

Verification of Programs via Intermediate Interpretation

Alexei P. Lisitsa

Department of Computer Science,
The University of Liverpool
a.lisitsa@liverpool.ac.uk

Andrei P. Nemytykh

Program Systems Institute,
Russian Academy of Sciences*
nemytykh@math.botik.ru

We explore an approach to verification of programs via program transformation applied to an interpreter of a programming language. A specialization technique known as Turchin’s supercompilation is used to specialize some interpreters with respect to the program models. We show that several safety properties of functional programs modeling a class of cache coherence protocols can be proved by a supercompiler and compare the results with our earlier work on direct verification via supercompilation not using intermediate interpretation.

Our approach was in part inspired by an earlier work by E. De Angelis et al. (2014-2015) where verification via program transformation and intermediate interpretation was studied in the context of specialization of constraint logic programs.

Keywords: program specialization, supercompilation, program analysis, program transformation, safety verification, cache coherence protocols

1 Introduction

We show that a well-known program specialization technique called the first Futamura projection [14, 45, 20] can be used for indirect verification of some safety properties. We consider functional programs modeling a class of non-deterministic parameterized computing systems specified in a language that differs from the object programming language treated by the program specializer. Let a specializer transforming programs be written in a language \mathcal{L} and an interpreter $\text{Int}_{\mathcal{M}}$ of a language \mathcal{M} , which is also implemented in \mathcal{L} , be given. Given a program p_0 written in \mathcal{M} , the task is to specialize the interpreter $\text{Int}_{\mathcal{M}}(p_0, d)$ with respect to its first argument, while the data d of the program p_0 is unknown.

Our interest in this task has been inspired by the following works [3, 5]. The authors work in terms of *constraint logic programming* (CLP), where the constraint language is the linear arithmetic inequalities imposed on integer values of variables. They use partial deduction [25] and CLP program specialization [11, 4] methods for specializing an interpreter of a C-like language with respect to given programs, aiming at verification of the C-like imperative specifications with respect to the postconditions defined in CLP and defining the same functions (relations) as done by the corresponding C-like programs. Additionally to the CLP program specialization system developed by E. De Angelis et al. and called VeriMAP [4] they use also external satisfiability modulo theories (SMT) solvers. We would also refer to an earlier work by J. P. Gallagher et al. [15] proposing a language-independent method for analyzing the imperative programs via intermediate interpretation by a logic programming language. Note that the transformation examples given in the papers [15, 11, 4] presenting the above mentioned approaches deal with neither function nor constructor application stack in the interpreted programs.

In this paper we focus our attention on self-sufficient methods for specialization of functional programs, aiming at proving some *safety* properties of the programs. We consider a program specialization

*The second author was supported by RFBR, research project No. 17-07-00285_a, and Russian Academy of Sciences, research project No. 0077-2014-0030.

method called Turchin’s supercompilation [46, 45, 44, 21] and study potential capabilities of the method for verifying the safety properties of the functional programs modeling a class of non-deterministic parameterized cache coherence protocols [7]. We use an approach to functional modeling of non-deterministic computing systems, first presented by these authors in [29, 30, 28]. The simple idea behind the approach is as follows. Given a program modeling a deterministic computing system, whose behavior depends on and is controlled by an input parameter value, let us call for an oracle producing the input value. Then the meta-system including both the program and the external oracle becomes non-deterministic one. And vice versa, given a non-deterministic system, one may be concerned about the behavior of the system only along one possible path of the system evaluation. In such a case, the path of interest may be given as an additional input argument of the system, forcing the system to follow along the path. Dealing with an unknown value of the additional parameter one can study any possible evolution of the system, for example, aiming at verifying some properties of the system.

Viability of such an approach to verification has been demonstrated in previous works using supercompilation as a program transformation and analysis technique [29, 30, 28, 31, 22], where it was applied to safety verification of program models of parameterized protocols and Petri nets models. Furthermore, the functional program modeling and supercompilation have been used to specify and verify cryptographic protocols, and in the case of insecure protocols a supercompiler was utilized in an interactive search for the attacks on the protocols [1, 39]. In these cases the supercompiler has been used for specializing the corresponding program models, aiming at moving the safety properties of interest from the semantics level of the models to simplest syntactic properties of the residual programs produced by the supercompiler. Later this approach was extended by G. W. Hamilton for verifying a wider class of temporal properties of reactive systems [16, 17].

Given a specializer transforming the program written in a language \mathcal{L} and used for program model verification, in order to mitigate the limitation of the specification language \mathcal{L} , in this paper we study potential abilities of the corresponding specialization method for verifying the models specified in another language \mathcal{M} . We analyze the supercompilation algorithms allowing us crucially to remove the interpretation layer and to verify indirectly the safety properties. The corresponding experiments succeeded in verifying some safety properties of the series of parameterized cache coherence protocols specified, for example, in the imperative WHILE language by N. D. Jones [19]. Nevertheless, in order to demonstrate that our method is able to deal with non-imperative interpreted programs, we consider the case when a modelling language \mathcal{M} is a non-imperative subset of the basic language \mathcal{L} . On the other hand, that allows us to simplify the presentation. In order to prove the properties of interest, some of the program models used in the experiments require one additional supercompilation step (i.e., the corresponding residual programs should be supercompiled once again¹).

The considered class of cache coherence protocols effectively forms a benchmark on which various methods for parameterized verification have been tried [42, 7, 9, 13, 30, 28, 27]. In [30, 28] we have applied direct verification via supercompilation approach without intermediate interpretation. The corresponding models of these and others parameterized protocols may be very large and the automatic proofs of their safety properties may have very complicated structures. See, for example, the structure of the corresponding proof [31] produced by the supercompiler SCP4 [35, 36, 38] for the functional program model of the parameterized Two Consumers - Two Producers (2P/2C) protocol for multithreaded Java programs [2]. Taking that into account, the experiments presented in this paper can also be considered as a partial verification of the intermediate interpreters $\text{Int}_{\mathcal{M}}(p, d)$ used in the experiments. That is to say,

¹Note that the method presented in the papers [3, 5] mentioned above sometimes requires a number of iterations of the specialization step and the number is unknown.

a verification of the interpreters with respect to the subset of the input values of the argument p , being the program models of the cache coherence protocols.

The program examples given in this paper were specialized by the supercompiler SCP4 [35, 36, 38], which is a program specilalizer based on the supercompilation technique. We present our interpreter examples in a variant of a pseudocode for a functional program while real supercompilation experiments with the programs were done in the strict functional programming language Refal [48]², [49] being both the object and implementation language of the supercompiler SCP4. One of advantages of using supercompilation, instead of other forms of partial evaluation or CLP specialization, is the use of Turchin's relation (Section 4.2, see also [47, 36, 41]) defined on function-call stacks, where the function calls are labeled by the times when they are generated by the unfold-fold loop. This relation is responsible for accurate generalization of the stack structures of the unfolded program configurations. It is *based on global properties* of the path in the corresponding unfolded tree rather than on the structures of two given configurations in the path. Turchin's relation both stops the loop unfolding the tree and *provides a guidance of how a given call-stack structure has to be generalized*. Proposition 1 proven in this paper shows that a composition of the Turchin and Higman-Kruskal relations may prevent generalization of two given interpreter configurations encountered inside one big-step of the interpreter. Such a prevention from generalization is crucial for optimal specialization of any interpreter w.r.t. a given program.

This paper assumes that the reader has basic knowledge of concepts of functional programming, pattern matching, term rewriting systems, and program specialization.

The contributions of this paper are: (1) We have developed a method aiming at uniform reasoning on properties of configurations' sequences that are encountered in specializing an interpreter of a Turing complete language. (2) In particular, we have proved the following statement. Consider specialization of the interpreter with respect to any interpreted program from an infinite program set that is large enough to specify a series of parameterized cache coherence protocols, controlled by a composition of the Turchin (Section 4.2) and Higman-Kruskal (Section 4.1) relations. Given a big-step of the interpreter to be processed by the unfold-fold loop, we assume that neither generalization nor folding actions were done by this loop up to the moment considered. Then any two non-transitive (Section 4.3) big-step internal configurations C_1, C_2 are prevented from both generalization and folding actions. (3) We have shown that supercompilation controlled by the composition of the relations above is able to verify some safety properties of the series of parameterized cache coherence protocols via intermediate interpretation of their program models. Note that these program specifications include both the function call and constructor application stacks, where the size of the first one is uniformly bounded on the value of the input parameter while the second one is not. Unlike VeriMAP [4] our indirect verification method involves no post-specialization unfold-fold.

The paper is organized as follows. In Section 2 we describe the syntax and semantics of a pseudocode for a subset of the strict functional language Refal which will be used throughout this paper. We give also the operational semantics of the subset, defining its "self-interpreter". In Section 3 we outline our approach for specifying non-deterministic systems by an example used through this paper. In Section 4 we shortly introduce an unfold-fold program transformation method known as Turchin's supercompilation that is used in our experiments. We describe the strategy controlling the unfold-fold loop. The corresponding relation is a composition of Turchin's relation and a variant of the Higman-Kruskal relation. This composition plays a central role in verifying the safety properties of the cache coherence protocols' models via intermediate interpretation. In Section 5 we prove in a uniform way a number of properties

²The reader is welcome to execute several sample Refal programs and even any program written by the user directly from the electronic version of the Turchin book.

of a huge number of the complicated configurations generated by specialization of the self-interpreter with respect to the given program modeling a cache coherence protocol. The argumentations given in the section are applicable for the whole series of the protocols mentioned in Section 6. Developing the method of such argumentations is the main aim of the paper and *both Proposition 1 and the proof of this proposition are the main results of the paper*. The statement given in Proposition 1 can be applied to a wide class of interpreters of Turing complete programming languages. This statement is a theoretical basis for the explanation why the approach suggested in the paper does succeed in verifying the safety properties of the series of the cache coherence protocols via intermediate interpretation. Finally, in Section 6 we report on some other experimental results obtained by using the approach, discuss the results presented in the paper, and compare our experiments with other ones done by existing methods.

2 An Interpreter for a Fragment of the SCP4 Object Language

A first interpreter we consider is an interpreter of a subset \mathcal{L} of the SCP4 object language, which we aim to put in between the supercompiler SCP4 and programs modeling the cache coherence protocols to be verified. We will refer to this interpreter as a “self-interpreter”.

2.1 Language

$prog ::=$	$def_1 \dots def_n$	Program
$def ::=$	$f(ps_1) \Rightarrow exp_1; \dots; f(ps_n) \Rightarrow exp_n;$	Function Definition
$exp ::=$	v	Variable
	$term : exp$	<i>Cons</i> Application
	$f(exp_1, \dots, exp_n)$	Function Application
	$exp_1 ++ exp_2$	<i>Append</i> Application
	\square	<i>Nil</i>
$term ::=$	$s.name$	Symbol-Type Variable
	(exp)	Constructor Application
	σ	Symbol
$ps ::=$	p_1, \dots, p_n	Patterns
$p ::=$	v	
	$s.name : p$	
	$(p_1) : p_2$	
	$\sigma : p$	
	\square	
$v ::=$	$s.name \mid e.name$	Variable

Programs in \mathcal{L} are *strict* term rewriting systems based on pattern matching.

The rules in the programs are ordered from the top to the bottom to be matched. To be closer to Refal we use two kinds of variables: *s*.variables range over *symbols* (i.e., characters and identifiers, for example, 'a' and *True*), while *e*.variables range over the whole set of the S-expressions.³ Given a rule

³This fragment of Refal is introduced for the sake of simplicity. The reader may think that the syntactic category *exp* of list expressions and the parentheses constructor are Lisp equivalents. Actually Refal does not include *Cons* constructor but, instead of *Cons*, *Append* is used as an associative constructor. Thus the Refal data set is wider as compared with the Lisp data set: the first one is the set of finite sequences of *arbitrary* trees, while the second one is the set of binary trees. See [48] for details.

$l \Rightarrow r$, any variable of r should appear in l . Each function f has a fixed arity, i.e., the arities of all left-hand sides of the rules of f and any expression $f(exp_1, \dots, exp_n)$ must equal the arity of f . The parentheses constructor (\bullet) is used without a name. *Cons* constructor is used in infix notation and may be omitted. The patterns in a function definition are not exhaustive. If no left-hand side of the function rules matches the values assigned to a call of the function, then executing the call is interrupted and its value is undefined. In the sequel, the expression set is denoted by \mathbb{E} ; \mathbb{D} and \mathbb{S} stand for the data set, i.e., the patterns containing no variable, and the symbols set, respectively. The name set of the functions of an arity n is denoted by \mathbb{F}_n while \mathbb{F} stands for $\bigcup_{n=0}^{\infty} \mathbb{F}_n$. \mathcal{V}_e and \mathcal{V}_s stand for the e - and s -variable sets, respectively, and \mathcal{V} denotes $\mathcal{V}_e \cup \mathcal{V}_s$. For an expression exp $\mathcal{V}_e(exp)$, $\mathcal{V}_s(exp)$, $\mathcal{V}(exp)$ denote the corresponding variable sets of exp . $\mu_v(exp)$ denotes the multiplicity of $v \in \mathcal{V}$ in exp , i.e., the number of all the occurrences of v in exp . exp is called passive if no function application occurs in exp otherwise it is called an active expression. \mathcal{T} stands for the term set, σ stands for a symbol. Given an expression exp and a variable substitution θ , $exp\theta$ stands for $\theta(exp)$.

2.2 Encoding

In our experiments considered in this paper the protocol program models have to be input values of the interpreter argument with respect to which the interpreter is specialized. Thus the program models should be encoded in the data set of the implementation language of the interpreter. The program models used in this paper are written in a fragment of the language described in Section 2.1, where `++` constructor is not allowed and only unary functions may appear.

Now we have to define the corresponding encoding function denoted by the underline, where the function \mathfrak{A} groups the program rules belonging to the same function as it is shown in the second definition line.

$\underline{prog} = (\mathfrak{A}(\underline{prog}))$;	Program
$\underline{f \{rules\} defs} = (\underline{f \ rules}) : \underline{defs}$;	Function Definitions
$\underline{rule}; \underline{rules} = (\underline{rule}) : \underline{rules}$;	Rules
$\underline{f(\underline{pattern})} \Rightarrow \underline{exp} = (\underline{pattern}) : '=' : (\underline{exp})$;	Rule
$\underline{term} : \underline{exp} = \underline{term} : \underline{exp}$;	Here $term ::= (exp) \mid s.name \mid \sigma$
$\underline{(exp)} = (' * ' \ \underline{exp})$;	Applications
$\underline{e.name} = (\text{Var } 'e' \ name)$;	Variables
$\underline{\square} = \square$;	$\underline{\sigma} = \sigma$;
	Nil and Symbol

Note that any pattern is an expression.

Supercompiler SCP4 in its processing dealing with programs as input data uses this encoding function and utilizes its properties. The image of \mathbb{D} under the encoding is a proper subset of \mathbb{D} , i.e., $\underline{\mathbb{D}} \neq \mathbb{D}$. For example, $(\text{Var } 's' \ name) \notin \underline{\mathbb{D}}$.

2.3 The Interpreter

The self-interpreter used in the experiments is given in Figure 1.

The entry point is of the form $\text{Int}((\text{Call } s.f \ e.d), e.P)$. Here the first argument is the application constructor of the main function to be executed. The second argument provides the name of a program to be interpreted. The encoded source of the program will be returned by a function call of *Prog* whenever it is asked by *EvalCall*. E.g., interpretation of program model Synapse N+1 given in Section 3.1 starts with the following application $\text{Int}((\text{Call } \text{Main } \underline{d_0}), (\text{Prog } \text{Synapse}))$, where d_0 is an input data given to program Synapse. Due to the large size of the encoded programs we omit the definition of function *Prog*.

$Int((Call\ s.f\ e.d), e.P) \Rightarrow Eval(EvalCall(s.f, e.d, e.P), e.P);$
 $Eval((e.env) : (Call\ s.f\ e.q) : e.exp, e.P)$
 $\Rightarrow Eval(EvalCall(s.f, Eval((e.env) : e.q, e.P), e.P), e.P) ++ Eval((e.env) : e.exp, e.P);$
 $Eval((e.env) : (Var\ e.var) : e.exp, e.P) \Rightarrow Subst(e.env, (Var\ e.var)) ++ Eval((e.env) : e.exp, e.P);$
 $Eval((e.env) : ('*' e.q) : e.exp, e.P) \Rightarrow ('*' Eval((e.env) : e.q, e.P)) : Eval((e.env) : e.exp, e.P);$
 $Eval((e.env) : s.x : e.exp, e.P) \Rightarrow s.x : Eval((e.env) : e.exp, e.P);$
 $Eval((e.env) : [], e.P) \Rightarrow [];$
 $EvalCall(s.f, e.d, (Prog\ s.n)) \Rightarrow Matching(F, [], LookFor(s.f, Prog(s.n)), e.d);$

$Matching(F, e.old, ((e.p) : '=' : (e.exp)) : e.def, e.d)$
 $\Rightarrow Matching(Match(e.p, e.d, ([])), e.exp, e.def, e.d);$
 $Matching((e.env), e.exp, e.def, e.d) \Rightarrow (e.env) : e.exp;$
 $Match((Var\ 'e' s.n), e.d, (e.env)) \Rightarrow PutVar((Var\ 'e' s.n) : e.d, (e.env));$
 $Match((Var\ 's' s.n) : e.p, s.x : e.d, (e.env))$
 $\Rightarrow Match(e.p, e.d, PutVar((Var\ 's' s.n) : s.x, (e.env)));$
 $Match(('*' e.q) : e.p, ('*' e.x) : e.d, (e.env))$
 $\Rightarrow Match(e.p, e.d, Match(e.q, e.x, (e.env)));$
 $Match(s.x : e.p, s.x : e.d, (e.env)) \Rightarrow Match(e.p, e.d, (e.env));$
 $Match([], [], (e.env)) \Rightarrow (e.env);$
 $Match(e.p, e.d, e.fail) \Rightarrow F;$

$PutVar(e.assign, (e.env)) \Rightarrow CheckRepVar(PutV((e.assign), e.env, []));$
 $PutV(((Var\ s.t\ s.n) : e.val), ((Var\ s.t\ s.n) : e.pval) : e.env, e.penv)$
 $\Rightarrow (Eq(e.val, e.pval)) : ((Var\ s.t\ s.n) : e.pval) : e.env;$
 $PutV((e.assign), (e.passign) : e.env, e.penv)$
 $\Rightarrow PutV((e.assign), e.env, (e.passign) : e.penv);$
 $PutV((e.assign), [], e.penv) \Rightarrow (T) : (e.assign) : e.penv;$

$CheckRepVar((T) : e.env) \Rightarrow (e.env);$
 $CheckRepVar((F) : e.env) \Rightarrow F;$

$Eq(s.x : e.xs, s.x : e.ys) \Rightarrow Eq(e.xs, e.ys);$
 $Eq(('*' e.x) : e.xs, ('*' e.y) : e.ys) \Rightarrow ContEq(Eq(e.x, e.y), e.xs, e.ys);$
 $Eq([], []) \Rightarrow T;$
 $Eq(e.xs, e.ys) \Rightarrow F;$
 $ContEq(F, e.xs, e.ys) \Rightarrow F;$
 $ContEq(T, e.xs, e.ys) \Rightarrow Eq(e.xs, e.ys);$

$LookFor(s.f, (s.f : e.def) : e.P) \Rightarrow e.def;$
 $LookFor(s.f, (s.g : e.def) : e.P) \Rightarrow LookFor(s.f, e.P);$

$Subst(((Var\ s.t\ s.n) : e.val) : e.env, (Var\ s.t\ s.n)) \Rightarrow e.val;$
 $Subst((e.assign) : e.env, e.var) \Rightarrow Subst(e.env, e.var);$

Figure 1: Self-Interpreter

$EvalCall((Call\ s.f\ e.d), e.P)$ asks for the definition of function $s.f$, calling $LookFor$, and initiates matching the data given by $e.d$ against the patterns of the definition rules. In order to start this pattern matching, it imitates a fail happening in matching the data against a previous nonexistent pattern.

Function $Matching$ runs over the definition rules, testing the result of matching the input data ($e.d$) against the current pattern considered. In the case if the result is F the function calls function $Match$, asking for matching the input data against the next rule pattern. The environment $e.env$ is initialized by $[\]$. If the pattern matching succeeds then $Matching$ returns $(e.env) : e.exp$, where expression $e.exp$ is the right-hand side of the current rule and the environment includes the variable assignments computed by the pattern matching. Function $Match$ is trying to match the input data given in its second argument, step by step, against the pattern given in its first argument. It computes the environment containing the variable substitution defined by the matching. If a variable is encountered then function $PutVar$ calls $PutV$ looking for an assignment to the same variable and, if such an assignment exists, the function tests a possible coincidence of the new and old values assigned to the variable. The third rule of function $Match$ deals with the tree structure, calling this function twice.

Function $Eval$ passes through an expression given in its second argument. The second $Eval$ rule deals with a variable and calls function $Subst$ looking the environment for the variable value and replacing the variable with its value.

We intend to specialize interpreter Int with respect to its second argument. The corresponding source code of the self-interpreter may be found in <http://refal.botik.ru/protocols/Self-Int-Refal.zip>.

3 Specifying Cache Coherence Protocols

We illustrate our method [30, 28] for specifying non-deterministic systems by an example used through this paper. The Synapse N+1 protocol definition given below is borrowed from [6]. The parameterized version of the protocol is considered and *counting abstraction* is used in the specification. The protocol has to react to five external non-deterministic events by updating its states being three integer counters. The initial value of counter *invalid* is parameterized (so it could be any positive integer), while the other two counters are initialized by zero. The primed state names stand for the updated state values. The empty updates mean that nothing happened.

$$\begin{array}{ll}
 \text{(rh)} & \textit{dirty} + \textit{valid} \geq 1 \rightarrow . \\
 \text{(rm)} & \textit{invalid} \geq 1 \rightarrow \textit{dirty}' = 0, \textit{valid}' = \textit{valid} + 1, \textit{invalid}' = \textit{invalid} + \textit{dirty} - 1 . \\
 \text{(wh2)} & \textit{valid} \geq 1 \rightarrow \textit{valid}' = 0, \textit{dirty}' = 1, \textit{invalid}' = \textit{invalid} + \textit{dirty} + \textit{valid} - 1 . \\
 \text{(wm)} & \textit{invalid} \geq 1 \rightarrow \textit{valid}' = 0, \textit{dirty}' = 1, \textit{invalid}' = \textit{invalid} + \textit{dirty} + \textit{valid} - 1 .
 \end{array}
 \qquad
 \begin{array}{l}
 \text{(wh1)} \quad \textit{dirty} \geq 1 \rightarrow .
 \end{array}$$

Specification of Safety Properties Any state reached by the protocol should not satisfy any of the two following properties: (1) $\textit{invalid} \geq 0, \textit{dirty} \geq 1, \textit{valid} \geq 1$; (2) $\textit{invalid} \geq 0, \textit{dirty} \geq 2, \textit{valid} \geq 0$.

3.1 Program Model of the Synapse N+1 Cache Coherence Protocol

The program model of Synapse N+1 protocol is given in Figure 2. The idea behind the program specifications modeling the reactive systems is given in Introduction 1 above. The *finite* stream of events is modeled by a *time* value. The time ticks are labeled by the events. The counters' values are specified in the unary notation. The unary addition is directly defined by function $Append$, i.e., without referencing to the corresponding macros. Function $Loop$ exhausts the event stream, step by step, and calls for $Test$ verifying the safety property required from the protocol. Thus function $Main$ is a predicate. Note that given input values the partial predicate terminates since the event stream is finite. The termination is normal, if the final protocol state asked by the input stream is reachable one, otherwise it is abnormal.

$$\begin{aligned}
& \text{Main}((e.time) : (e.is)) \Rightarrow \text{Loop}((e.time) : (\text{Invalid } I e.is) : (\text{Dirty}) : (\text{Valid})); \\
& \text{Loop}(([]) : (\text{Invalid } e.is) : (\text{Dirty } e.ds) : (\text{Valid } e.vs)) \\
& \quad \Rightarrow \text{Test}((\text{Invalid } e.is) : (\text{Dirty } e.ds) : (\text{Valid } e.vs)); \\
& \text{Loop}((s.t : e.time) : (\text{Invalid } e.is) : (\text{Dirty } e.ds) : (\text{Valid } e.vs)) \\
& \quad \Rightarrow \text{Loop}((e.time) : \text{Event}(s.t : (\text{Invalid } e.is) : (\text{Dirty } e.ds) : (\text{Valid } e.vs))); \\
& \text{Event}(rm : (\text{Invalid } I e.is) : (\text{Dirty } e.ds) : (\text{Valid } e.vs)) \\
& \quad \Rightarrow (\text{Invalid Append}((e.ds) : (e.is))) : (\text{Dirty}) : (\text{Valid } I e.vs); \\
& \text{Event}(wh2 : (\text{Invalid } e.is) : (\text{Dirty } e.ds) : (\text{Valid } I e.vs)) \\
& \quad \Rightarrow (\text{Invalid Append}((e.vs) : (\text{Append}((e.ds) : (e.is))))) : (\text{Dirty } I) : (\text{Valid}); \\
& \text{Event}(wm : (\text{Invalid } I e.is) : (\text{Dirty } e.ds) : (\text{Valid } e.vs)) \\
& \quad \Rightarrow (\text{Invalid Append}((e.vs) : (\text{Append}((e.ds) : (e.is))))) : (\text{Dirty } I) : (\text{Valid}); \\
& \text{Append}(([]) : (e.vs)) \Rightarrow e.vs; \\
& \text{Append}((s.x : e.xs) : (e.vs)) \Rightarrow s.x : \text{Append}((e.xs) : (e.vs)); \\
& \text{Test}((\text{Invalid } e.is) : (\text{Dirty } I e.ds) : (\text{Valid } I e.vs)) \Rightarrow \text{False}; \\
& \text{Test}((\text{Invalid } e.is) : (\text{Dirty } I I e.ds) : (\text{Valid } e.vs)) \Rightarrow \text{False}; \\
& \text{Test}((\text{Invalid } e.is) : (\text{Dirty } e.ds) : (\text{Valid } e.vs)) \Rightarrow \text{True};
\end{aligned}$$

Figure 2: Model of the Synapse N+1 cache coherence protocol

4 On Supercompilation

In this paper we are interested in one particular approach in program transformation and specialization, known as supercompilation⁴. Supercompilation is a powerful semantics-based program transformation technique [44, 46] having a long history well back to the 1960-70s, when it was proposed by V. Turchin. The main idea behind a supercompiler is to observe the behavior of a functional program p running on a *partially* defined input with the aim to define a program, which would be equivalent to the original one (on the domain of the latter), but having improved properties. Given a program and its parameterized entry point, supercompilation is performed by an *unfold-fold cycle* unfolding this entry point to a potentially infinite tree of all its possible computations. It reduces the redundancy that could be present in the original program. It folds the tree into a finite graph of states and transitions between possible parameterized configurations of the computing system. And, finally, it analyses global properties of the graph and specializes this graph with respect to these properties (without additional unfolding steps).⁵ The resulting program definition is constructed solely based on the meta-interpretation of the source program rather than by a (step-by-step) transformation of the program. The result of supercompilation may be a specialized version of the original program, taking into account the properties of partially known arguments, or just a re-formulated, and sometimes more efficient, equivalent program (on the domain of the original).

Turchin's ideas have been studied by a number of authors for a long time and have, to some extent, been brought to the algorithmic and implementation stage [38]. From the very beginning the development

⁴From *supervised compilation*.

⁵See also Appendix to the extended version of this paper [34].

of supercompilation has been conducted mainly in the context of the programming language Refal [35, 36, 37, 48]. A number of model supercompilers for subsets of functional languages based on Lisp data were implemented with the aim of formalizing some aspects of the supercompilation algorithms [21, 23, 44]. The most advanced supercompiler for Refal is SCP4 [35, 36, 38].

The verification system VeriMAP [4] by E. De Angelis et al. [3, 5] uses nontrivial properties of integers recognized by both CLP built-in predicates and external SMT solvers. We use also a nontrivial property of the configurations. The property is the associativity of the built-in append function ++ supported by the supercompiler SCP4 itself⁶, rather than by an external solver.

4.1 The Well-Quasi-Ordering on \mathbb{E}

The following relation is a variant of the Higman-Kruskal relation and is a well-quasi-ordering [18, 24] (see also [26]).

Definition 1 *The homeomorphic embedding relation $\underline{\simeq}$ is the smallest transitive relation on \mathbb{E} satisfying the following properties, where $f \in \mathbb{F}_n$, $\alpha, \beta, \tau, s, t, t_1, \dots, t_n \in \mathbb{E}$ and $\alpha, \beta, \tau \in \mathcal{T}$.*

- (1) $\forall x, y \in \mathcal{V}_e. x \underline{\simeq} y, \forall u, v \in \mathcal{V}_s. u \underline{\simeq} v$; (2) $[\] \underline{\simeq} t, t \underline{\simeq} t, t \underline{\simeq} f(t_1, \dots, t, \dots, t_n), t \underline{\simeq} (t), t \underline{\simeq} \alpha : t$;
 (3-4) if $s \underline{\simeq} t$ and $\alpha \underline{\simeq} \beta$, then both $(s) \underline{\simeq} (t), \alpha : s \underline{\simeq} \beta : t$ and $f(t_1, \dots, s, \dots, t_n) \underline{\simeq} f(t_1, \dots, t, \dots, t_n)$.

Note that the definition takes into account function *append*, since its infix notation $exp_1 ++ exp_2$ stands for `append(exp1, exp2)`. We use relation $\underline{\simeq}$ modulo associativity of ++ and the following equalities $term : exp_1 = term ++ exp_1, exp ++ [\] = exp$ and $[\] ++ exp = exp$.

Given an infinite sequence of expressions t_1, \dots, t_n, \dots , relation $\underline{\simeq}$ is relevant to approximation of loops increasing the syntactical structures in the sequence; or in other words to looking for the regular similar cases of mathematical induction on the structure of the expressions. That is to say the cases, which allow us to refer one to another by a step of the induction. An additional restriction separates the basic cases of the induction from the regular ones. The restriction is: $\forall \sigma. ([\]) \not\underline{\simeq} (\sigma) \ \& \ \forall v \in \mathcal{V}_s. ([\]) \not\underline{\simeq} (v)$.

We impose this restriction on the relation $\underline{\simeq}$ modulo the equalities above and denote the obtained relation as \preceq . It is easy to see that such a restriction does not violate the quasi-ordering property. Note that the restriction may be varied in the obvious way, but for our experiments its simplest case given above is used to control generalization and has turned out to be sufficient. In the sequel, $t_1 \prec t_2$ stands for the following relation $t_1 \preceq t_2$ and $t_1 \neq t_2$, which is also transitive.

Definition 2 *A parameterized configuration is a finite sequence of the form*

let $e.h = f_1(exp_{11}, \dots, exp_{1m})$ in ... let $e.h = f_k(exp_{k1}, \dots, exp_{kj})$ in exp_{n+1} , where exp_{n+1} is passive, for all $i > 1$ $\mu_{e.h}(f_i(\dots)) = \mu_{e.h}(exp_{n+1}) = 1$, and $\mu_{e.h}(f_1(\dots)) = 0$; for all i and all j variable $e.h$ does not occur in any function application being a sub-expression of exp_{ij} . In the sequel, we refer to such a function application $f_i(\dots)$ given explicitly in the configuration as an upper function application.

The configurations represent the function application stacks, in which all constructors' applications not occurring in arguments of the upper function applications are moved to the rightmost expressions. Every expression of \mathcal{L} can be rewritten into an equivalent composition of the configurations connected with *let*-construct (see Section 5.1 for an example). Here the *append* ++ is treated as a complex constructor⁷, rather than a function. The rightmost expression is the bottom of the stack. Since the value of $e.h$ is reassigned in each *let* in the stack, for brevity sake, we use the following presentation of the

⁶As well as by the real programming language in terms of which the experiments described in this paper were done.

⁷I.e., we use nontrivial properties of configurations containing the *append* ++. See the remark given in the footnote on p. 57.

configurations:

$f_1(\text{exp}_{11}, \dots, \text{exp}_{1m}), \dots, f_k(\text{exp}_{k1}, \dots, \text{exp}_{kj}), \text{exp}_{n+l}$, where variable $e.h$ is replaced with bullet \bullet . I.e., the bullet is just a placeholder. The last expression may be omitted if it equals \bullet . An example follows: $f(a : e.xs ++ e.ys), g(\bullet ++ e.ys, (\text{Var } b\ c), []), f(s.x : \bullet), s.x : \bullet ++ t(s.x : e.zs), \bullet$.

4.2 The Well-Disordering on Timed Configurations

Let a program to be specialized and a path starting at the root of the tree unfolded by the unfold-fold loop widely used in program specialization be given. The vertices in the path are labeled by the program parameterized configurations. These configurations form a sequence. Given a configuration from such a sequence and a function application from the configuration, we label the application by the time when it is generated by the unfold-fold loop. Such a labeled function application is said to be a timed application. A configuration is said to be timed if all upper function applications in the configuration are timed. Given a timed configuration, all its timed applications have differing time-labels. Given two different configurations C_1, C_2 , if the unfold-fold loop copies an upper function application from C_1 and uses this copy in C_2 , then C_1, C_2 share this timed application. In the sequel, a sequence of the timed configurations generated by the unfold-fold loop is also called just a path. In this section we define a binary relation \triangleleft on the timed configurations in the path. The relation is originated from V.F. Turchin [47] (see also [36, 40, 41]). It is *not* transitive⁸, but like the well-quasi-ordering it satisfies the following crucial property used by supercompilation to stop the loop unfolding the tree. For any infinite path $C_1, C_2, \dots, C_n, \dots$ there exist two timed configurations C_i, C_j such that $i < j$ and $C_i \triangleleft C_j$ (see [47, 41]). For this reason we call relation \triangleleft a well-disordering relation. In the sequel, the time-labels are denoted with subscripts.

Definition 3 *Given a sequence of timed configurations C_1, \dots, C_n, \dots ; C_i and C_j are elements of the sequence such that $i < j$ and $C_i ::= f_{t_1}^1(\dots), \dots, f_{t_k}^k(\dots), \text{exp}_1, C_j ::= g_{\tau_1}^1(\dots), \dots, g_{\tau_m}^m(\dots), \text{exp}_2$, where $f_{t_s}^s$ and $g_{\tau_q}^q$, $1 \leq s \leq k$ and $1 \leq q \leq m$, stand for function names f^s, g^q labeled with times t_s and τ_q , respectively, and $\text{exp}_1, \text{exp}_2$ are passive expressions.*

If $k \leq m$, $\delta = m - k$ and $\exists l. (1 < l \leq k)$ such that $\forall s. (0 \leq s \leq k - l) f_{t_{l+s}}^{l+s} = g_{\tau_{\delta+l+s}}^{\delta+l+s}$ (i.e., $f^{l+s} = g^{\delta+l+s}$ and $t_{l+s} = \tau_{\delta+l+s}$ hold), $f_{t_{l-1}}^{l-1} \neq g_{\tau_{\delta+l-1}}^{\delta+l-1}$, and $\forall s. (0 < s < l) f_{t_s}^s \simeq g_{\tau_s}^s$ (i.e., $f^s = g^s$), then $C_i \triangleleft C_j$.

We say that configurations C_i, C_j are in Turchin's relation $C_i \triangleleft C_j$. This longest coincided suffix of the configurations are said to be the context, while the parts equal one to another modulo their time-labels are called prefixes of the corresponding configurations.

The idea behind this definition is as follows. The function applications in the context never took a part in computing the configuration C_j , in this segment of the path, while any upper function application in the prefix of C_i took a part in computing the configuration C_j . Since the prefixes of C_i, C_j coincide modulo their time-labels, these prefixes approximate a loop in the program being specialized. The prefix of C_i is the entry point in this loop, while the prefix of C_j initiates the loop iterations. The common context approximates computations after this loop. Note that Turchin's relation does not impose any restriction on the arguments of the function applications in C_i, C_j .

For example, consider the following two configurations $C_1 ::= f_4(\dots), f_3(\dots), g_2(\dots), t_1(\dots), \bullet$ and $C_2 ::= f_{10}(\dots), f_7(\dots), g_5(\dots), \dots, t_1(\dots), \bullet$, then $C_1 \triangleleft C_2$ holds. Here the context is $t_1(\dots)$, the prefix of C_1 is $f_4(\dots), f_3(\dots), g_2(\dots)$, and the prefix of C_2 is $f_{10}(\dots), f_7(\dots), g_5(\dots)$, where the subscripts of the application names stand for the time-labels. See also Appendix to the extended version of this paper [34] for a detailed example regarding Turchin's relation.

⁸See Appendix to the extended version of this paper [34] for an example demonstrating the nontransitivity of relation \triangleleft .

4.3 The Strategy Controlling the Unfolding Loop

Now we describe the main relation controlling the unfold-fold loop. That is to say, given a path starting at the root of the unfolded tree, and two timed configurations C_1, C_2 in the path such that C_1 was generated before C_2 , this relation stops the loop unfolding the tree and calls the procedures responsible for folding this path. These tools, firstly, attempt to fold C_2 by a previous configuration and, if that is impossible, then attempt to generalize this configuration pair. The relation is a composition of relations \triangleleft and \preceq . It is denoted with $\triangleleft \circ \preceq$ and is a well-disordering (see [47, 41]).

Thus we are given two timed configurations C_1, C_2 from a path, such that C_1 is generated before C_2 , and C_2 is the last configuration in the path. If relation $C_1 \triangleleft C_2$ does not hold, then the unfold-fold loop unfolds the current configuration C_2 and goes on. In the case relation $C_1 \triangleleft C_2$ holds, these configurations are of the forms (see Section 4.2 for the notation used below):

$$C_1 = f_{t_1}^1(\dots), \dots, f_{t_{l-1}}^{l-1}(\dots), f_{t_l}^l(\dots), \dots, f_{t_k}^k(\dots), exp_1,$$

$$C_2 = f_{\tau_1}^1(\dots), \dots, f_{\tau_{l-1}}^{l-1}(\dots), g_{\tau_l}^l(\dots), \dots, g_{\tau_m}^m(\dots), f_{t_l}^l(\dots), \dots, f_{t_k}^k(\dots), exp_2,$$

where the context starts at $f_{t_l}^l(\dots)$. Let C_i^p stand for the prefix of C_i , and C_i^c stand for the context of C_i followed by exp_i .

Now we compare the prefixes as follows. If there exists i ($1 \leq i < l$) such that $f_{t_i}^i(\dots) \preceq f_{\tau_i}^i(\dots)$ does not hold, then C_2 is unfolded and the unfold-fold loop goes on. Otherwise, the sub-tree rooted in C_1 is removed and the specialization task defined by the C_1 is decomposed into the two specialization tasks corresponding to C_1^p and C_1^c . Further the attempts to fold C_2^p by C_1^p and C_2^c by C_1^c do work. If some of these attempts fail, then the corresponding configurations are generalized. Note that the context may be generalized despite the fact that it does not take a part in computing the current configuration C_2 , since a narrowing of the context parameters may have happened.

A program configuration is said to be a transitive configuration if one-step unfolding of the configuration results in a tree containing only the vertices with at most one outgoing edge. For example, any function application of the form $f(d_1, \dots, d_n)$, where any $d_i \in \mathbb{D}$, is transitive. For the sake of simplicity, in the experiments described in this paper, the following strategy is used. The unfold-fold loop skips all transitive configurations encountered and removes them from the tree being unfolded. In the sequel, we refer to the strategy described in this section, including relation $\triangleleft \circ \preceq$, as the $\triangleleft \circ \preceq$ -strategy.

5 Indirect Verifying the Synapse N+1 Program Model

In this section we present an application of our program verification method based on supercompilation of intermediate interpretations. In general the method may perform a number of program specializations⁹, but all the cache coherence protocol program models that we have tried to verify by supercompiler SCP4 require at most two specializations.

Given a program partial predicate modeling both a cache coherence protocol and a safety property of the protocol, we use supercompilation aiming at moving the property hidden in the program semantics to a simple syntactic property of the residual program generated by supercompilation, i.e., this syntactic property should be easily recognized. In the experiments discussed in this paper we hope the corresponding residual programs will include no operator *return False*. Since the original direct program model terminates on any input data (see Section 3.1), this property means that the residual predicate never returns *False* and always *True*. Thus we conclude the original program model satisfies the given safety property. In the terms of functional language \mathcal{L} presented in Section 2.1 the corresponding syntactic

⁹I.e., the iterated ordinary supercompilation, which does not use any intermediate interpretation, of the residual program produced by one indirect verification.

property is “No rule’s right-hand side contains identifier *False*”.¹⁰

We can now turn to the program modeling the Synapse N+1 protocol given in Section 3.1. In order to show that the Synapse program model is safe, below we specialize the self-interpreter *Int* (see Section 2.3) with respect to the Synapse program model rather than the program model itself. Since program Synapse terminates, the self-interpreter terminates when it interprets any call of the *Main* entry function of Synapse. Since Synapse is a partial predicate, the calls of the form $Int((Call\ Main\ e.d), (Prog\ Synapse))$, where *e.d* takes any data, define the same partial predicate. Hence, the self-interpreter restricted to such calls is just another program model of protocol Synapse N+1. This indirect program model is much more complicated as compared with the direct model. We intend now to show that supercompiler SCP4 [35, 36, 38] is able to verify this model. Thus our experiments show potential capabilities of the method for verifying the safety properties of the functional programs modeling some complex non-deterministic parameterized systems. In particular, the experiments can also be considered as a partial verification of the intermediate interpreter used. In other words, verifying the interpreter with respect to a set of the interpreted programs that specify the cache coherence protocols. This specialization by supercompilation is performed by following the usual *unfold-fold cycle* controlled by the $\triangleleft \circ \preceq$ -strategy described in Section 4.3. Note that this program specification includes both the function call and constructor application stacks, where the size of the first one is uniformly bounded on the value of the input parameter while the second one is not.

We start off by unfolding the initial configuration $Int((Call\ Main\ e.d), (Prog\ Synapse))$, where the value of *e.d* is unknown. The safety property will be proved if supercompilation is able to recognize all rules of the interpreted program model, containing the *False* identifier, as being unreachable from this initial configuration.

In our early work [28] we have given a formal model of the verification procedure above by supercompilation. Let a program model and its safety property be given as described above, i.e., a partial program predicate. Given an initial parameterized configuration of the partial predicate, it has been shown that the unfold-fold cycle may be seen as a series of proof attempts by structural induction over the program configurations encountered during supercompilation aiming at verification of the safety property. Here the initial configuration specifies the statement that we have to prove. There are too many program configurations generated by the unfold-fold cycle starting with the initial configuration given above and the self-interpreter configurations are very large. As a consequence it is not possible to consider all the configurations in details. We study the configurations’ properties being relevant to the proof attempts and the method for reasoning on such properties.

5.1 On Meta-Reasoning

Let a program P_0 written in \mathcal{L} and a function application $f(d_0)$, where $d_0 \in \mathbb{D}$ is its input data, be given. Let π_0 stand for expression $(Prog\ N_{P_0})$, where N_{P_0} is the program name. The unfolding loop standing alone produces a computation path $C_0, C_1, \dots, C_n, \dots$ starting off $f(d_0)$. If $f(d_0)$ terminates then the path is finite C_0, C_1, \dots, C_k . In such a case, for any $0 \leq i < k$ C_i is a configuration not containing parameters, while C_k is either a passive expression, if partial function f is defined on the given input data, or the abnormal termination sign \perp otherwise. The unfolding iterates function $step(\cdot)$ such that

¹⁰Actually *False* is never encountered at all in any residual program generated by repeated launching the supercompiler SCP4 verifying the cache coherence protocol models considered in this paper. I.e., the property is simpler than the formulated one.

Given a safety property required from a protocol, in order to look for witnesses violating the property, the method above can be extended by deriving *False* by unfolding, using a specializer in an interactive mode. See [31, 32, 33, 39] for examples of bugged protocols and the corresponding witnesses constructed by means of the supercompiler SCP4.

$step(C_i) = C_{i+1}$.

Now let us consider the following non-parameterized configuration $K_0 = Eval([\] : \underline{f}(d_0), \pi_0)$ of the self-interpreter. If $f(d_0)$ terminates then the loop unfolding the configuration K_0 results in the encoded passive configuration produced by the loop unfolding $f(d_0)$.

$$K_1 = step(K_0) = Eval(EvalCall(\underline{f}, Eval([\] : \underline{d_0}, \pi_0), \pi_0), \pi_0) ++ Eval([\] : [\], \pi_0)$$

Expression K_1 is not a configuration. According to the strategy described in Section 2.3 the unfolding has to decompose expression K_1 in a sequence of configurations connected by the let-variables. This decomposition results in

$$\{ Eval([\] : \underline{d_0}, \pi_0), EvalCall(\underline{f}, \bullet, \pi_0), Eval(\bullet, \pi_0), \bullet \}, \\ \text{let } e.x = \bullet \text{ in } \{ Eval([\] : [\], \pi_0), e.x ++ \bullet \}, \text{ where } e.x \text{ is a fresh parameter.}$$

Hence, considering modulo the arguments, the following holds. Given a function-call stack element f , this $step$ maps the interpreted stack element to this segment of the interpreting function-call stack represented by the first configuration above, when this stack segment will be computed then its result is declared as a value of parameter $e.x$ and the last configuration will be unfolded. Note that (1) these two configurations separated with the `let`-construct will be unfolded completely separately one from the other, i.e., the first configuration becomes the input of the unfolding loop, while the second configuration is postponed for a future unfolding call; (2) built-in function `append ++` is not inserted in the stack at all, since it is treated by the supercompiler as a kind of a special constructor, which properties are known by the supercompiler handling this special constructor on the fly. The sequence $step(K_i), \dots, step(K_{j-1})$ between two consecutive applications $step(K_i), step(K_j)$ of the first $Eval$ rewriting rule unfolds the big-step of the interpreter, interpreting the regular step corresponding to the application of a rewriting rule of the f definition interpreted.

Given an expression exp to be interpreted by the interpreter, exp defines the current state of the P_0 function-call stack. Let C be the configuration representing this stack state (see Section 2.3). Let $f(exp_1)$ be the application on the top of the stack. Then the current $step(K_i)$ corresponding to the application of the first $Eval$ rewriting rule maps f to the stack segment $Eval_1, EvalCall_2, Eval_3$ of the interpreter, considering modulo their arguments, and this stack segment becomes the leading segment of the interpreting function-call stack. The remainder of the interpreted stack is encoded in the arguments of $Eval_3, Eval_4$.

This remark allows us to follow the development of these two stacks in parallel. Given the following two parameterized configurations $f(exp_1)$ and $Eval((exp_{env}) : \underline{f}(exp_1), \pi_0)$ we are going to unfold these configurations in parallel, step by step. The simpler logic of unfolding $f(exp_1)$ will provide hints on the logic of unfolding $Eval((exp_{env}) : \underline{f}(exp_1), \pi_0)$.

Now we consider the set of the configuration pairs that may be generated by the unfold-fold loop and are in the relation $\triangleleft \circ \preceq$.

5.2 Internal Properties of the Interpreter Big-Step

In this section we consider several properties of the configurations generated by the unfold-fold loop inside one big-step of the self-interpreter. In order to prove indirectly that the program model is safe, we start off by unfolding the following initial configuration $Int((Call\ Main\ e.d), (Prog\ Synapse))$, where the value of $e.d$ is unknown. Let π_s stand for $(Prog\ Synapse)$.

Consider any configuration C_b generated by the unfold-fold loop and initializing a big-step of the interpreter. Firstly, we assume that C_b is not generalized and no configuration was generalized by this loop before C_b . In such a case, C_b is of the form $Eval((env) : \underline{arg}, \pi_s), EvalCall(\underline{f}, \bullet, \pi_s), Eval(\bullet, \pi_s), \dots$ where \underline{arg} stands for the formal *syntactic* argument taken from the right-hand side of a rewriting rule where $\underline{f}(\underline{arg})$ originates from, $env ::= (var : val) : env \mid [\], var ::= \underline{s.n} \mid \underline{e.n}$, and val stands for a partially known value of variable var . Since application $\underline{f}(\underline{arg})$ is on the top of the stack, argument \underline{arg}

includes no function application. As a consequence, the leading *Eval* application has only to look for variables and to call substitution *Subst* if a variable is encountered.

Thus, excluding all the transitive configurations encountered before the substitution, we consider the following configuration: $\{Subst(env, (Var \underline{var}_{t.n})), \bullet\}$, $\text{let } e.x_1 = \bullet \text{ in } \{\{Eval(env) : \underline{arg}_1, \pi_s), \bullet\}$,

$$\text{let } e.x_2 = \bullet \text{ in } \{EvalCall(\underline{f}, e.x_1 ++ e.x_2, \pi_s), Eval(\bullet, \pi_s), \dots\}$$

where $e.x_1, e.x_2$ are fresh parameters, \underline{arg}_1 stands for a part of \underline{arg} above to be processed, and $\underline{var}_{t.n}$ denotes the type and the name of the variable encountered.

We turn now to the first configuration to be unfolded. All configurations unfolded, step by step, from the first configuration are transitive (see Section 4.3) since *Subst* tests only types and names of the environment variables. Function *Subst* is tail-recursive and returns value \underline{val} asked for.

We skip transforming these transitive configurations and continue with the next one.

$\{Eval(env) : \underline{arg}_1, \pi_s), \bullet\}$, $\text{let } e.x_2 = \bullet \text{ in } \{EvalCall(\underline{f}, \underline{val} ++ e.x_2, \pi_s), Eval(\bullet, \pi_s), \dots\}$
By our assumption above, the loop unfolding this first configuration never generates a function application. So the leading configuration proceeds to look for the variables in the same way shown above.

When \underline{arg} is entirely processed and all variables occurring in \underline{arg} are replaced with their partially known values from the environment, then the current configuration looks as follows:

$$EvalCall(\underline{f}, \underline{arg}_2, \pi_s), Eval(\bullet, \pi_s), \dots$$

Here expression \underline{arg}_2 is $\underline{arg}\theta$, where θ is the substitution defined by environment env . I. e., \underline{arg}_2 may include parameters standing for unknown data, while \underline{arg} does not. Any application of *EvalCall* function is one-step transitive. Recalling π_s , we turn to the next configuration:

$$Prog(Synapse), LookFor(\underline{f}, \bullet), Matching(F, [], \bullet, \underline{arg}_2), Eval(\bullet, \pi_s), \dots$$

$Prog(Synapse)$ returns the source code of the interpreted program *Synapse*, while the *LookFor* application returns the definition of the function called by the interpreter, using the known name \underline{f} . Skipping the corresponding transitive configurations, we have:

$$Matching(F, [], ((\underline{p}_1) : '?' : (\underline{exp}_1)) : \underline{def}_{r_1}, \underline{arg}_2), Eval(\bullet, \pi_s), \dots$$

Here the third *Matching* argument is the \underline{f} definition, where $\underline{p}_1, \underline{exp}_1, \underline{def}_{r_1}$ stand for the pattern, the right-hand side of the first rewriting rule of the definition, and the rest of this definition, respectively. This *Matching* application transitively initiates matching the parameterized data \underline{arg}_2 against pattern \underline{p}_1 and calls another *Matching* application. This second *Matching* application is provided with the \underline{f} definition rest and \underline{arg}_2 for the case this pattern matching will fail. The next configuration is as follows.

$$(\checkmark) Match(\underline{p}_1, \underline{arg}_2, ()), Matching(\bullet, \underline{exp}_1, \underline{def}_{r_1}, \underline{arg}_2), Eval(\bullet, \pi_s), \dots$$

Remark 1 *By now all the configurations generated by the unfolding loop were transitive. The steps processing syntactic structure of the function application considered might meet constructor applications. These constructor applications are accumulated in the second EvalCall argument. The analysis above did not use any particular property of the interpreted program despite the fact that the source code of the interpreted program has been received and processed.*

Now we start to deal with function *Match* playing the main role in our analysis. In order to unfold the configuration \checkmark , we have now to use some particular properties of the interpreted program.

Since for any pattern \underline{p}_1 in program *Synapse* and any $v \in \mathcal{V}$ $\mu_v(\underline{p}_1) < 2$ holds, Proposition 1 below implies that the unfold-fold loop never stops unfolding the configuration \checkmark until the *Match* application on the top of the stack will be completely unfolded to several passive expressions, step by step. These expressions will appear on different possible computation paths starting at the configuration above. Skipping the steps unfolding the tree rooted in this stack-top configuration, we turn to the configurations that appear on the leaves of this tree. Each path starting at the top configuration leads to a configuration of

one of the following two forms. These configurations are transitive:

$$\begin{aligned} & \text{Matching}((env_1), \underline{exp_1}, \underline{def_{r_1}}, \underline{arg_3}), \text{Eval}(\bullet, \pi_s), \dots \\ & \text{Matching}(F, \underline{exp_1}, \underline{def_{r_1}}, \underline{arg_3}), \text{Eval}(\bullet, \pi_s), \dots \end{aligned}$$

In the first case, the pattern matching did succeed and function *Matching* replaces the current function application with the right-hand side of the chosen rewriting rule, provided with the constructed environment. The big-step being considered has been finished. In order to launch the next big-step, interpreter *Int* has now to update the top of the interpreting function application stack.

In the second case, the pattern matching fails and function *Matching* once again calls *Match*, aiming to match the parameterized data against the pattern of the next rewriting rule of function *f*. The next configuration is of the form \surd above, in which the third *Matching* argument value is decremented with the rewriting rule has been considered. If this value is empty and cannot be decremented, then, according to the language \mathcal{L} semantics, see Section 2.1, we have the abnormal deadlock state and the interpreter work is interrupted. Starting off from this configuration, the unfold-fold loop proceeds in the way shown above.

Proposition 1 *For any pattern p_0 such that for any $v \in \mathcal{V}$ $\mu_v(p_0) < 2$ and any parameterized passive expression d , the unfold-fold loop, starting off from configuration $\text{Match}(p_0, d, ([]))$ and controlled by the $\triangleleft \circ \preceq$ -strategy, results in a tree program¹¹ such that any non-transitive vertex in the tree is labeled by a configuration of the form $\text{Match}(\underline{p}_i, d_i, (env_i))$, Given a path in the tree and any two configurations M_i, M_j of the forms $\text{Match}(\underline{p}_i, d_i, (env_i))$, ... and $\text{Match}(\underline{p}_j, d_j, (env_j))$, ... , respectively, belonging to the path, such that M_j is a descendant of M_i , then $\underline{p}_j \prec \underline{p}_i$ holds.*

Proof If all descendants of configuration $\text{Match}(p_0, d, ([]))$ are transitive then the unfold-fold loop results in a tree being a root and this tree satisfies the property required. Now consider non-transitive descendants of $\text{Match}(p_0, d, ([]))$ that may be generated by the unfold-fold loop before the first generalization or folding action happened. The patterns of the *PutV* rewriting rules never test unknown data, hence any application of *PutV* is transitive. Since for any $v \in \mathcal{V}$ relation $\mu_v(p_0) < 2$ holds, function *Eq* will be never applied and the application of *CheckRepVar* is transitive. As a consequence, all the non-transitive descendants are of the forms $\text{Match}(\underline{p}_i, d_i, (env_i))$, Only the paths originated by applications of the 2-nd, 3-rd, 4-th *Match* rewriting rules may contain configurations of such forms.

Consider a configuration $\text{Match}(\underline{p}_i, d_i, (env_i))$, Below M_i denotes such a configuration. Since p_i is a constant, one-step unfolding this configuration by the 2-nd rewriting rule leads to configuration $\text{PutVar}(\underline{s.n}: d_s, (env_i))$, $\text{Match}(\underline{p}_{i+1}, d_{i+1}, \bullet)$, ... such that \underline{p}_{i+1} is a proper part of \underline{p}_i , in which at least one constructor is removed. Hence, $\underline{p}_{i+1} \prec \underline{p}_i$ holds. Since the *PutVar* application is transitive, a number of unfolding steps lead transitively to $\text{Match}(\underline{p}_{i+1}, d_{i+1}, (env_{i+1}))$, ... such that $\underline{p}_{i+1} \prec \underline{p}_i$.

One-step unfolding the configuration $\text{Match}(\underline{p}_i, d_i, (env_i))$, ... by the 3-rd rewriting rule leads to configuration $\text{Match}(\underline{p}_{i+1}, d_{i+1}, (env_i))$, $\text{Match}(\underline{p}_{i+2}, d_{i+2}, \bullet)$, ... such that \underline{p}_{i+1} and \underline{p}_{i+2} are proper parts of constant \underline{p}_i . Hence, $\underline{p}_{i+1} \prec \underline{p}_i$ and $\underline{p}_{i+2} \prec \underline{p}_i$ hold.

One-step unfolding the configuration $\text{Match}(\underline{p}_i, d_i, (env_i))$, ... by the 4-th rewriting rule leads to configuration $\text{Match}(\underline{p}_{i+1}, d_{i+1}, (env_i))$, ... such that \underline{p}_{i+1} is a proper part of \underline{p}_i , in which at least one constructor is removed. Hence, $\underline{p}_{i+1} \prec \underline{p}_i$ holds.

Now consider any two configurations of the forms $\text{Match}(\underline{p}_i, d_i, (env_i))$, ... and $\text{Match}(\underline{p}_j, d_j, (env_j))$, ... such that the second configuration belongs to a path originating from the first one and is encountered before any generalization. Hence, \underline{p}_i and \underline{p}_j are constants.

Given a configuration C , the length of C , denoted by $ln_c(C)$, is the number of the *upper* function applications in C . (See Definition 2 above.)

¹¹I.e., without any function application, except an entry point of this residual program.

If $ln_c(M_j) < ln_c(M_i)$ then $M_i \not\triangleleft M_j$ and the Turchin relation prevents M_j, M_i from generalization and M_j from any folding action. (See Section 4.3.)

If $ln_c(M_j) \geq ln_c(M_i)$ and $M_i \triangleleft M_j$ then we consider the first shortest configuration M_k in the path segment being considered. By definition of the stack, $Match(\underline{p}_k, d_k, \bullet)$ is from M_i and $Match(\underline{p}_j, d_j, (env_j))$ is a descendent of $Match(\underline{p}_k, d_k, (env_k))$.

Since $Match(\underline{p}_k, d_k, \bullet)$ took a part in computing M_j , it is the last function application of the stack prefix defined by the following Turchin relation $M_i \triangleleft M_j$, which holds. On the other hand, by the reasoning given above, for any p_{j_i} from the prefix of M_j , defined by *this* Turchin relation, $p_{j_i} \prec \underline{p}_k$ holds, and, as a consequence, the following relation $Match(\underline{p}_{j_i}, d_{j_i}, (env_{j_i})) \prec Match(\underline{p}_k, d_k, \bullet)$ holds. According to the $\triangleleft \circ \prec$ -strategy this relation prevents M_j, M_i from generalization and M_j from any folding action. (See Section 4.3.) The proposition has been proven. \square

5.3 Dealing with the Interpreter Function-Application Stack

Firstly, assume that no generalization happened in the unfold-fold loop up to now. Given a right-hand side exp_0 of a rewriting rule returned by function *Matching* as described in Section 5.2, then the following configuration has to map, step by step, the segment of the interpreted function-application stack, that is defined by known exp_0 , on the top of the interpreting stack: $Eval((env) : \underline{exp}_0, \pi_s), \dots$

According to the call-by-value semantics, function *Eval* looks for the function application, whose right bracket is the leftmost closing bracket, in completely known exp_0 . It moves from left to right along exp_0 , substitutes transitively the values of the variables encountered, as shown in Section 5.2, pushes the interpreted function application in the interpreting stack, mapping it into an *EvalCall* application, whenever the interpreted application should be pushed in the interpreted stack. See Section 5.1 for the details. Finally the depth first *EvalCall* application initiates the next big-step.

Since for any pattern p of program Synapse and any $v \in \mathcal{V}$ $\mu_v(p) < 2$ holds, all applications of *Eq*, *ContEq* and *PutV* are transitive. This note together with Proposition 1 implies Proposition 2 below.

Let p_i stand for sub-patterns of a pattern of program Synapse, arg_i and def stand for partially known parameterized expressions and several, maybe zero, rewriting rules being a rest of a function definition of Synapse, respectively. Let exp stand for the right-hand side of a rewriting rule from this definition. $env ::= (var : val) : env \mid [], var ::= \underline{s.n} \mid \underline{e.n}$, and val stands for a partially known value of variable var . Let Int_0 denote $Int((Call\ Main\ e.d), \pi_s)$, where the value of $e.d$ is unknown.

Proposition 2 *Let the unfold-fold loop be controlled by the $\triangleleft \circ \prec$ -strategy. Let it start off from the initial configuration Int_0 . Then the first generalized configuration generated by this loop, if any, will generalize two configurations of the following form and any configuration folded, before this generalization, by a previous configuration is of the same form, where $n > 0$,*

(\heartsuit) $Match(\underline{p}_1, arg_1, (env)), \dots, Match(\underline{p}_n, arg_n, \bullet), Matching(\bullet, \underline{exp}, \underline{def}, arg_0), Eval(\bullet, \pi_s), \dots$

Now consider any application of the form $Match(\underline{p}_1, arg_1, (env))$ staying on the stack top. Let $Match_1$ denote this application. Only the third rewriting rule of *Match* increases the stack. In this case the next state after the stack \heartsuit is of the form $Match_3, Match_2 \dots$. Then, by Proposition 1, along any path originating from $Match_2$ the application $Match_2$ is not replaced until application $Match_3$ will be completely unfolded. Hence, for any stack state of the form $f_i, \dots, Match_2, \dots$ on such a path, where f_i denotes any function application, the following relation $Match_2, \dots \triangleleft f_i, \dots, Match_2, \dots$ holds. That proves the following corollary using the notation given above.

Corollary 1 *Given a timed application $Matching_{t_0}(\bullet, \underline{exp}, \underline{def}, arg_0)$, any two timed configurations of the form $\dots, Matching_{t_0}(\bullet, \underline{exp}, \underline{def}, arg_0), Eval(\bullet, \pi_s), \dots$ can neither be generalized nor folded one by the other.*

Since any application $Matching(\dots, \underline{exp}, \underline{def}, \dots)$ decreases, step by step, the list \underline{def} of the rewriting rules, the following corollary holds.

Corollary 2 *Given a big-step of Int, a timed pair $\underline{exp}_{t_i}, \underline{def}_{t_i}$ ¹², and a timed application $Matching_{\tau_0}(\bullet, \underline{exp}_{t_i}, \underline{def}_{t_i}, arg_0)$ inside this big-step, then any two timed configurations of the form $\dots, Matching_{\tau_i}(\bullet, \underline{exp}_{t_i}, \underline{def}_{t_i}, arg_0), Eval(\bullet, \pi_s), \dots$ can neither be generalized nor folded one by the other.*

Given a function definition F and a rewriting rule r of the definition, let $exp_{F,r}$ stand for the right-hand side of r , while $def_{F,r}$ stand for the rest of this definition rules following the rule r . The following is a simple syntactic property of the Synapse program model given in Section 3.

(5.1) For any $F_1, F_2 \in \{Main, Loop, Event, Append, Test\}$ and for any two *distinct* rewriting rules r_1, r_2 of F_1, F_2 , respectively, $(exp_{F_1, r_1}, def_{F_1, r_1}) \neq (exp_{F_2, r_2}, def_{F_2, r_2})$ holds.

That together with Corollaries 1, 2 imply:

Proposition 3 *Given two configurations C_1 and C_2 of the form \blacktriangledown to be generalized or folded one by the other. Then (1) C_1 and C_2 cannot belong to the same big-step of Int; (2) there are two functions F_1, F_2 of a cache coherence protocol model from the series mentioned in Section 6 (and specified in the way shown in Section 3) and rewriting rules r_1, r_2 of F_1, F_2 , respectively, such that $(exp_{F_1, r_1}, def_{F_1, r_1}) = (exp_{F_2, r_2}, def_{F_2, r_2})$, where functions F_1, F_2 and rules r_1, r_2 may coincide, respectively.*

Remark 2 *Proposition 3 depends on Property 5.1. Nevertheless, the restriction imposed by 5.1 is very weak and the most of programs written in \mathcal{L} satisfy 5.1. It can easily be overcome by providing the interpreting function $Matching$ with an additional argument that is the interpreted function name. The second statement of this proposition is crucial for the expectation of removing the interpretation overheads. Despite the fact that the reasoning above follows the Synapse program model, it can be applied to any protocol model used in our experiments described in this paper.*

We conclude that any configuration encountered by the loop unfolding the given big-step is neither generalized with another configuration generated in unfolding this big-step nor folded by such a configuration.

5.4 On Generalizing the Interpreter Configurations

The unfold-fold loop processes the paths originating from the initial configuration Int_0 , following the corresponding interpreted paths. The latter ones are processed according to the order of the rewriting rules in the Synapse function definitions. The configuration Int_0 is unfolded according to the *Main* function definition of the interpreted program. Application of this function leads transitively to the following function application: $Loop((time):(Invalid I is):(Dirty):(Valid))$. The interpreter has to match this call against the left-hand side of the first rewriting rule of the *Loop* definition. The first corresponding pattern is: $([]):(Invalid e.is):(Dirty e.ds):(Valid e.vs)$.

The pattern matching processes the pattern and argument from the left to the right, by means of function *Match*. The known part of the tree structure is mapped into the interpreting stack by the third rule of *Match*. The number of *Match* applications in the stack is increased by this rewriting rule.

In the given context of specialization the values of *time* and *is* are unknown. Hence, the prefix of this stack that is responsible for matching the argument constant structure on the left-hand side of *time* will

¹²Here we use the notation given above and the timed expressions, which are defined in the same way as the timed applications (Section 4.2).

transitively disappear and the unfolding loop will stop at the configuration of the following form

$$\text{Match}(\ [], \text{time}, (\ [])), \text{Match}(p_2, \text{arg}_2, \bullet), \text{Matching}(\bullet, \underline{\text{exp}}, \underline{\text{def}}, \text{arg}_0), \text{Eval}(\bullet, \pi_s), \bullet.$$

Here the leading *Match* application meets the unknown data *time* and has to match it against \square given in the first argument. The environment in the third argument is empty since no variable was still assigned up to now. The second *Match* application is responsible for matching the suspended part of the input data arg_2 against the rest p_2 of the pattern. I. e., arg_2 equals $(\text{Invalid } I \text{ is}):(\text{Dirty }):(\text{Valid })$ and p_2 equals $(\text{Invalid } e \text{ is}):(\text{Dirty } e \text{ ds}):(\text{Valid } e \text{ vs})$.

Now we note that the arguments of all applications of the recursive *Loop* and *Append* functions have exactly the same constant prefix as considered above. That leads to the following proposition.

Proposition 4 *Let the unfold-fold loop be controlled by the $\triangleleft \circ \preceq$ -strategy. Let it start off from the initial configuration Int_0 . Then the first generalized configuration generated by this loop, if any, will generalize two configurations of the following forms and any configuration folded, before this generalization, by a previous configuration is of the same form*

$$\text{Match}(\ [], \text{arg}_1, (\ [])), \text{Match}(p_2, \text{arg}_2, \bullet), \text{Matching}(\bullet, \underline{\text{exp}}, \underline{\text{def}}, \text{arg}_0), \text{Eval}(\bullet, \pi_s), \dots$$

In the given context of specialization Turchin's relation plays a crucial role in preventing the encountered configurations from generalization (see Proposition 1). It never forces decomposing the generalized configurations as might do, in general (see Section 4.2). Both the crucial configurations of the unfolding history leading to verification of this program model and some properties of the corresponding generalized configurations may be found in the extended version of this paper published as a preprint [34].

6 Conclusion

We have shown that a combination of the verification via supercompilation method and the first Futamura projection allows us to perform verification of the program being interpreted. We discussed the crucial steps of the supercompilation process involved in the verification of a parameterized cache coherence protocol used as a case study. In the same way we were able to verify all cache coherence protocols from [7, 8], including MSI, MOSI, MESI, MOESI, Illinois University, Berkley RISC, DEC Firefly, IEEE Futurebus+, Xerox PARC Dragon, specified in the interpreted language \mathcal{L} . Furthermore, we were able to verify the same protocols specified in the language WHILE [19]. The complexity of involved processes is huge and further research is required for their better understanding.

Our experimental results show that Turchin's supercompilation is able to verify rather complicated program models of non-deterministic parameterized computing systems. The corresponding models used in our experiments are constructed on the base of the well known series of the cache coherence protocols mentioned above. So *they might be new challenges to be verified by program transformation rather than an approach for verifying the protocols themselves*. This protocol series was early verified by Delzanno [7] and Esparza et al. [10] in abstract terms of equality and inequality constraints. Using unfold-fold program transformation tools this protocol series was early verified by the supercompiler SCP4 [29, 30, 28, 31] in terms of a functional programming language, several of these protocols were verified in terms of logic programming [43, 12]. One may consider the indirect protocol models presented in this paper as a new collection of tests developing the state-of-the-art unfold-fold program transformation.

The intermediate interpreter considered in this paper specifies the operational semantics of a *Turing-complete* language \mathcal{L} . We have proved several statements on properties of the configurations generated by the unfold-fold cycle in the process specializing *Int* with respect to the cache coherence protocols specified as shown in Section 3. The main of them is Proposition 1. Some of these properties *do not depend on specific protocols* from the considered series, i. e., they hold for any protocol specified in

the way shown in Section 3. That allows us to reason, *in a uniform way*, about a huge number of complicated configurations. Note that the programs specifying the protocols include both the function call and constructor application stacks, where the size of the first one is uniformly bounded on the value of the input parameter while the second one is not.

As a future work, we would like to address the issue of the description of suitable properties of interpreters to which our uniform reasonings demonstrated in this paper might be applied.

Acknowledgements: We would like to thank Antonina Nepeivoda and anonymous reviewers for helping to improve this work.

References

- [1] A. Ahmed, A.P. Lisitsa & A.P. Nemytykh (2013): *Cryptographic Protocol Verification via Supercompilation (A Case Study)*. In: *VPT 2013, EPiC Series 16*, EasyChair, pp. 16–29. Available at URL <http://www.easychair.org/publications/paper/147590>.
- [2] L. van Begin: *The BABYLON Project: A Tool for Specification and Verification of Parameterized Systems to Benchmark Infinite-State Model Checkers*. [online]. <http://www.ulb.ac.be/di/ssd/lvbegin/CST/>.
- [3] E. De Angelis, F. Fioravanti, A. Pettorossi & M. Proietti (2014): *Program verification via iterated specialization*. *Science of Computer Programming 95 / Selected and extended papers from Partial Evaluation and Program Manipulation 2013(Part 2)*, pp. 149–175, doi:10.1016/j.scico.2014.05.017.
- [4] E. De Angelis, F. Fioravanti, A. Pettorossi & M. Proietti (2014): *VeriMAP: A Tool for Verifying Programs through Transformations*. In: *Proc. of (TACAS 2014), LNCS 8413*, Springer, pp. 568–574, doi:10.1007/978-3-642-54862-8_47.
- [5] E. De Angelis, F. Fioravanti, A. Pettorossi & M. Proietti (2015): *Proving correctness of imperative programs by linearizing constrained Horn clauses*. *Theory and Practice of Logic Programming 15 / 31st International Conference on Logic Programming(4–5)*, pp. 635–650, doi:10.1017/S1471068415000289.
- [6] G. Delzanno: *Automatic Verification of Cache Coherence Protocols via Infinite-state Constraint-based Model Checking*. [online]. <http://www.disi.unige.it/person/DelzannoG/protocol.html>.
- [7] G. Delzanno (2000): *Automatic Verification of Parameterized Cache Coherence Protocols*. In: *The Proc. of the 12th Int. Conference on Computer Aided Verification, LNCS 1855*, pp. 53–68, doi:10.1007/10722167_8.
- [8] G. Delzanno (2000): *Verification of Consistency Protocols via Infinite-state Symbolic Model Checking, A Case Study: the IEEE Futurebus+ Protocol*. In: *The Proc. of FORTE/PSTV*, Springer US, pp. 171–186, doi:10.1007/978-0-387-35533-7_11.
- [9] E. A. Emerson & V. Kahlon (2003): *Exact and Efficient Verification of Parameterized Cache Coherence Protocols*. In: *the Proc. of CHARME 2003: 12th IFIP WG 10.5 Advanced Research Working Conference*, Springer, pp. 247–262, doi:10.1007/978-3-540-39724-3_22.
- [10] J. Esparza, A. Finkel & R. Mayr (1999): *On the verification of broadcast protocols*. In: *the Proc. of 14th Symposium Logic in Computer Science*, IEEE, pp. 352–359, doi:10.1109/LICS.1999.782630.
- [11] F. Fioravanti, A. Pettorossi & M. Proietti (2001): *Verifying CTL properties of infinite state systems by specializing constraint logic programs*. In: *the Proc. of VCL01, Tech. Rep. DSSE-TR-2001-3*, University of Southampton, UK, pp. 85–96. Available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.15.9752>.
- [12] F. Fioravanti, A. Pettorossi, M. Proietti & V. Senni (2013): *Generalization Strategies for the Verification of Infinite State Systems*. In W. Faber & N. Leone, editors: *Theory and Practice of Logic Programming*, 13 / Special Issue 02 (25th GULP annual conference), Cambridge University Press, pp. 175–199, doi:10.1017/S1471068411000627.

- [13] M. Fisher, B. Konev & A. Lisitsa (2006): *Practical Infinite-State Verification with Temporal Reasoning*. In: *Proc. of the VISSAS 2005*, IOS Press, pp. 91–100.
- [14] Y. Futamura (1971): *Partial evaluation of computing process an approach to a compiler-compiler*. *Systems, Computers, Controls* 2(5), pp. 45–50, doi:10.1023/A:1010095604496.
- [15] J. P. Gallagher, J. C. Peralta & H. Saglam (1998): *Analysis of imperative programs through analysis of Constraint Logic Programs*. In B. Demoen & V. Lifschitz, editors: *SAS 1998*, LNCS 1503, Springer, pp. 246–261, doi:10.1007/3-540-49727-7_15.
- [16] G. W. Hamilton (2015): *Verifying Temporal Properties of Reactive Systems by Transformation*. *Electronic Proceedings in Theoretical Computer Science* 199, pp. 33–49, doi:10.4204/EPTCS.199.3.
- [17] G. W. Hamilton (2016): *Generating Counterexamples for Model Checking by Transformation*. *Electronic Proceedings in Theoretical Computer Science* 216, pp. 65–82, doi:10.4204/EPTCS.216.4.
- [18] G. Higman (1952): *Ordering by divisibility in abstract algebras*. *Proc. London Math. Soc.* 2(7), pp. 326–336, doi:10.1112/plms/s3-2.1.326.
- [19] N. D. Jones (1997): *Computability and Complexity from a Programming Perspective*. The MIT Press. ISBN:9780262100649, Revised version (2007) is available at <http://www.diku.dk/~neil/comp2book2007/book-whole.pdf>.
- [20] N. D. Jones (2004): *Transformation by Interpreter Specialisation*. *Science of Computer Programming* 52, pp. 307–339, doi:10.1016/j.scico.2004.03.010.
- [21] P. A. Jonsson & J. Nordlander (2009): *Positive Supercompilation for a higher order call-by-value language*. *ACM SIGPLAN Notices* 44(1), pp. 277–288, doi:10.1145/1480881.1480916.
- [22] A. V. Klimov (2012): *Solving Coverability Problem for Monotonic Counter Systems by Supercompilation*. In: *the Proc. of PSI'11*, LNCS 7162, pp. 193–209, doi:10.1007/978-3-642-29709-0_18.
- [23] I. Klyuchnikov (2010): *Supercompiler HOSC 1.5: homeomorphic embedding and generalization in a higher-order setting*. Technical Report 62, Keldysh Institute of Applied Mathematics, Moscow.
- [24] J. B. Kruskal (1960): *Well-quasi-ordering, the tree theorem, and Vazsonyi's conjecture*. *Trans. Amer. Math. Society* 95, pp. 210–225, doi:10.2307/1993287.
- [25] H. Lehmann & M. Leuschel (2004): *Inductive Theorem Proving by Program Specialisation: Generating Proofs for Isabelle Using Ecce*. In: *Proceedings of LOPSTR03*, LNCS 3018, pp. 1–19, doi:10.1007/978-3-540-25938-1_1.
- [26] M. Leuschel (1998): *On the Power of Homeomorphic Embedding for Online Termination*. In: *Proc. of the SAS'98*, LNCS 1503, pp. 230–245, doi:10.1007/3-540-49727-7_14.
- [27] A. Lisitsa (2010): *Reachability as Derivability, Finite Countermodels and Verification*. In: *ATVA 2010*, LNCS 6252, pp. 233–244, doi:10.1007/978-3-642-15643-4_18.
- [28] A. Lisitsa & A. P. Nemytykh (2008): *Reachability Analysis in Verification via Supercompilation*. *International Journal of Foundations of Computer Science* 19(4), pp. 953–970, doi:10.1142/S0129054108006066.
- [29] A. P. Lisitsa & A. P. Nemytykh (2007): *A Note on Specialization of Interpreters*. In: *The 2-nd International Symposium on Computer Science in Russia (CSR-2007)*, LNCS 4649, pp. 237–248, doi:10.1007/978-3-540-74510-5_25.
- [30] A. P. Lisitsa & A. P. Nemytykh (2007): *Verification as Parameterized Testing (Experiments with the SCP4 Supercompiler)*. *Programmirovanie, (In Russian)* 1, pp. 22–34, doi:10.1134/S0361768807010033. English translation in *J. Programming and Computer Software*, Vol. 33, No.1, pp: 14–23, 2007,.
- [31] A. P. Lisitsa & A. P. Nemytykh (2007–2009): *Experiments on Verification via Supercompilation*. [online]. <http://refal.botik.ru/protocols/>.
- [32] A. P. Lisitsa & A. P. Nemytykh (2008): *Extracting Bugs from the Failed Proofs in Verification via Supercompilation*. In Bernhard Beckert & Reiner Hahnle, editors: *Tests and Proofs: Papers Presented at the Second International Conference TAP 2008*, Reports of the Faculty of Informatics, Univesitat Koblenz-Landau

- 5/2008, pp. 49–65. Available at <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.160.7790&rep=rep1&type=pdf#page=57>.
- [33] A. P. Lisitsa & A. P. Nemytykh (2014): *A Note on Program Specialization. What Can Syntactical Properties of Residual Programs Reveal?* In: *VPT 2014, EPiC Series 28*, EasyChair, pp. 52–65. Available at URL <http://www.easychair.org/publications/paper/183895>.
- [34] A. P. Lisitsa & A. P. Nemytykh (2017): *Verifying Programs via Intermediate Interpretation (Extended version)*. arXiv:1705.06738.
- [35] A. P. Nemytykh (2003): *The Supercompiler SCP4: General Structure. (Extended abstract)*. In: *the Proc. of PSI'03, LNCS 2890*, pp. 162–170, doi:10.1007/978-3-540-39866-0_18.
- [36] A. P. Nemytykh (2007): *The Supercompiler SCP4: General Structure*. URSS, Moscow. (Book in Russian: <http://urss.ru/cgi-bin/db.pl?cp=&page=Book&id=55806&lang=Ru&blang=ru&list=>).
- [37] A. P. Nemytykh, V. A. Pinchuk & V. F. Turchin (1996): *A Self-Applicable Supercompiler*. In: *PEPM'96, LNCS 1110*, Springer-Verlag, pp. 322–337, doi:10.1007/3-540-61580-6_16.
- [38] A. P. Nemytykh & V. F. Turchin (2000): *The Supercompiler SCP4: Sources, On-line Demonstration*. [online]. <http://www.botik.ru/pub/local/scp/refal5/>.
- [39] Antonina Nepeivoda (2013): *Ping-Pong Protocols as Prefix Grammars and Turchin Relation*. In: *VPT 2013, EPiC Series 16*, EasyChair, pp. 74–87. Available at URL <http://www.easychair.org/publications/paper/147593>.
- [40] Antonina Nepeivoda (2014): *Turchin's Relation and Loop Approximation in Program Analysis. Proceedings on the Functional Language Refal (in Russian) 1*, pp. 170–192. ISBN:978-5-9905410-1-6, available at http://refal.botik.ru/library/refal2014_issue-I.pdf.
- [41] Antonina Nepeivoda (2016): *Turchin's Relation for Call-by-Name Computations: A Formal Approach. Electronic Proceedings in Theoretical Computer Science 216*, pp. 137–159, doi:10.4204/EPTCS.216.8.
- [42] F. Pong & M. Dubois (1997): *Verification Techniques for Cache Coherence Protocols*. *ACM Comput. Surv.* 29(1), pp. 82–126, doi:10.1145/248621.248624.
- [43] A. Roychoudhury & I. V. Ramakrishnan (2004): *Inductively verifying invariant properties of parameterized systems*. *Automated Software Engineering* 11(2), pp. 101–139, doi:10.1023/B:AUSE.0000017740.35552.88.
- [44] M. H. Sørensen, R. Glück & N. D. Jones (1996): *A Positive Supercompiler*. *Journal of Functional Programming* 6(6), pp. 811–838, doi:10.1017/S0956796800002008.
- [45] V. F. Turchin (1980): *The language Refal – The Theory of Compilation and Metasystem Analysis*. Technical Report 20, Courant Institute of Mathematical Sciences, New York University. Available at <https://pdfs.semanticscholar.org/75cb/a161e01f2920c8d4668fa5c979adc7422461.pdf>.
- [46] V. F. Turchin (1986): *The Concept of a Supercompiler*. *ACM Transactions on Programming Languages and Systems* 8(3), pp. 292–325, doi:10.1145/5956.5957.
- [47] V. F. Turchin (1988): *The Algorithm of Generalization in the Supercompiler*. In: *Proc. of the IFIP TC2 Workshop, Partial Evaluation and Mixed Computation*, North-Holland Publishing Co., pp. 531–549.
- [48] V. F. Turchin (1989): *Refal-5, Programming Guide and Reference Manual*. New England Publishing Co., Holyoke, Massachusetts. Electronic version: <http://www.botik.ru/pub/local/scp/refal5/>, 2000.
- [49] V. F. Turchin, D. V. Turchin, A. P. Konyshov & A. P. Nemytykh (2000): *Refal-5: Sources, Executable Modules*. [online]. <http://www.botik.ru/pub/local/scp/refal5/>.