData$, the previous judgment means interpretation or transient driving.
$p : s \rightarrow z$	Interpretation as a subset of specialization: This is a shortcut notation for the particular case of the above judgment of the form $p : l \mid \{l \mapsto s\} \rightsquigarrow \emptyset : z$.
$p : s \overset{\circ}{\rightarrow} z$	Interpretation (semantics): This judgement is equivalent to $p : s \rightarrow z$ with the requirement that it be deduced using only the interpretation and transient driving rules marked with \circ . This is proper interpretation when $FVars(s) = \emptyset$ and hence z is a ground value, $z \in Data$.

Fig. 4. Judgments

3 Interpretation as a Subset of Specialization

A specialization relation is an extension of the semantics of a language, which is usually a function (for deterministic languages). We could give the language semantics, develop separately the specialization relation, and then prove the statement that the specialization relation includes the semantics.

However, to save space and mental effort we follow another line. We define the specialization relation by inference rules in such a way that a subset of the rules

TR-DRV	$\frac{p : l \mid \{l \mapsto s\} \rightsquigarrow \emptyset : z}{p : s \rightarrow z}$	$z \in CData$
INT	$\frac{p : s \rightarrow z}{p : s \overset{\circ}{\rightarrow} z}$	if inferred with $\overset{\circ}$ -rules only $z \in CData$

Fig. 5. Transient driving and interpretation as a subset of specialization

defines the semantics, interpretation. The interpretation rules have labels marked with $\overset{\circ}$. The rules with unmarked labels extend interpretation to specialization.

The specialization judgments $p : s \Rightarrow q : r$ and $p : z \mid \Delta \rightsquigarrow m : r$ contain a residual program q and an auxiliary mapping m (explained in Section 5.1 below), which has the same domain, $\text{Dom}(q) = \text{Dom}(m) \subset F\text{Name}$, the set of residual function names. When there are no residual functions, that is, $p = q = \emptyset$, the residual code is merely a tree represented by the term r that cannot contain **call** terms. The particular case where the residual term r is a value virtually defines the semantics of the language. It can be proven that judgments of this form are deducible by means of the interpretation rules only.

We define a shortcut notation $p : s \rightarrow z$ for this case by rule TR-DRV, and denote by a circle over the arrow the fact that it is inferred by the interpretation rules only: $p : s \overset{\circ}{\rightarrow} z$ (see rule INT in Fig. 5).

The interpretation rules also define *transient driving*, which is the basic case of driving where the configuration and liaison variables do not prevent a specializer from unambiguously performing a step. The case of proper interpretation can be distinguished from the case of transient driving by the restriction that the initial configuration does not contain configuration variables, $F\text{Vars}(s) = \emptyset$.

Definition 1 (Interpretation, semantics). *A source program p with an initial term s (in which arguments have been substituted) without configuration variables, $F\text{Vars}(s) = \emptyset$, evaluates to a term $x \in \text{Data}$ if the following judgment is deducible:*

$$p : s \overset{\circ}{\rightarrow} x$$

4 Interpretation and Specialization of Primitives

4.1 Interpretation and Transient Driving of Primitives

A judgment of the form $d \rightsquigarrow z$ means *interpretation or transient driving* of a primitive monoconfiguration d (i.e., a source program primitive term, in which program variables have been replaced with their values) produces a value z .

In the rules we distinguish between the case where configuration values can occur and the case where only ground values occur by the names of free variables: $k \in \text{Atom} \subset \text{Data}$, $x \in \text{Data}$, $z \in C\text{Data}$, that is, k and x range over ground

I-VALUE [◦]	$z \rightsquigarrow z$	$z \in CData$
I-FST [◦]	$\mathbf{fst} (\text{Cons } z_1 z_2) \rightsquigarrow z_1$	
I-SND [◦]	$\mathbf{snd} (\text{Cons } z_1 z_2) \rightsquigarrow z_2$	
I-CONS-T [◦]	$\mathbf{cons?} (\text{Cons } z_1 z_2) \rightsquigarrow \text{True}$	
I-CONS-F [◦]	$\mathbf{cons? } k \rightsquigarrow \text{False}$	$k \in Atom$
I-EQ-T [◦]	$\mathbf{equ? } z z \rightsquigarrow \text{True}$	
I-EQ-F [◦]	$\mathbf{equ? } x_1 x_2 \rightsquigarrow \text{False}$	$x_1 \neq x_2$ $x_1, x_2 \in Data$
D-EQ-CK	$\mathbf{equ?} (\text{Cons } z_1 z_2) k \rightsquigarrow \text{False}$	$k \in Atom$
D-EQ-CC	$\frac{\mathbf{equ? } z_{1i} z_{2i} \rightsquigarrow \text{False}}{\mathbf{equ?} (\text{Cons } z_{11} z_{12}) (\text{Cons } z_{21} z_{22}) \rightsquigarrow \text{False}}$	$i \in \{1, 2\}$

Fig. 6. Interpretation and transient driving of primitives

values and z ranges over configuration values that may contain configuration and liaison variables. Note that only rules I-CONS-F[◦] and I-EQ-F[◦], which defines inequality, require values to be ground.

The last two rules in Fig. 6, D-EQ-CK and D-EQ-CC, and the rules in Fig. 7 that infer judgments of the form $d \rightsquigarrow z$ define the cases of transient driving that are not covered by the interpretation rules.

The rules in Fig. 7 mean:

- D-CONS-F and D-EQ-UCF — returning False in the case where a configuration variable u contains in its negative set the symbol Cons and hence is unequal to any Cons term;
- D-EQ-COM — commutativity of **equ?**;
- D-EQ-UKF — returning False in the case where a configuration variable u contains in its negative set the atom k it is compared to;
- D-EQ-UUF — returning False in the case where a configuration variable u_1 contains in its negative set the configuration variable u_2 it is compared to.

4.2 Driving with Branching

A judgment of the form $d \prec \mathcal{T}(c_1, c_2)$ means *driving*⁴ of a primitive monoconfiguration d produces a branching in residual code represented by a conditional term $\mathcal{T}(-, -)$ with two free positions for positive and negative branches and two contractions c_1 and c_2 . The contractions c_1 and c_2 , being applied as substitutions to the configuration d , divide it into two subconfigurations, which are initial configurations for positive and negative branches respectively.

For the sake of notation brevity, the contractions c_1 and c_2 occupy in $\mathcal{T}(-, -)$ the positions where the residual terms for the positive and negative branches will occur in the final residual code.

Figure 7 contains the rules that infer branching in residual code for the source primitives **cons?** and **equ?**. The branching happens when a configuration variable u (or two variables u_1 and u_2 in the case of the **equ?** term) prohibits from performing an evaluation step. Notice the branching rules perform no evaluation step of the source program, but just produce a residual **if** term and contractions. The proper evaluation step is performed by transient driving rules for the same primitive after contractions c_1 and c_2 has been substituted into the current configuration by rule PS-BRANCH, which produces initial configurations for branches, and transient driving of the primitive has been “invoked” by rule PS-PRIM^o in Fig. 9.

The correctness of this deduction is based on the perfectness [3] of contractions c_1 and c_2 and on the fact that after substitution of c_i into d some transient driving rule for the judgment $d^{c_i} \rightsquigarrow z$ is applicable.

In each of the three use cases of the term $\mathcal{T}(-, -)$ in Fig. 7 it meets the following property: the value of $\mathcal{T}(x_1, x_2)$ is either x_1 for the configuration obtained by contraction c_1 , or x_2 for the configuration obtained by contraction c_2 . The correctness of rule PS-BRANCH relies on this property.

In Fig. 7, the rules of driving of **equ?** and **cons?** terms that infer judgments of the form $d \prec \mathcal{T}(c_1, c_2)$ mean:

- Rule D-CONS defines the branching in a residual program that corresponds to a monoconfiguration of the form **cons?** u . The right-hand side consists of the residual **if** term that tests the value of the variable u , the assignments of the parts of the u value to fresh variables u_1 and u_2 on the positive branch, and two complementary contractions $\{u \mapsto \text{Cons } u_1 \ u_2\}$ and $\{u \mapsto u^{-\text{Cons}}\}$ occupying in the term $\mathcal{T}(-, -)$ the positions of the positive and negative branches respectively.
- Rules D-EQ-UK and D-EQ-UU analogously define the branchings corresponding to monoconfigurations **equ?** $u \ k$ and **equ?** $u_1 \ u_2$ in the case where there is no information in negative sets of configuration variables about the equalities under the respective tests.

⁴ What is usually called driving in the theory of supercompilation is unfolding an *infinite process tree* [17,3,1]. This sense could be captured by the ^o-rules together with rule PS-BRANCH if we consider *infinite* residual terms r . However, it would be another theory.

D-CONS-F	$\mathbf{cons?} \ u \rightsquigarrow \text{False}$	$u = u^{-\text{Cons}}$ $u \in \text{Var}$
D-CONS	$\mathbf{cons?} \ u \prec \mathbf{let} \ u_0 = \mathbf{cons?} \ u \mathbf{ in}$ $\mathbf{if} \ u_0$ $\mathbf{then let} \ u_1 = \mathbf{fst} \ u \mathbf{ in}$ $\mathbf{let} \ u_2 = \mathbf{snd} \ u \mathbf{ in}$ $\{u \mapsto \text{Cons } u_1 \ u_2\}$ $\mathbf{else} \ \{u \mapsto u^{-\text{Cons}}\}$	$u \neq u^{-\text{Cons}}$ $u_0, u_1, u_2 \text{ new}$ $u, u_i \in \text{Var}$
D-EQ-COM	$\frac{\mathbf{equ?} \ u \ z \rightsquigarrow \mathcal{T}}{\mathbf{equ?} \ z \ u \rightsquigarrow \mathcal{T}}$	$u \in \text{Var}$
D-EQ-UKF	$\mathbf{equ?} \ u \ k \rightsquigarrow \text{False}$	$u = u^{-k}$ $u \in \text{Var}$ $k \in \text{Atom}$
D-EQ-UK	$\mathbf{equ?} \ u \ k \prec \mathbf{let} \ u_0 = \mathbf{equ?} \ u \ k \mathbf{ in}$ $\mathbf{if} \ u_0$ $\mathbf{then} \ \{u \mapsto k\}$ $\mathbf{else} \ \{u \mapsto u^{-k}\}$	$u \neq u^{-k}$ $u_0 \text{ new}$ $u, u_0 \in \text{Var}$ $k \in \text{Atom}$
D-EQ-UUF	$\mathbf{equ?} \ u_1 \ u_2 \rightsquigarrow \text{False}$	$u_1 = u_1^{-u_2}$ $u \in \text{Var}$
D-EQ-UU	$\mathbf{equ?} \ u_1 \ u_2 \prec \mathbf{let} \ u_0 = \mathbf{equ?} \ u_1 \ u_2 \mathbf{ in}$ $\mathbf{if} \ u_0$ $\mathbf{then} \ \{u_1 \mapsto u_2\}$ $\mathbf{else} \ \{u_1 \mapsto u_1^{-u_2}, \ u_2 \mapsto u_2^{-u_1}\}$	$u_1 \neq u_1^{-u_2}$ $u_0 \text{ new}$ $u_i \in \text{Var}$
D-EQ-UCF	$\mathbf{equ?} \ u \ (\text{Cons } z_1 \ z_2) \rightsquigarrow \text{False}$	$u = u^{-\text{Cons}}$ $u \in \text{Var}$ or $u \in \text{FVars}(z_1)$ or $u \in \text{FVars}(z_2)$
D-EQ-UC	$\frac{\mathbf{cons?} \ u \prec \mathcal{T}}{\mathbf{equ?} \ u \ (\text{Cons } z_1 \ z_2) \prec \mathcal{T}}$	$u \neq u^{-\text{Cons}}$ $u \in \text{Var}$ $u \notin \text{FVars}(z_1)$ $u \notin \text{FVars}(z_2)$

Fig. 7. Driving of primitives

- Rule D-EQ-UC reduces the test of equality of a variable u and a Cons term, $\mathbf{equ?} u (\text{Cons } z_1 z_2)$, to the test of whether u is a Cons term or not, $\mathbf{cons?} u$. Recall the judgment $d \prec \mathcal{T}(c_1, c_2)$ defines no evaluation step, only a branching. So, the rule reads as follows: to advance driving of the configuration $\mathbf{equ?} u (\text{Cons } z_1 z_2)$ residualize the same branching as for the configuration $\mathbf{cons?} u$. After that, the evaluation step will be performed by the driving rules for the $\mathbf{equ?}$ term.

5 Interpretation and Specialization of Control Terms

A judgment of the form $p : z \mid \Delta \rightsquigarrow m : r$ inferred by the rules in Fig. 9 asserts that a residual term r is a specialization of a source program p with an initial polyconfiguration $z \mid \Delta$ with respect to a mapping m of residual function names to polyconfigurations.

5.1 Correspondence between Residual Functions and Configurations

The mapping $m : \text{Dom}(q) \rightarrow P\text{Conf}$ assigns meaning to the residual **call** terms occurring in r . This can be explained in terms of the language semantics as follows. For all values of the configuration variables of the configuration $z \mid \Delta$ and the residual term r ,⁵ evaluation of the configuration $z \mid \Delta$ with the source program p gives the same result as evaluation of the term r in the following two steps:

- 1) for each subterm of the form **call** $f_i b_i$ occurring in r , evaluate the configuration $f_i^{mb_i} = \overline{z_i \mid \Delta_i}^{b_i}$ with the source program p , where $z_i \mid \Delta_i$ is the configuration bound to the function name f_i by the mapping m ;
- 2) evaluate the term r using the thus obtained values of the **call** terms.

In other words, each residual function body in q is equivalent to the corresponding configuration in m . In its turn, the equivalence can be stated by the specialization relation.

Notice the sense of the mapping m is inherently recursive: on the one hand, m is used in the definition of the specialization relation; on the other hand, its sense is based on the semantics of the language, which is a subset of the specialization relation.

To escape from the vicious circle, we define a relation between m and source and residual programs p, q . We refer to it as *consistency of m with p and q* .

Definition 2 (Consistency). *A mapping $m : \text{Dom}(q) \rightarrow P\text{Conf}$ of residual function names to configurations is consistent with source and residual programs p and q if for every residual function name $f \in \text{Dom}(q)$ the following judgment is deducible:*

$$p : f^m \rightsquigarrow m : f^q$$

provided the judgment is not inferred immediately from axiom PS-GEN.

⁵ Note $\text{FVars}(z \mid \Delta) = \text{FVars}(r)$.

$$\text{SPEC} \quad \frac{p : l \mid \{l \mapsto s\} \rightsquigarrow m : r}{p : s \Rightarrow q : r} \quad \text{if } m \text{ is consistent with } p \text{ and } q$$

Fig. 8. Specialization relation

The equivalence of configurations and residual terms represented by the mapping m is virtually an inductive hypotheses, the deduction of the judgments in the definition being an induction step. When proving by induction, care must be taken to avoid premature use of the inductive hypothesis. This is the role of the provision in the definition. Otherwise, the residual program may contain a loop that is absent in the source program, and the residual program may not terminate when the source program terminates.

5.2 Specialization Relation

Now we are ready to define the specialization relation, a relation between pairs consisting of a program and an initial term. We denote the pairs by $p : s$ and $q : r$ for a source program pair and a residual program pair respectively, $p, q \in \text{Prog}$, $s, r \in \text{CData}$. Only such terms s and r that have the same configuration variables can relate by the specialization relation, $\text{FVars}(s) = \text{FVars}(r)$.

For the sake of uniformity of input and output, we consider the specialization relation over plain terms (monoconfigurations) rather polygenetic configurations. An initial term s corresponds to the initial polyconfiguration $l \mid \{l \mapsto s\}$.

Definition 3 (Specialization). *A pair $q : r$ of a residual program q and an initial term r is a specialization of a pair $p : s$ of a source program p and an initial term s if there exist a mapping m of residual function names to configurations, consistent with p and q , such that the following judgment is deducible:*

$$p : l \mid \{l \mapsto s\} \rightsquigarrow m : r.$$

We denote the fact that pairs $p : s$ and $q : r$ satisfy the specialization relation by judgment $p : s \Rightarrow q : r$. Formally it is defined by rule SPEC in Fig. 8.

5.3 Rules for Control Terms

Figure 9 contains the main part of the specialization relation definition for the control terms.

Interpretation of Control Terms. Axiom PS-BASE^o asserts the evident fact that a constructor term $z \in \text{CData}$ is equivalent to itself considered as either a configuration, or a residual term.

PS-BASE [◦]	$p : z \mid \{\} \rightsquigarrow m : z$	$z \in CData$
PS-PRIM [◦]	$\frac{d \rightsquigarrow z'}{p : z \mid \Delta \rightsquigarrow m : r}$	$d \in Prim$ $z' \in CData$
PS-BRANCH	$\frac{\frac{d \prec T(c_1, c_2)}{p : z \mid \{l \mapsto d\} \Delta \rightsquigarrow m : r_1} \quad p : z \mid \{l \mapsto d\} \Delta \rightsquigarrow m : r_2}{p : z \mid \{l \mapsto d\} \Delta \rightsquigarrow m : T(r_1, r_2)}$	$FVars(d) \subseteq Dom(\rho)$ $d \in Prim$
PS-IF-T [◦]	$\frac{p : z \mid \{l \mapsto s_1\} \Delta \rightsquigarrow m : r}{p : z \mid \{l \mapsto \mathbf{if\ True\ then\ } s_1 \mathbf{\ else\ } s_2\} \Delta \rightsquigarrow m : r}$	
PS-IF-F [◦]	$\frac{p : z \mid \{l \mapsto s_2\} \Delta \rightsquigarrow m : r}{p : z \mid \{l \mapsto \mathbf{if\ False\ then\ } s_1 \mathbf{\ else\ } s_2\} \Delta \rightsquigarrow m : r}$	
PS-LET [◦]	$\frac{p : z \mid \{l' \mapsto s_1, l \mapsto s_2^{\{v \mapsto l'\}}\} \Delta \rightsquigarrow m : r}{p : z \mid \{l \mapsto \mathbf{let\ } v = s_1 \mathbf{\ in\ } s_2\} \Delta \rightsquigarrow m : r}$	l' new $l' \in LVar$
PS-CALL [◦]	$\frac{p : z \mid \{l \mapsto f^{pb}\} \Delta \rightsquigarrow m : r}{p : z \mid \{l \mapsto \mathbf{call\ } f \ b\} \Delta \rightsquigarrow m : r}$	$FVars(f^p) \subseteq Dom(b)$ $f \in Dom(\rho)$ $f \in FName$ $b \in Var \rightarrow Arg$
PS-GEN	$p : f^{mb} \rightsquigarrow m : \mathbf{call\ } f \ b$	$FVars(f^m) \subseteq Dom(b)$ $f \in Dom(m)$ $f \in FName$ $b \in Var \rightarrow Arg$

Fig. 9. Polygenetic specialization

Rule PS-PRIM[◦] describes the case where a monoconfiguration d taken from a polyconfiguration $z \mid \{l \mapsto d\} \Delta$ can be evaluated to a value z' , that is, $d \rightsquigarrow z'$.

Rules PS-IF-T[◦] and PS-IF-F[◦] define the semantics of the **if** term.

Rule PS-LET[◦] defines the **let** term by decomposing it into two parts and reducing to a configuration where the parts s_1 and s_2 are bound to separate liaison variables l' and l respectively.

Rule PS-CALL[◦] defines the term **call** $f b$ by picking up the body f^p of a function f from a program p and applying the argument substitution to it.

Other rules specify specialization proper.

Residualization of Conditional Term. Rule PS-BRANCH uses the result of driving of a Boolean primitive to build a branching in residual code. It was commented in Section 4.2 above.

Notice the case where a value of a conditional a is a configuration variable, is absent. It is useless due to the syntactic restriction on the term a (see Section 2.1).

Generalization. The most interesting rule is axiom PS-GEN that defines the notion of generalization of a configuration together with folding into a residual **call** term.

Reading the judgment

$$p : f^{mb} \rightsquigarrow m : \mathbf{call} f b$$

from left to right in terms of production of residual code rather than its specification, we say that some configuration $f^{mb} = z \mid \Delta$ is generalized to configuration $f^m = z' \mid \Delta'$ with the substitution b such that $z \mid \Delta = \overline{z' \mid \Delta'}^b$. The term **call** $f b$ is residualized, where the function f has such a body $r = f^q$ that satisfies the specialization relation

$$p : z' \mid \Delta' \rightsquigarrow m : r$$

which is implied by the consistency requirement to the mapping m of residual function names to configurations.

Splitting a Configuration. The definition of polygenetic specialization is incomplete without a rule that allows for *composition* of configurations if inference rules are read forwards, or *splitting* a configuration into two ones if inference rules are read backwards. Such rule PS-SPLIT is presented in Fig. 10.

Rules PS-GEN and PS-SPLIT are the only rules that do not allow for unambiguously constructing the specialization algorithm from the inference rules. They reveal the place in construction of supercompilers where decision-taking strategies when and how to generalize and when and how to split configurations are to be developed. This is the most sophisticated part of supercompilers. The *quality* of residual programs depends mainly on them, while the *correctness* is guaranteed by the mere fact the result matches these inference rules.

PS-SPLIT	
$\frac{\frac{p : z \mid \Delta_1 \rightsquigarrow m : r_1}{p : l \mid \{l \mapsto s\} \Delta_2 \rightsquigarrow m : r_2}}{p : z \mid \Delta_1 \{l \mapsto s\} \Delta_2 \rightsquigarrow m : \mathbf{let} \ u = r_2 \ \mathbf{in} \ r_1}$	$\begin{aligned} V_1 \cap V_2 &= \{l\} \\ u &\text{ new} \\ u &\in \text{Var} \end{aligned}$
<p>where $V_1 = \text{LVars}(z \mid \Delta_1)$ $V_2 = \text{LVars}(\{l \mapsto s\} \Delta_2)$</p>	

Fig. 10. Splitting a Configuration

PS-COLLAPSE	$\frac{\frac{p : z \mid \{l \mapsto s\} \Delta \rightsquigarrow m : r}{p : z \mid \{l \mapsto s, l' \mapsto s\} \Delta \rightsquigarrow m : r}}{p : z \mid \{l \mapsto s, l' \mapsto s\} \Delta \rightsquigarrow m : r}$	
PS-NONSTRICT [◦]	$\frac{p : z \mid \Delta \rightsquigarrow m : r}{p : z \mid \{l \mapsto s\} \Delta \rightsquigarrow m : r}$	$\begin{aligned} l &\notin \text{LVars}(z \mid \Delta) \\ l &\in \text{LVar} \end{aligned}$

Fig. 11. Extensions of interpretation and specialization

5.4 Extensions

The specialization relation can be infinitely extended by adding more and more rules that describe additional equivalences between source and residual programs. Figure 11 demonstrates two of them.

The first extension is *collapsed-jungle driving* [12,13] defined by rule PS-COLLAPSE. It avoids multiple evaluation of equal terms by deleting one of two liaison variable bindings of the form $\{l \mapsto s, l' \mapsto s\}$ and replacing all occurrences of the deleted variable l' by the variable l . The classic example of application of this rule is transformation of the naive recursive definition of the Fibonacci function with exponential complexity to the definition with linear complexity.

The second rule PS-NONSTRICT[◦] extends the relation to *non-strict semantics*, which means the possibility of evaluation of a function call without evaluation of its arguments. The specialization rules allow for arbitrary order of evaluation. All of the other rules preserve unevaluated terms even if their results are unneeded. With rule PS-NONSTRICT[◦], which removes a binding of an unused liaison variable, the relation allows for both *non-strict semantics* and *lazy evaluation* in an interpreter as well as in a specializer.

Notice rule PS-COLLAPSE is considered a specialization relation rule (having no [◦] mark), while PS-NONSTRICT[◦] is marked with [◦] as an interpretation rule. The reason for the difference is that the former rule does not change the language denotational semantics, while the latter does. The collapsed-jungle driving allows

Idempotency $\frac{p : s \Rightarrow q : r}{q : r \Rightarrow q : r}$	Transitivity $\frac{p : s \Rightarrow q : r \quad q : r^c \Rightarrow q' : r'}{p : s^c \Rightarrow q' : r'}$	Preservation of semantics by specialization $(p : s \overset{\circ}{\rightarrow} x) \Leftrightarrow (p : s \rightarrow x)$
Completeness $\frac{p : s \Rightarrow q : r \quad p : s^c \rightarrow x}{q : r^c \rightarrow x}$	Soundness $\frac{p : s \Rightarrow q : r \quad q : r^c \rightarrow x}{p : s^c \rightarrow x}$	Correctness $\frac{p : s \Rightarrow q : r}{(p : s^c \overset{\circ}{\rightarrow} x) \Leftrightarrow (q : r^c \overset{\circ}{\rightarrow} x)}$

Fig. 12. Properties of the specialization relation

for merely achieving additional speed-up by specialization, while the non-strict extension changes the termination behavior of a program.

If rule PS-NONSTRICT^o were applied only to specialization then a residual program might have a larger domain than the source program. This is often the case of practical supercompilers for strict languages (e.g. for Refal [9]), since supercompilers use lazy evaluation for achieving better results. This is the reason for the widespread myth that supercompilation in essence violates the termination behavior. However, if both source and residual programs as well as a specializer use the same semantics—either strict, or non-strict—the semantics is preserved (provided some other rules do not violate it).

6 Properties of Specialization Relation

As it is common in mathematics, relations obey more interesting properties than functions. In Fig. 12 the most important properties of the specialization relation are summed up. For readability, the statements are written out in form of inference rules. Their validity can be proven from the specialization rules.

The properties are rather natural for such a relation. Some of the properties are mandatory: preservation of semantics, soundness, completeness, and their corollary—correctness. Others—idempotency and transitivity—are additional nice properties that allow for simpler and more natural reasoning about specialization in form of a relation.

6.1 Idempotency

Intuitively, we consider returning an unchanged program to be a trivial case of specialization. One may expect that $p : s \Rightarrow p : s$ is true, that is, the specialization relation is *reflexive*. However, our rules require some specialization of function bodies always be performed, and hence many programs cannot occur in the right-hand sides of deducible judgments. In principle, it is easy to add

several rules that would allow for “doing nothing”, but we prefer the present version, which guide us in construction of non-trivial specializers. Nevertheless, any residual program is allowed to be its own specialization by our relation. Such property is referred to as *idempotency*.

Proposition 1 (Idempotency). *For all p, s, q, r such that*

$$p : s \Rightarrow q : r$$

the following judgment is deducible:

$$q : r \Rightarrow q : r$$

6.2 Transitivity

Specialization can be performed stepwise: specialization of a source program $p : s$ with respect to a part of arguments (let them be already substituted into s) followed by specialization of the residual program $q : r$ with respect to a part of the rest arguments (let them be represented by a substitution c) produces the second residual program $q' : r'$, which may be expected to be a specialization of the source program with respect to all information known so far. This property (referred to as *transitivity*) is not generally true when specializer *functions* are concerned, but it may hold for a specialization *relation*. This is indeed our case.

Proposition 2 (Transitivity). *For all p, s, c, q, r, q', r' such that*

$$\begin{aligned} p : s &\Rightarrow q : r \\ q : r^c &\Rightarrow q' : r' \end{aligned}$$

the following judgment is deducible:

$$p : s^c \Rightarrow q' : r'$$

6.3 Soundness

Consider the special case of transitivity where the second specialization is interpretation, that is, the residual program is empty, $q' = \emptyset$, and the residual term x is a configuration value, $x \in CData$. In this case transitivity means: if a residual program $q : r$ when run with arguments given by a contraction c produces some result x , the source program $p : s$ run with the same arguments also terminates and gives the same result. This property is *soundness* of specialization.

Proposition 3 (Soundness). *For all p, s, c, q, r, x such that*

$$\begin{aligned} p : s &\Rightarrow q : r \\ q : r^c &\rightarrow x \end{aligned}$$

the following judgment is deducible:

$$p : s^c \rightarrow x$$

Soundness is an immediate corollary of transitivity and a necessary condition for correctness.

6.4 Completeness

The converse property to soundness is *completeness*: if a source program $p : s$ when run with arguments c produces some result x , a residual program $q : r$ run with the same arguments also terminates and gives the same result.

Proposition 4 (Completeness). *For all p, s, q, r, c, x such that*

$$\begin{aligned} p : s &\Rightarrow q : r \\ p : s^c &\rightarrow x \end{aligned}$$

the following judgment is deducible:

$$q : r^c \rightarrow x$$

Completeness is one more necessary condition for correctness.

6.5 Preservation of Semantics

Recall we use a subset of the specialization rules as the definition of the language semantics. Hence, the fact that specialization includes interpretation is trivial. However, we must ensure that the specialization rules do not occasionally extend the semantics. Formally speaking, the following proposition must hold and it holds for our specialization relation indeed.

Proposition 5 (Preservation of semantics). *For all p, s, x such that*

$$p : s \rightarrow x$$

the following judgment is deducible:

$$p : s \xrightarrow{\circ} x$$

6.6 Correctness

Since the semantics of the object language is represented by a part of the inference rules, the correctness of the specialization relation is its internal property that can be expressed as follows.

Proposition 6 (Correctness). *For all p, s, q, r such that*

$$p : s \Rightarrow q : r$$

it holds that for all c and x the following judgments are deducible or not deducible simultaneously:

$$\begin{aligned} p : s^c &\xrightarrow{\circ} x \\ q : r^c &\xrightarrow{\circ} x \end{aligned}$$

The last two judgments mean interpretation of source and residual programs $p : s$ and $q : r$ with values supplied by a contraction c produces equal results x .

The correctness is an immediate corollary of soundness, completeness and preservation of semantics.

7 Conclusion and Related Work

This paper presents a formal specification of a class of specializers by inference rules in the style of Natural Semantics [6]. The rules define a relation between source and residual programs, which partial evaluation and supercompilation obey. The proposed intensional relation lies between algorithmic definitions of specializers and the extensional equivalence of programs.

The specialization relation definition declaratively captures the essential notions of supercompilation: *configuration*, *driving*, *generalization of a configuration*, *splitting a configuration*, as well as advanced notions like *collapsed-jungle driving* and variations of *strictness* and *laziness* of semantics, while abstracting from algorithmic problems of when, what and how to generalize and split, and when to terminate. It provides a basis for correctness proofs of supercompilers and for construction of an alternative proof of the correctness of partial evaluators [4,2]. To prove the correctness of a particular specializer we just need to prove that its inputs and outputs satisfy the specialization relation. By nature of Natural Semantics, the definition in form of inference rules allows for automated derivation of specializers that satisfy it as well as checkers of the correctness of residual programs.

An earlier version of specialization relation definition was presented at the Dagstuhl Seminar on Partial Evaluation, where only abstract [7] was published. It continues the work started in [3] and aimed at clarifying and formalizing the ideas of supercompilation. This paper gives a generalization of the definition presented in [8] from monogenetic to polygenetic case.

In Turchin's original papers [16,17] and others, the essential ideas of supercompilation and technical details of algorithms were not separated enough to give their short formal definition. Later on, several works have been done to fill this gap, e.g. [3,13,14,15]. All of them formalize the function of the supercompiler, while our work is, to our knowledge, the first attempt to define an input-output relation, which specializers based on both supercompilation and partial evaluation satisfy. The closest related work is [13,14] where the notion of the graph of configurations is formalized by inference rules that deduce the arcs of the graph.

The specialization relation obeys a number of nice properties: idempotency, transitivity and its corollary soundness, completeness, correctness, and others.

Future work will include development of specialization relation definitions for more sophisticated languages, including object-oriented ones, further investigation into their properties, and construction of supercompilers that satisfy the specialization relation and hence are provably correct.

8 Acknowledgments

The material of this paper is a refinement of Valentin Turchin's ideas of supercompilation. I would like to express my sincere gratitude to him for introducing me into this exciting field more than three decades ago and deeply influencing my life and work during all years.

I am very grateful to Sergei Romanenko and Yuri Klimov for their invaluable comments and suggestions to make the formalism simpler and more readable, and to Alexei Lisitsa for helpful advices to improve the quality of the paper.

References

1. S. M. Abramov. *Metavychislenija i ikh prilozhenija (Metacomputation and its applications)*. Nauka, Moscow, 1995. (In Russian).
2. C. Consel and S. C. Khoo. On-line and off-line partial evaluation: semantic specifications and correctness proofs. *Journal of Functional Programming*, 5(4):461–500, Oct. 1995.
3. R. Glück and A. V. Klimov. Occam’s razor in metacomputation: the notion of a perfect process tree. In P. Cousot, M. Falaschi, G. Filè, and G. Rauzy, editors, *Static Analysis Symposium. Proceedings*, volume 724 of *Lecture Notes in Computer Science*, pages 112–123. Springer-Verlag, 1993.
4. C. K. Gomard. A self-applicable partial evaluator for the lambda calculus: correctness and pragmatics. *ACM TOPLAS*, 14(2):147–172, 1992.
5. N. D. Jones, C. K. Gomard, and P. Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice-Hall, 1993.
6. G. Kahn. Natural Semantics. In F. G. Brandenburg, G. Vidal-Naquet, and W. Wirsing, editors, *STACS87. 4th Annual Symposium on Theoretical Aspects of Computer Science, Passau, Germany*, volume 247 of *Lecture Notes in Computer Science*, pages 22–39. Springer-Verlag, 1987.
7. A. V. Klimov. A specification of a class of supercompilers. In O. Danvy, R. Glück, and P. Thiemann, editors, *Draft Proceedings of the Dagstuhl Seminar on Partial Evaluation*, page 232. Technical Report WSI-96-6, Universität Tübingen, Germany, 1996.
8. A. V. Klimov. Specifying monogenetic specializers by means of a relation between source and residual programs. In *Perspectives of Systems Informatics (Proc. 6th International Andrei Ershov Memorial Conference, PSI 2006, Novosibirsk, Russia, June 27-30, 2006)*, volume 4378 of *Lecture Notes in Computer Science*, pages 248–259. Springer-Verlag, 2007.
9. A. P. Nemytykh. *Superkompilyator SCP4: Obschaya struktura (Supercompiler SCP4: general structure)*. Editorial URSS, Moscow, 2007. (In Russian).
10. S. A. Romanenko. Arity raiser and its use in program specialization. In N. D. Jones, editor, *ESOP’90, Copenhagen, Denmark*, volume 432 of *Lecture Notes in Computer Science*, pages 341–360. Springer-Verlag, May 1990.
11. D. Sands. Proving the correctness of recursion-based automatic program transformations. *Theoretical Computer Science*, 167(10), October 1996.
12. J. P. Secher. Driving in the jungle. In O. Danvy and A. Filinski, editors, *Programs as Data Objects*, volume 2053 of *Lecture Notes in Computer Science*, pages 198–217, Aarhus, Denmark, May 2001. Springer-Verlag.
13. J. P. Secher and M. H. B. Sørensen. On perfect supercompilation. In D. Bjørner, M. Broy, and A. Zamulin, editors, *Proceedings of Perspectives of System Informatics*, volume 1755 of *Lecture Notes in Computer Science*, pages 113–127. Springer-Verlag, 2000.
14. M. H. Sørensen and R. Glück. Introduction to supercompilation. In J. Hatcliff, T. Mogensen, and P. Thiemann, editors, *Partial Evaluation: Practice and Theory*, volume 1706 of *Lecture Notes in Computer Science*, pages 246–270. Springer-Verlag, 1999.

15. M. H. Sørensen, R. Glück, and N. D. Jones. A positive supercompiler. *Journal of Functional Programming*, 6(6):811–838, 1996.
16. V. F. Turchin. The language Refal, the theory of compilation and metasystem analysis. Courant Computer Science Report 20, Courant Institute of Mathematical Sciences, New York University, 1980.
17. V. F. Turchin. The concept of a supercompiler. *Transactions on Programming Languages and Systems*, 8(3):292–325, 1986.
18. P. Wadler. Deforestation: transforming programs to eliminate trees. *Theoretical Computer Science*, 73:231–248, 1990.