

Basic Research in Computer Science

BRICS RS-00-39 Møller & Schwartzbach: The Pointer Assertion Logic Engine

The Pointer Assertion Logic Engine

Anders Møller
Michael I. Schwartzbach

BRICS Report Series

RS-00-39

ISSN 0909-0878

December 2000

**Copyright © 2000, Anders Møller & Michael I. Schwartzbach.
BRICS, Department of Computer Science
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work
is permitted for educational or research use
on condition that this copyright notice is
included in any copy.**

**See back inner page for a list of recent BRICS Report Series publications.
Copies may be obtained by contacting:**

**BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK-8000 Aarhus C
Denmark
Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk**

**BRICS publications are in general accessible through the World Wide
Web and anonymous FTP through these URLs:**

`http://www.brics.dk`
`ftp://ftp.brics.dk`
This document in subdirectory RS/00/39/

The Pointer Assertion Logic Engine

Anders Møller Michael I. Schwartzbach

BRICS*

Department of Computer Science
University of Aarhus, Denmark
{amoeller,mis}@brics.dk

Abstract

We present a new framework for verifying partial specifications of programs in order to catch type and memory errors and check data structure invariants. Our technique can verify a large class of data structures, namely all those that can be expressed as *graph types*. Earlier versions were restricted to simple special cases such as lists or trees. Even so, our current implementation is as fast as the previous specialized tools.

Programs are annotated with partial specifications expressed in Pointer Assertion Logic, a new notation for expressing properties of the program store. We work in the logical tradition by encoding the programs and partial specifications as formulas in monadic second-order logic. Validity of these formulas is checked by the MONA tool, which also can provide explicit counterexamples to invalid formulas.

Other work with similar goals is based on more traditional program analyses, such as shape analysis. That approach requires explicit introduction of an appropriate operational semantics along with a proof of correctness whenever a new data structure is being considered. In comparison, our approach only requires the data structure to be abstractly described in Pointer Assertion Logic.

1 Introduction

We present a new contribution to the area of *pointer verification*, which is concerned with verifying partial specifications of programs that make explicit use of pointers. In practice, there is an emphasis on catching type and memory errors and checking data structure invariants.

*Basic Research in Computer Science,
Centre of the Danish National Research Foundation.

For data type implementations, standard type checking systems, as in C or Java, are not sufficiently expressive. For example, the type of binary trees is identical to the one for doubly-linked lists. Both are just records with pairs of pointers, which makes the type checker fail to catch many common bugs. In contrast, pointer verification can validate the underlying data structure invariants, for instance, to guarantee that doubly-linked lists maintain their shapes after pointer manipulations. Memory errors, such as dereference of `null` pointers or dangling references, and creation of memory leaks are also beyond the scope of standard type checking.

There have been several different approaches to pointer verification, but not many that are as expressive as the one we propose in this paper. Clearly there is a trade-off between expressiveness and complexity, since less detailed analyses will be able to handle larger programs. Our approach is designed to verify a single abstract data type at a time. Since such implementations often contain intricate pointer manipulations and are trusted implicitly by programmers, they are a fair target for detailed scrutiny.

We work in the logical tradition by encoding programs and partial specifications as formulas in monadic second-order logic. Formulas are processed by the MONA tool [27, 34] which reduces them to equivalent tree automata from which it is simple to conclude validity or to extract concrete counterexamples. Translated back into the underlying programming language, a counterexample is an initial store that causes the given program to fail. Program annotations, in the form of assertions and invariants, are allowed and may prove necessary to obtain the desired degree of precision. This approach can be viewed both as lightweight program verification, since the full behavior of the program is not considered, and as heavyweight type checking, since properties well beyond the expressiveness of standard type systems can be checked.

We have reported on our approach in two earlier works. In the first we introduce the basic technique applied to linear lists [25]. In the second we provide a generalization to tree-shaped data structures and introduce a new encoding to make the analysis feasible [14]. The current paper takes a leap forward in generalizing the class of data structures that can be considered, without sacrificing precision or efficiency. Our new framework can handle all data structures that can be described as *graph types* [29]. These include data structures that are well-known from folklore or literature, such as doubly-linked lists, trees with parent pointers, threaded trees, two-dimensional range trees, and endless customized versions such as trees in which all leaves are linked in a cyclic list. Our framework is also designed to handle the common situation where a data structure invariant must be temporarily violated at some program points.

Our contributions are:

- An extension of the results in [25, 14] to the whole class of graph types;
- a language for expressing data structures and operations along with correctness specifications;

- an full implementation exploiting intricate parts of the MONA tool to obtain an efficient decision procedure, together with a range of non-trivial examples.

To verify a data type implementation, the desired data structure is specified in an abstract notation, and the program is annotated with assumptions and assertions. It is not necessary to customize or optimize the implementation, and no proof obligations are left to be dealt with manually.

We rely on a new formal notation, *Pointer Assertion Logic (PAL)* to specify the structural invariants of graph types, to state pre- and post-conditions for procedures, and to formulate invariants and assertions that are given as hints to the system. The PAL notation is essentially a monadic second-order logic in which the universe of discourse contains records, pointers, and booleans. Programs with PAL annotations are verified with the tool PALE, the *Pointer Assertion Logic Engine*. The “secret” behind the PALE implementation is using the MONA tool to decide validity of Hoare triples based on PAL over loop-free code. Code with loops or recursion is handled by splitting it into loop-free fragments using invariants, as in classical Hoare logic. While the MONA logic has an inherent non-elementary complexity [33], we demonstrate that it can efficiently handle many real programs. Furthermore, the ability to insert assertions to break larger triples into smaller ones suggests that the overall approach is modular and thus can scale reasonably.

A framework for pointer verification, such as ours, should be evaluated on four different criteria. First, how precise is the analysis? Second, it is fast and scalable? Third, does it allow or require programs to be annotated? Fourth, which data structures can be considered and how are they described? In the following sections, we will describe a programming language that uses Pointer Assertion Logic for expression of store properties, describe the decision procedure based on Hoare logic and MONA, and through a number of experiments argue that the Pointer Assertion Logic approach provides a productive compromise between expressibility and efficiency.

A Tiny Example

Consider the type of linked lists with tail pointers, which as a graph type is expressed as:

```

type Head = {
  data first: Node;
  pointer
  last: Node[this.first<next*.[pos.next=null]>last];
}
type Node = {
  data next: Node;
}

```

The notation is explained in the following section, but intuitively the `last` pointer is annotated with a formula that constrains its destination to be the last `Node` in

the list. A candidate for verification is the following procedure which concatenates two such structures:

```
proc concat(data l1,l2: Head): Head
{
  if (l1.last!=null) { l1.last.next = l2.first; }
  else { l1.first = l2.first; }
  if (l2.first!=null) { l1.last = l2.last; }
  l2.first = null;
  l2.last = null;
  return l1;
}
```

These are tedious pointer manipulations that are easy to get wrong. However, if we annotate the procedure with the pre-condition that `l1` and `l2` are not `null` and run PALE, it will in half a second report that no memory errors occur and, importantly, that the data structure invariant is maintained.

Related Work

General theorem provers, such as HOL [5], may consider the full behavior of programs but are slow and not fully automated. Tools such as ESC [12] and LCLint [15] consider memory errors among other undesirable behaviors but usually ignore data structure invariants. Also, they trade soundness for efficiency and hence may flag false errors.

Model checkers such as Bebop [2] and Bandera [9] abstract away the heap and only verify properties of control flow. The JPF [20] model checker verifies simple assertions for a subset of Java, but does not consider structural invariants.

The constraint solver Alloy has been used to verify properties about bounded initial segments of computation sequences [24]. While this is not a complete decision procedure even for straight-line code, it finds many errors and can produce counterexamples. With this technique, data structure invariants can be expressed in first-order logic with transitive closure. However, since it assumes computation bounds, absence of error reports does not imply a guarantee of correctness, and the technique does not appear to scale.

The symbolic executor PREFIX [7] simulates unannotated code through possible executions paths and detects a large class of errors. Again, this is not a complete or sound decision procedure, and data structure invariants are not considered. However, PREFIX gives useful results on huge source programs.

Verification based on static analysis has culminated with shape analysis. The goals of the shape analyzer TVLA [32, 38, 31] are closest to ours but its approach is radically different. Rather than encoding programs in logic, TVLA performs fix-point iterations on abstract descriptions of the store. Regarding precision and speed, PALE and TVLA seem to be at the same level. TVLA can handle some data abstractions and hence reason about sorting algorithms; we show in Section 6 that

we can do the same. TVLA analyzes programs with only pre- and post-conditions, where PALE often uses loop invariants and assertions. This seems like an undisputed advantage for TVLA; however, not having invariants can cause a loss in precision making TVLA falsely reject a program. Regarding the specification of new data structures we claim an advantage. Once a graph type has been abstractly described with PAL, the PALE tool is ready to analyze programs. In TVLA it is necessary to specify in three-valued logic an operational semantics for a collection of primitive actions specific to the data structure in question. Furthermore, to guarantee soundness of the analysis, this semantics should be proven correct by hand. TVLA is applicable also to data structures that are not graph types, but so far all their examples have been in that class. Unlike PALE, TVLA cannot produce explicit counterexamples when programs fail to verify.

There exists a variety of assertion languages designed to express properties of data structures, such as ADDS [21], L_r [3], and Shape Types [18]. We rely on PAL since it provides a high degree of expressiveness while still having a decision procedure that works in practice.

2 Pointer Assertion Logic

In this section, we informally present the components of our framework. First, we describe the underlying *store model*. Second, we use the notion of *graph types* to describe data structures. Third, we employ a simple *programming language* to express data structure operations. And, finally, we use *program annotations* in the form of Pointer Assertion Logic formulas, for expressing properties of the program store. In the following, we present these concepts informally and refer the reader to [35] for formal definitions.

Store Model

In our model, the store consists of a *heap* and some *program variables*. The heap contains *records* whose *fields* are either *pointers* or *boolean* values. A pointer either has the value `null` or points to a record. Program variables are either *data variables* or *pointer variables*. A data variable is the root of a data structure, whereas a pointer variable may point to any record in the heap. This is a very concrete representation. We only abstract away arithmetic values and the actual addresses of records.

Graph Types

Collections of records and pointers can form any number of interesting data structures, which are generally expressed through an invariant on the allowed shapes. We wish to explicitly declare such data structures so that their invariants can be verified by our system. For this purpose, we use *graph types* [29] which is an intuitive notation that makes it feasible to describe complex structures. Invariants of graph

type structures can be expressed in monadic second-order logic on finite trees, which allows us to use the MONA tool to verify correctness.

A graph type is a tree-shaped data structure with extra pointers. The underlying tree is called the *backbone*. The constituent records have two kinds of fields: *data fields* which define the backbone, and *pointer fields* which may point anywhere in the backbone. To describe a structural invariant, a pointer field is annotated with a *routing expression* which restricts its destination. In the current work, we have generalized the annotations to be arbitrary formulas that may contain routing expressions as basic predicates. Another difference to [29] is that instead of building types from unions and records, we only use records and nullable pointers. Clearly, the two variations can encode each other; we choose the more primitive version, since it turns out to lead to a more efficient decision procedure. Our syntax and semantics of graph type declarations is described in the next section.

Surprisingly many data structures can be described as graph types. As a simple example, consider the type of binary trees where all nodes contain pointers to the root. In our notation, it looks like:

```
type Tree = {
  data left,right:Tree;
  pointer root:Tree[root<(left+right)*>this &
    empty(root^Tree.left union
      root^Tree.right)];
}
```

The syntax for formulas is presented below, but the restriction on the source, **this**, and destination, **root**, of the pointer is read as follows: **this** must be reachable from **root** by following a sequence of **left** or **right** pointers, and the set of **Tree** records having **left** or **right** pointers to the **root** must be empty. Another example is doubly-linked lists with boolean values:

```
type Node = {
  bool value;
  data next:Node;
  pointer prev:Node[this^Node.next={prev}];
}
```

Here, the set of nodes that can reach the **this** node through a **next** pointer must only contain the **prev** node. The convention that **{null}** is interpreted as the empty set handles the first node in the list.

Our benchmark programs cover a variety of data structures expressed as graph types, including singly-linked lists, doubly-linked lists with tail pointers, red-black search trees, and post-order threaded trees with parent pointers. Additional examples are presented in [29].

The Programming Language

A *program* consists of a set of declarations of types, variables, and procedures, specified by the following grammar:

```

typedecl  →  type T = { ( field ; ) * }
field     →  data p⊕: T
          |  pointer p⊕: T [ form ]
          |  bool b⊕
progvar   →  data p⊕: T
          |  pointer p⊕: T
          |  bool b⊕
procedure →  proc n ( progvar⊗ ) : ( T | void )
          ( logicvar ; ) *
          [ form ] { ptrdir⊗ }
          ( { ( progvar ; ) * stm } ) ?
          [ form ] { ptrdir⊗ }

```

We use the notation \oplus and \otimes for comma-separated lists with one-or-more elements and zero-or-more elements, respectively. T , p , b , and n range over names of types, pointer variables or fields, boolean variables or fields, and procedures, respectively. Ignore for now all *logicvars*, *forms*, and *ptrdirs*; they are introduced in the next section. A **type** consists of a number of fields of kind **data**, **pointer**, or **bool**. The **data** fields span the tree value and the **pointer** fields define extra pointers whose destinations are constrained by a formula. The **bool** fields are used to model finite values. A **procedure** has a name, formal parameters, a return type, and a body consisting of local variable declarations and statements. If the body is omitted, the declaration is considered a *prototype*. A statement is one of the following constructs; the **assert** and **split** statements are described later:

```

stm  →  stm stm
      |  asn⊕;
      |  if ( condexp ) { stm } ( else { stm } ) ?
      |  assert [ form ] { ptrdir⊗ }
      |  while [ form ] { ptrdir⊗ } ( condexp ) { stm }
      |  return progexp ;
      |  split [ form ] { ptrdir⊗ } [ form ] { ptrdir⊗ }
      |  proccall ;
asn  →  lbexp = ( condexp | proccall )
      |  lptexp = ( ptexp | proccall )

```

The language permits multiple-assignment statements where all right-hand sides are evaluated before assigning—these are useful for certain program transformations. Expressions have the following form:

<i>condexp</i>	→	<i>bexp</i>		?		[<i>form</i>]
<i>bexp</i>	→	(<i>bexp</i>)		!		<i>bexp</i>
		<i>bexp</i> & <i>bexp</i>				<i>bexp</i> <i>bexp</i>
		<i>bexp</i> => <i>bexp</i>				<i>bexp</i> <=> <i>bexp</i>
		<i>bexp</i> = <i>bexp</i>				<i>ptrexp</i> = <i>ptrexp</i>
		<i>bexp</i> != <i>bexp</i>				<i>ptrexp</i> != <i>ptrexp</i>
		<i>true</i> <i>false</i>				<i>lbexp</i>
<i>lbexp</i>	→	<i>b</i>		<i>ptrexp</i> .		<i>b</i>
<i>ptrexp</i>	→	null		<i>lptrexp</i>		
<i>lptrexp</i>	→	<i>p</i>		<i>ptrexp</i> .		<i>p</i>
<i>proccall</i>	→	<i>n</i> ((<i>condexp</i> <i>ptrexp</i>) [⊗])				[<i>formula</i>]

The “?” operator stands for nondeterministic boolean choice, which is used to model arithmetic conditions that we cannot capture precisely. The operator “.” dereferences a pointer, and the other constructs have the expected meanings.

The language does not contain arithmetic, since our approach focuses on the structural aspects of data types. However, as described in a later section, the technique does permit abstractions of arithmetic properties, for instance for specifying certain ordered data structures.

Program Annotations

Pointer Assertion Logic is a *monadic second-order logic on graph types*. It allows quantification over heap records, both of individual elements and of sets of elements, and uses routing expressions [29] for convenient navigation in the heap. Formulas are used in pointer fields to constrain their destinations, in **while**-loops and procedure calls as invariants, in procedure declarations as pre- and post-conditions, and in **assert** and **split** statements. The syntax of formulas is as follows:

<i>form</i>	→	(<i>existpos</i> <i>allpos</i>) <i>p</i> [⊕] of <i>T</i> : <i>form</i>
		(<i>existset</i> <i>allset</i>) <i>s</i> [⊕] of <i>T</i> : <i>form</i>
		(<i>existptr</i> <i>allptr</i>) <i>p</i> [⊕] of <i>T</i> : <i>form</i>
		(<i>existbool</i> <i>allbool</i>) <i>s</i> [⊕] : <i>form</i>
		(<i>form</i>) ! <i>form</i>
		<i>form</i> & <i>form</i> <i>form</i> <i>form</i>
		<i>form</i> => <i>form</i> <i>form</i> <=> <i>form</i>
		<i>ptrexp</i> in <i>setexp</i> <i>setexp</i> sub <i>setexp</i>
		<i>setexp</i> = <i>setexp</i> <i>setexp</i> != <i>setexp</i>
		empty (<i>setexp</i>) <i>bexp</i>
		return <i>n</i> . <i>b</i>
		<i>m</i> ((<i>form</i> <i>ptrexp</i> <i>setexp</i>) [⊗])
		<i>ptrexp</i> < <i>routingexp</i> > <i>ptrexp</i>
<i>predicate</i>	→	pred <i>m</i> (<i>logicvar</i> [⊗]) = <i>form</i>

The identifiers *m* and *s* denote predicates and set variables, respectively. The **pos** and **ptr** quantifiers differ in that the former range over heap records while the latter

also includes the `null` value. A routing expression formula $p_1 \langle r \rangle p_2$ is satisfied by a given model if there is a path from p_1 to p_2 satisfying r , as defined below. For reuse of formulas, predicates can be defined as top-level declarations.

Logical variables can be associated to procedures to allow the pre- and post-conditions to be related, as commonly seen in the literature [19, 11]:

$$\begin{array}{l} \text{logicvar} \rightarrow \text{pointer } p^\oplus : T \\ \quad | \quad \text{bool } b^\oplus \\ \quad | \quad \text{set } s^\oplus : T \end{array}$$

In formulas, *ptrexp* has two additional forms allowing access in procedure post-condition to the returned value and in procedure call formulas to the logical variables of the called procedure:

$$\text{ptrexp} \rightarrow \dots \mid \text{return} \mid n . p$$

Set expressions can contain the usual set operators, along with the *up* operation $x \hat{\ } T . p$ which denotes the set of records of type T having a p successor to x :

$$\begin{array}{l} \text{setexp} \rightarrow s \\ \quad | \quad \text{ptrexp } \hat{\ } T . p \\ \quad | \quad \{ \text{ptrexp}^\oplus \} \\ \quad | \quad \text{setexp union setexp} \\ \quad | \quad \text{setexp inter setexp} \\ \quad | \quad \text{setexp minus setexp} \end{array}$$

The syntax of routing expressions is a slightly generalized version of that in [29]. A routing expression is a regular expression over routing directives, each being a step *down* or *up* a pointer or data field, or a formula with the extra free variable *pos* filtering away those records that cause the formula to evaluate to false when *pos* denotes one of them:

$$\begin{array}{l} \text{routingexp} \rightarrow p \mid \hat{\ } T . p \mid [\text{form}] \\ \quad | \quad \text{routingexp} . \text{routingexp} \\ \quad | \quad \text{routingexp} + \text{routingexp} \\ \quad | \quad (\text{routingexp}) \mid \text{routingexp} * \end{array}$$

By default, a pointer field must satisfy the formula given in its type declaration. This can be overridden with *pointer directives* of the form:

$$\text{ptrdir} \rightarrow T . p [\text{form}]$$

They allows pointer fields to be constrained differently at different program points. This is important because temporary but intentional invalidation of data structure invariants often occurs in imperative programs, as noted for instance in [21]. Pointer directives are required to be *well-formed*. This means that in any store and for any record, the directives associated to the pointer fields must denote *exactly one* record.

Fortunately, as proved in [29] this is decidable.

A pair consisting of a formula and a set of pointer directives:

$$[\textit{form}] \{ \textit{ptrdir}, \dots \}$$

is called a *property* and denotes the set of stores where:

- the formula is satisfied;
- the `data` variables denote disjoint acyclic backbones spanning the heap; and
- each `pointer` field satisfies its pointer directive.

Properties occur as procedure *pre-* and *post-conditions*, as *while loop invariants*, as *split assertions* and *assumptions* (`split` contains two properties), and as *assert assertions*.

Semantics of Annotations

The program annotations are invariants of the program that must be interpreted as follows:

- the pre-condition of a procedure may be assumed to hold when evaluating the procedure body;
- the post-condition must hold upon termination of the procedure body;
- every `while` loop invariant must hold upon entry and after each iteration, and may be assumed to hold when the loop terminates;
- assertions specified with `assert` must hold at those program points;
- for `split` statements, the assertion properties must hold, and the assumption properties may be assumed to hold; and
- at every procedure call, the invariant conjoined with the pre-condition of the called procedure must hold for some valuation of its logical variables, and the invariant conjoined with the post-condition may be assumed upon return, also for some valuation of the logical variables.

In later sections, we show that the requirements imposed by the annotations can be verified automatically, provided that valid and sufficiently detailed invariants are given.

3 Example: Threaded Trees

Before describing our decision procedure, we show a larger example of using PAL. A *threaded tree* is a binary tree in which all nodes contain a pointer to its cyclic successor in a post-order traversal. As a further complication, we equip all nodes with a parent pointer as well. This corresponds to the following graph type:

```
type Node = {
  data left,right:Node;
  pointer post:Node[POST(this,post)];
  pointer parent:Node[PARENT(this,parent)];
}
```

where POST and PARENT are predicates that spell out these relationships. For example, PARENT(a,b) abbreviates the formula:

```
a^Node.left union a^Node.right={b}
```

The POST predicate is more involved and makes use of auxiliary predicates LEAF, ROOT, and LESSEQ.

We consider a procedure `fix(x)` that assigns the correct value to `x.post` assuming that this field initially contains the value `null` and that `x` is non-`null`. This is a non-trivial operation that looks like:

```
proc fix(pointer x: Node): void
{
  if (x.left=null & x.right=null) {
    if (x.parent=null) { x.post = x; }
    else {
      if (x.parent.right=null | x.parent.right=x) {
        x.post = x.parent;
      }
      else {
        x.post = findsmallest(x.parent.right)
      }
    }
  }
  else { x.post = findsmallest(x) }
}
```

where the auxiliary procedure `findsmallest` is:

```
proc findsmallest(pointer t: Node): Node
  pointer T: Node;
{
  while (t.left!=null | t.right!=null) {
    if (t.left!=null) { t = t.left; }
  }
```

```

    else { t = t.right; }
  }
  return t;
}

```

The question is: Does this code verify? Does the resulting tree always satisfy the data structure invariant? Can type or memory errors ever occur? PALE can provide the answers with some help from us. First, since the argument to `fix` is not a proper threaded tree, we must state a suitable pre-condition as the property:

```
[x!=null]{Node.post[ALMOSTPOST(this,post,x)]}
```

Here we require that the argument is not `null` and that the data structure invariant can be temporarily violated. The `ALMOSTPOST` predicate is:

```
(this!=x => POST(this,post)) & (this=x => post=null)
```

which simply states the exception that we allow. Second, the `while`-loop in `findsmallest` needs an invariant, which is the property:

```
[INV]{Node.post[ALMOSTPOST(this,post,x)]}
```

where the pointer directive states that the threaded tree is still messed up, and the proper invariant `INV` equals:

```
T<(left+right)*>t &
allpos c of Node: LESSEQ(c,t,T) => t<(left+right)*>c
```

which states that `t` is a descendant of `T` and all its post-order successors are further descendants. See [35] for the full code with all post-conditions. In total, six annotations are required. In less than 8 seconds PALE verifies that the code contains no errors.

4 Hoare Logic Revisited

Given an annotated program, we wish to decide whether the program is correct with respect to the annotations. The first step in our decision procedure is to split the given program into Hoare triples [22, 1, 11]. The idea of modeling transformations of the heap with Hoare logic has received much attention [23, 37, 17]. The novelty of our approach is the choice of PAL as assertion language. Our Hoare triples have a nonstandard form:

$$\text{triple} \rightarrow [\text{form}] \{ \text{ptrdir}, \dots \} \text{stm}$$

The statement *stm* is not allowed to contain `while`-loops, `split` statements, or procedure calls. As opposed to normal Hoare triples, these have no explicit post-condition, but the *stm* part may contain `assert` sub-statements. This simple generalization allows many assertions to be made without always breaking triples into smaller parts, as was often the case in [25] and [14]. Also, it simplifies the encoding in monadic second-order logic described in Section 5. A triple is *valid* if executing *stm* in a store where the property defined by *form* and the *ptrdirs* is satisfied cannot violate any assertions specified by `assert` statements occurring in *stm*.

We define the *cut-points* of a program (according to [16]) as the following set of program points: the beginning and end of procedure bodies and `while` bodies, the `split` statements (these do not affect the computation and are considered single program points), and before each procedure call.

For each cut-point, we generate a Hoare triple from the property associated with that point and the code that follows until reaching other cut-points. Recall that we do allow `if` statements in the Hoare triples. If one branch contains a cut-point, we require syntactically that the other also contains one, for instance a `split` statement. Thus, the statement part of a Hoare triple has in general a tree shape with one cut-point in the root and one in each leaf. Extra `assert` statements are automatically inserted for the leaf cut-points, reflecting the assertions they define. In case of `split` statements, we use the assertion property. For procedure calls, we use the post-condition property of the called procedure conjoined with the call invariant formula. See [35] for more details.

We claim without proof that this reduction is semantically sound, with two exceptions:

- For `split` statements, the assertion property may not be implied by the assumption property, thereby causing a “gap” between the Hoare triples. This is intentional, because it allows to recover from situations where the required properties are beyond what is expressible in Pointer Assertion Logic, such as arithmetical properties. Using `split` statements at a few selected places, one can then still verify properties of the remaining parts of the code. However, none of the examples shown in Section 7 require this feature.
- Procedure calls are known to cause complications for Hoare logic [11]. In our case, there is in general no guarantee that the call invariant is actually a valid invariant. However, in most situations, simple syntactic requirements suffice, since recursive calls in data type operations typically follow the recursive structure of the graph type backbones. A sufficient condition is that the call invariant only accesses variables and record fields that are not assigned to in the procedure. Such requirements ensure that the invariant and the procedure’s pre- and post-conditions express properties of disjoint parts of the store, reminiscent of the “independent conjunctions” in [37]. All the examples shown in Section 7 can be handled by simple rules, which we plan to build into PALE.

In PALE, this phase is implemented as a desugaring process reducing all procedures, `while`-loops, `split` statements, and procedure calls to *transduction* declarations having the form “`transduce triple`”. In the following section we describe how validity of these simpler `transduce` constructs can be decided.

In contrast to techniques based on generating the weakest preconditions for all procedures, each program or procedure is not turned into one single verification condition; instead we use the annotations to split the program into Hoare triples that are verified independently. Also, as opposed to [17], we will not rely on fixpoint iterations. This means that detailed invariants may be required; however, it has the advantage that the technique becomes highly modular and hence scalable.

5 Deciding Hoare Triples in MONA

We need to decide validity of a Hoare triple of the form:

$$[\textit{form}] \{ \textit{ptrdir}, \dots \} \textit{stm}$$

where the statement is without loops and procedure calls. The question is whether any execution of *stm* starting from a store satisfying the formula and the pointer directives is guaranteed to satisfy the assertions given by `assert` statements and to result in stores with disjoint, acyclic backbones spanning the heap in which all pointer directives hold. A result in [30] shows in a very general setting that this is a decidable question.

Similarly to the previous implementations [25, 14] we use a particular *transduction* technique. This idea allows us to avoid an explicit construction of weakest pre-conditions working backwards through the statement sequence. Instead, we directly simulate (transduce) the statements and mirror their effect by updating a fixed collection of *store predicates* which abstractly describes a set of stores. It is shown in [30] that any question about the resulting set of stores can be answered by phrasing it in terms of the transduced store predicates and checking for validity of the resulting formula.

The store predicates describe a set of stores in MONA logic. They can be thought of as an interface for asking questions about a store. There are 11 kinds of predicates:

- `bool_T_b(v)` gives the value of the `bool` field *b* in a record *v* of type *T*;
- `succ_T_d(v,w)` holds if the record *w* is reachable from the record *v* of type *T* along a data field named *d*;
- `null_T_d(v)` holds if the data field *d* in the record *v* of type *T* is `null`;
- `succ_T_p(v,w)` holds if the record *w* is reachable from the record *v* of type *T* along a pointer field named *p*;
- `null_T_p(v)` holds if the pointer field *p* in the record *v* of type *T* is `null`;

- $\text{ptr}_d(v)$ holds if the record v is the value of the data variable d ;
- $\text{null}_d()$ holds if the data variable d is `null`;
- $\text{ptr}_p(v)$ holds if the record v is the destination of the pointer variable p ;
- $\text{null}_p()$ holds if the pointer variable p is `null`;
- $\text{bool}_b()$ gives the value of the boolean variable b ;
- $\text{memfailed}()$ holds if a memory error has occurred.

All properties of a store can be expressed using these predicates in monadic second-order logic on finite trees. The transduction process generates a collection of such store predicates for each program point. For convenience, we describe this by indexing the predicates with program points; for example, for each program point i there is a version of the `bool` predicate called $\text{bool}_{T_b_i}(v)$.

An initial collection of store predicates is defined to reflect the formula and pointer directives that constitute the pre-condition of the Hoare triple. In the encoding into MONA code, the program variables are modeled as free variables, which are universally quantified in the final validity formula that is given to MONA. For example, a `bool` variable is modeled as a boolean variable `_bool_b` in MONA and the corresponding initial store predicate is:

$$\text{bool}_{b_0}() = _bool_b$$

Similarly, a pointer variable p of type T is modeled as a second-order variable $_ptr_{T_p}$ and the corresponding initial store predicate is:

$$\text{bool}_{T_p_0}(v) = v = _ptr_{T_p}$$

since records are modeled as singleton sets. A `bool` field b in a record of type T is modeled as a second-order variable $_bool_{T_b}$ containing the set of records in which b is true. Consequently, the corresponding initial store predicate is:

$$\text{bool}_{T_b_0}(v) = v \text{ sub } _bool_{T_b}$$

As a final example, we consider pointer fields whose initial store predicate is:

$$\text{succ}_{T_p_0}(\text{this}, p) = \text{form}'$$

where form' is the encoding of the formula associated with the p field of T . If the pre-condition of the Hoare triple contains the pointer directive $T.p[\text{form}]$, then that formula is form , otherwise the default formula from the type definition is used.

Across a simple statement, two collections of store predicates are related in a manner that reflects the semantics of that statement. Consider for example a type of linked lists:

```
type Node = { data next: Node; }
```

and a simple statement involving two pointer variables of type `Node`:

```
p = q.next;
```

If this statement is enclosed by program points i and j , then the store predicates are updated as follows in MONA code:

```
memfailed_j() = memfailed_i() | null_q_i()  
ptr_p_j(v) = ex2 w: succ_Node_next_i(w,v) & ptr_q_i(w)  
null_p_j() = ex2 w: ptr_q_i(w) & null_Node_next(w)
```

while the other store predicates remain unchanged. The PALE tool generates such store predicate updates for all Hoare triples and subsequently generates formulas to check the required properties. Between conditionals, routing expressions, and various primitive statements this is a complex translation reminiscent of generating machine code in a compiler. The details can be studied in [35]. The way assignments are handled without losing aliasing information, as in the example above, is essentially the same as in [36].

Checking that the resulting backbones are disjoint, acyclic, and span the heap is based on formulas for expressing transitive closure. Checking that a pointer directive holds is in [29] shown to be decidable in monadic second-order logic. This result generalizes easily to our extension of graph types, where arbitrary formulas rather than only routing expressions can be used as pointer directives.

The MONA tool [27, 26, 34] transforms the resulting formula, which can be quite large, into an equivalent minimal Guided Tree Automaton [4] represented as a BDD structure [6], and from that either deduces validity or generates a counterexample model. In the latter case, the PALE tool decompiles that model into a program store which causes the program to fail. The use of Guided Tree Automata rather than ordinary tree automata yields an exponential saving by factorizing the state space according to the recursive structure of the graph type backbones. Compared to the WSRT technique used in [14], our choice of describing the backbones as records with pointers rather than as recursive types allow a simpler and more efficient automaton guide to be constructed.

Note that a collection of store predicates is vaguely similar to the abstract store descriptions employed by TVLA. Consequently, it might seem that we could follow their approach and use a fixpoint process to transduce a `while`-loop. However, this is in general not possible, since such fixpoints may require transfinite induction. Hence, we resort to using invariants to break up loops.

This transduction approach introduces no imprecision; it is both sound and complete for individual Hoare triples.

6 Data Abstractions

In [38, 31], abstractions of the data contained in the heap records can be tracked by specifying suitable *instrumentation predicates*. As an example, a predicate $dle(x, y)$ is used to represent “the data in x is less than or equal to the data in y ”. To illustrate the power of PAL, we show that a similar approach works for our technique.

As an example, we instrument the ubiquitous linked-list `reverse` example to verify that reversal of a list ordered in increasing order results in a list ordered in decreasing order:

- We associate two boolean fields, `next_dle` and `next_dge`, to the `next` field in the linked-list type, with the intended meaning: `next_dle` is true in a given record if the data in the record denoted by the `next` pointer is certain to be *less than or equal* to the data in the given record – and likewise for `next_dge` with *greater than or equal*.
- Similarly, for each pair of program pointer variables, two boolean variables are added to keep track of the relative order of the records being pointed to. With a subsequent dead-code elimination, a total of three boolean variables suffice.
- For each pointer assignment, the new boolean fields and variables are updated accordingly. For instance,

```
list.next = res;
```

is replaced by the multiple-assignment statement:

```
list.next = res, list.next_dle = res_dle_list;
```

reflecting the change of the `next` field.

If arbitrary PAL formulas are allowed as right-hand sides of the new assignments, even complex reachability properties can be captured. For this example, simple assignment suffice, though. As in [31], this is also sufficient to verify for instance that `bubblesort` actually sorts the elements.

The intellectual effort needed to update the data abstraction bits seems to be the same as to define the required operational semantics in TVLA. As hinted in the example, some degree of automation is possible for our technique; however, we leave that for future work.

Note that many data structures, in particular variations of search trees, can be abstractly described by associating to every node a few of bits of information summarizing properties of the tree. Those data structures can also easily be verified using techniques like these.

7 Implementation and Evaluation

Our verification technique is implemented in a tool called PALE, the Pointer Assertion Logic Engine. Given an annotated program, PALE checks that:

- the pointer directives are well-formed;
- `null` pointer dereferences cannot occur;
- at each cut-point that the `data` variables contain disjoint, acyclic backbones spanning the heap and that the assertions and pointer directives are satisfied;
- all `assert` assertions are valid; and
- all cut-point properties are satisfiable.

There is not necessarily an error in the program if a cut-point property is unsatisfiable, but it usually indicates an error in the specification. As in [25] and [14], memory allocation can easily be expressed such that the tool would also check for memory leaks and dangling references.

Using PALE, we have evaluated the technique on a number of examples dealing with a variety of data structures. In all cases, we check for memory errors and possible violations of the data structure invariants:

- *Singly-linked lists* with the operations `reverse`, `search`, `zip`, `delete`, `insert`, and `rotate`. These examples have been scrutinized before [8, 25, 32]. We also include the `concat` operation on lists with tail pointers from Section 1. We have tried `bubblesort` as in [31] but with various degrees of abstraction of the data: In `bubblesort_simple`, the record values are abstracted away so only `null` pointer dereferences are checked for; in `bubblesort_boolean`, the values are abstracted to booleans which in the post-condition are checked to be properly sorted; and in `bubblesort_full`, the data abstraction technique from Section 6 is used as in [31] to conclude that the resulting lists are sorted. We also use data abstractions in `orderedreverse` to show that `reverse` switches the order of a sorted list. Finally, we try `recreverse`, which is a recursive version of `reverse`.
- *Doubly-linked lists with tail pointers* [29] with the operations `delete`, `search`, `insert`, and `concat`.
- *Red-black search trees* [10] with the standard operations `leftrotate`, `rightrotate`, `treeinsert`, and `redblackinsert`. We include the non-arithmetic part of the red-black search tree invariant, that is, that the root is black and red nodes have black children:

```
BLACK(root) &  
allpos q of Node: ROOT<(left+right)*>q =>  
  (RED(q) => BLACK(q.left) & BLACK(q.right));
```

Example name	Lines of code	GTA operations	Largest GTA		Time (sec.)	Memory (MB)
			States	BDD nodes		
reverse	16	1,262	35	142	0.61	3
search	12	1,150	27	78	0.35	2
zip	33	1,968	174	730	4.91	11
delete	22	966	73	349	1.38	5
insert	33	1,006	103	443	2.73	7
rotate	11	582	44	213	0.22	2
concat	24	1,134	48	177	0.53	3
bubblesort_simple	43	1,732	373	3,267	3.05	18
bubblesort_boolean	43	2,218	357	3,922	3.77	13
bubblesort_full	43	2,694	373	3,291	4.63	19
orderedreverse	24	1,318	32	105	0.59	3
rereverse	15	1,214	42	176	0.40	3
doublylinked	72	5,526	241	796	11.72	14
leftrotate	30	1,310	165	1,550	4.64	8
rightrotate	30	1,310	165	1,550	4.79	8
treeinsert	36	2,482	137	844	9.54	32
redblackinsert	57	6,946	298	2,419	36.77	44
threaded	54	9,348	50	248	7.77	10

Figure 1: Statistics from PALE experiments.

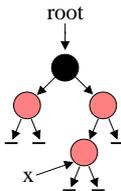
- *Threaded trees* [29], as shown in Section 3, where every node has a pointer to its post-order cyclic successor and a pointer to its parent, with a `fix` operation for reestablishing the correct `post` pointer for a given node.

The resources for translation into MONA code and for the automaton analysis are negligible. Figure 1 shows the time and space consumptions of the MONA automaton operations (on a 466MHz Celeron PC) for the examples, along with the number of GTA operations (here we count only the essential operations: minimization, projection, and product), the size of the largest intermediate minimized automaton (in number of states and in number of BDD nodes). Note that some examples implement individual operations while others implement whole data types. The lines of code measure the underlying program only, thus disregarding the PAL annotations. The table shows that the examples typically run in seconds despite requiring a quite large number of automaton operations. Since the complexity is non-elementary in the size of the program, intractable examples do exist but they do not seem to occur often in practice.

The code for the `bubblesort` examples (excluding annotations) is taken from [31]. Interestingly, PALE discovered a minor bug (a null-pointer dereference) even though the code had allegedly been verified by TVLA, which spent 245 seconds compared to 5 seconds for PALE. This huge speedup shows an instance where using invariants

is much faster than performing a fixpoint iteration. This suggests that PALE may be quite scalable. Another noteworthy point discovered by PALE is that in [10], the authors forget to require the root to be initially black in `redblackinsert`.

Versions with plausible bugs planted typically take roughly the same time to process as the correct programs. For such buggy versions, counterexamples are generated, which is crucial for determining whether the error is in the program, the assumptions, or the assertions. As an example, if a conditional in `redblackinsert` erroneously tests for a specific node to be black rather than red, PALE produces the following counterexample store for the Hoare triple containing the conditional:



Here, the root node is black and the others are red, and we omit field names and all pointer fields. Such a counterexample is clearly useful for locating the bug. Notice that for this bug, the approach in [24] would not find the bug for heap bounds of less than four records.

The experiments show that our approach does work in practice for non-trivial data structures, and with time and space requirements which are as good as or better than those for the previous more specialized versions [25, 14] and related approaches with similar goals [31, 24, 13, 17].

8 Conclusion

It is well known that developing formal program specifications is expensive, but for some safety critical applications a guarantee of partial correctness of data type implementations can be worth the effort. A tool such as PALE can be used to verify specifications expressible in Pointer Assertion Logic, and also to guide the programmer by the generation of counterexamples. With verification techniques based on undecidable logics, the programmer suffers the burden of guiding a theorem prover to the proofs, not even being certain that they exist. With less expressive techniques, important aspects of the data types may not be expressible and hence not verifiable. We believe that Pointer Assertion Logic is a fruitful compromise between expressibility and ease of use.

Although facing a non-elementary theoretical complexity, the examples we provide show that logic and automaton based program verification is feasible. Furthermore, we believe that the efficiency of the implementation can be improved by at least an order of magnitude by tuning the MONA tool using heuristics as proposed in [28]. As also suggested in [24, 20] we may benefit from an initial simplification phase that performs program slicing or partial evaluation of the source programs.

Future work will also examine the possibility of incorporating simple arithmetic to the language. The MONA tool can also be used as an efficient decision procedure for Presburger arithmetic [39, 27], which is sufficient for many properties.

The full source code for the PALE tool, the examples, and a detailed description of the desugaring and code generation to MONA are available from the PALE site at <http://www.brics.dk/PALE/>.

References

- [1] Krzysztof R. Apt. Ten years of Hoare’s logic: A survey—part I. *ACM Transactions on Programming Languages*, 3(4):431–483, 1981.
- [2] Thomas Ball and Sriram K. Rajamani. Bebop: A symbolic model checker for boolean programs. In *Proceedings of the SPIN Software Model Checking Workshop*, volume 1885 of *LNCS*, 2000.
- [3] Michael Benedikt, Thomas Reps, and Mooly Sagiv. A decidable logic for describing linked data structures. In *European Symposium On Programming*, volume 1576 of *LNCS*, 1999.
- [4] Morten Biehl, Nils Klarlund, and Theis Rauhe. Algorithms for guided tree automata. In *First International Workshop on Implementing Automata, WIA ’96*, volume 1260 of *LNCS*, 1997.
- [5] Paul E. Black and Phillip J. Windley. Inference rules for programming languages with side effects in expressions. In *International Conference on Theorem Proving in Higher Order Logics*, 1996.
- [6] Randal E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, Aug 1986.
- [7] William R. Bush, Jonathan D. Pincus, and David J. Sielaff. A static analyzer for finding dynamic programming errors. *Software: Practice and Experience*, 30(7), 2000.
- [8] Stephen A. Cook and Derek C. Oppen. An assertion language for data structures. In *2nd ACM Symposium on Principles of Programming Languages*, pages 160–166, 1975.
- [9] C. Corbett, Matthew B. Dwyer, John Hatcliff, and Robby. A language framework for expressing checkable properties of dynamic software. In *Proceedings of the SPIN Software Model Checking Workshop*, volume 1885 of *LNCS*, 2000.
- [10] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. MIT Press, 1990.

- [11] Patrick Cousot. Formal models and semantics. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 841–993. MIT Press/Elsevier, 1990.
- [12] David L. Detlefs, K. Rustan M. Leino, Greg Nelson, and James B. Saxe. Extended static checking. Research Report 159, Compaq Systems Research Center, December, 1998.
- [13] Nurit Dor, Michael Rodeh, and Mooly Sagiv. Checking cleanness in linked lists. In *Seventh International Static Analysis Symposium*, 2000.
- [14] Jacob Elgaard, Anders Møller, and Michael I. Schwartzbach. Compile-time debugging of C programs working on trees. In *European Symposium on Programming Languages and Systems*, volume 1782 of *LNCS*, 2000.
- [15] David Evans. Static detection of dynamic memory errors. In *Programming Language Design and Implementation, PLDI'96*, 1996.
- [16] Robert W. Floyd. Assigning meanings to programs. *Mathematical Aspects of Computer Science*, *American Math. Soc.*, 1967.
- [17] Pascal Fradet, Ronan Gaugne, and Daniel Le Métayer. Static detection of pointer errors: An axiomatisation and a checking algorithm. In *6th European Symposium on Programming*, volume 1058 of *LNCS*, 1996.
- [18] Pascal Fradet and Daniel Le Métayer. Shape types. In *Principles of Programming Languages, POPL'97*. ACM, 1997.
- [19] David Gries. *The Science of Programming*. Springer-Verlag, 1981.
- [20] Klaus Havelund and Thomas Pressburger. Model checking Java programs using Java PathFinder. *International Journal on Software Tools for Technology Transfer*, 2(4), April 2000.
- [21] Laurie Hendren, Joseph Hummel, and Alexandru Nicolau. Abstractions for recursive pointer data structures: Improving the analysis and transformation of imperative programs. In *Conference on Program Language Design and Implementation, PLDI'92*, 1992.
- [22] C.A.R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, October 1969.
- [23] C.A.R. Hoare. Proof of correctness of data representations. *Acta Informatica*, 1:271–281, 1971.
- [24] Daniel Jackson and Mandana Vaziri. Finding bugs with a constraint solver. In *International Symposium on Software Testing and Analysis, ISSTA'00*. ACM, 2000.

- [25] Jacob L. Jensen, Michael E. Jørgensen, Nils Klarlund, and Michael I. Schwartzbach. Automatic verification of pointer programs using monadic second-order logic. In *PLDI '97*, 1997.
- [26] Nils Klarlund. Mona & Fido: The logic-automaton connection in practice. In *Computer Science Logic, CSL '97*, volume 1414 of *LNCS*, 1998.
- [27] Nils Klarlund and Anders Møller. *MONA Version 1.3 User Manual*. BRICS Notes Series NS-98-3 (2nd revision), Department of Computer Science, University of Aarhus, October 1998.
- [28] Nils Klarlund, Anders Møller, and Michael I. Schwartzbach. MONA implementation secrets. In *Fifth International Conference on Implementation and Application of Automata, CIAA 2000*, LNCS, 2000.
- [29] Nils Klarlund and Michael I. Schwartzbach. Graph types. In *20th Symposium on Principles of Programming Languages*. ACM, 1993.
- [30] Nils Klarlund and Michael I. Schwartzbach. Graphs and decidable transductions based on edge constraints. In *Trees in Algebra and Programming, CAAP '94*, volume 787 of *LNCS*, 1994.
- [31] Tal Lev-Ami, Thomas Reps, Mooly Sagiv, and Reinhard Wilhelm. Putting static analysis to work for verification: a case study. In *International Symposium on Software Testing and Analysis, ISSTA '00*. ACM, 2000.
- [32] Tal Lev-Ami and Mooly Sagiv. TVLA: A system for implementing static analyses. In *Seventh International Static Analysis Symposium*, 2000.
- [33] Albert R. Meyer. Weak monadic second-order theory of successor is not elementary recursive. In R. Parikh, editor, *Logic Colloquium, (Proc. Symposium on Logic, Boston, 1972)*, volume 453 of *LNCS*, pages 132–154, 1975.
- [34] Anders Møller. MONA project home page. www.brics.dk/mona.
- [35] Anders Møller. PALE project home page. www.brics.dk/PALE.
- [36] Joseph M. Morris. A general axiom of assignment. In *Theoretical Foundations of Programming Methodology*, 1982.
- [37] John C. Reynolds. Intuitionistic reasoning about shared mutable data structure. In *Millennial Perspectives in Computer Science*, 2000.
- [38] Mooly Sagiv, Thomas Reps, and Reinhard Wilhelm. Parametric shape analysis via 3 valued logic. In *Symp. on Princ. of Prog. Lang., POPL'99*, 1999.
- [39] Thomas R. Shiple, James H. Kukula, and Rajeev K. Ranjan. A comparison of Presburger engines for EFSM reachability. In *Computer Aided Verification, CAV '98*, volume 1427 of *LNCS*, 1998.

Recent BRICS Report Series Publications

- RS-00-39** Anders Møller and Michael I. Schwartzbach. *The Pointer Assertion Logic Engine*. December 2000. 23 pp. To appear in *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '01 Proceedings, 2001*.
- RS-00-38** Bertrand Jeannot. *Dynamic Partitioning in Linear Relation Analysis: Application to the Verification of Synchronous Programs*. December 2000.
- RS-00-37** Thomas S. Hune, Kim G. Larsen, and Paul Pettersson. *Guided Synthesis of Control Programs for a Batch Plant using UP-PAAL*. December 2000. 29 pp. Appears in Hsiung, editor, *International Workshop in Distributed Systems Validation and Verification. Held in conjunction with 20th IEEE International Conference on Distributed Computing Systems (ICDCS '2000)*, DSVV '00 Proceedings, 2000.
- RS-00-36** Rasmus Pagh. *Dispersing Hash Functions*. December 2000. 18 pp. Preliminary version appeared in Rolim, editor, *4th. International Workshop on Randomization and Approximation Techniques in Computer Science, RANDOM '00, Proceedings in Informatics 8, 2000*, pages 53–67.
- RS-00-35** Olivier Danvy and Lasse R. Nielsen. *CPS Transformation of Beta-Redexes*. December 2000. 12 pp.
- RS-00-34** Olivier Danvy and Morten Rhiger. *A Simple Take on Typed Abstract Syntax in Haskell-like Languages*. December 2000. 25 pp. To appear in *Fifth International Symposium on Functional and Logic Programming, FLOPS '01 Proceedings, LNCS, 2001*.
- RS-00-33** Olivier Danvy and Lasse R. Nielsen. *A Higher-Order Colon Translation*. December 2000. 17 pp. To appear in *Fifth International Symposium on Functional and Logic Programming, FLOPS '01 Proceedings, LNCS, 2001*.
- RS-00-32** John C. Reynolds. *The Meaning of Types — From Intrinsic to Extrinsic Semantics*. December 2000. 35 pp.
- RS-00-31** Bernd Grobauer and Julia L. Lawall. *Partial Evaluation of Pattern Matching in Strings, revisited*. November 2000. 48 pp.