Scalable Specification and Reasoning: Challenges for Program Logic

Peter W. O'Hearn

Queen Mary, University of London

Abstract. If program verification tools are ever to be used widely, it is essential that they work in a modular fashion. Otherwise, verification will not scale. This paper discusses the scientific challenges that this poses for research in program logic. Some recent work on separation logic is described, and test problems that would be useful in measuring advances on modular reasoning are suggested.

1 Introduction

Software verification has seen an upsurge of interest in recent years. Partly this is a result of a convergence that has resulted from maturation of proof tools and lowering of aims, from full behavioural specifications to partial (often safety) properties of a system. Prominent examples include the SLAM model checker [2] and the ESC/Java static assertion checker [19,14]. But modularity is a problem.

Modularity is essential for scalable specification and reasoning. If we find ourselves in a position where the specification of one program component must talk about all other components in a system, or the states of other components, then we will very quickly be overwhelmed by the complexity of specifications. Programming features such as pointers (in various of their guises), concurrency and reflection raise particularly challenging problems for program logic. Simple methods for achieving modularity, such as listing the variables that might change (using "modifies" clauses), are not sufficient for common programs written in widely-used languages, which feature complex and dynamically changing interconnections between program components.

The problem faced by program logic is not an in-principle one – being able to describe behaviours at all – but rather is one of tractability. For example, when one considers programs with pointers and concurrency, reasoning with traditional program logic can become so complex as to be detached from computational intuition. The best way to illustrate this claim is with examples, and I consider three, describing what the more general technical challenges are as we go along. Some relevant work on separation logic [46,26,37,47] is described, and the promise of and problems for this approach are discussed. Finally, some wholly unresolved problems are mentioned.

There are many obstacles facing any Program Verifier challenge project [24] – particularly, the strength of theorem provers – and I am not saying that full solutions to the problems I discuss are necessary for it to have some success. My

B. Meyer and J. Woodcock (Eds.): Verified Software, LNCS 4171, pp. 116–133, 2008.

aim here is just to communicate some unsolved problems in program logic which, if progress were made on them, could have a considerable positive impact.

2 Framing and Indirection

I begin with a simple program and consider how one might specify it using traditional Floyd-Hoare logic. The specification is found to be unsatisfactory, and then is amended to provide a technically correct one. It is then argued that this technically correct specification is conceptually wrong.

2.1 An Incorrect Specification

Consider a procedure for disposing a tree, held as a linked structure in memory.

```
procedure DispTree(p)
local i, j;
if  p≠nil then
    i = p→l ; j:= p→r;
    DispTree(i);
    DispTree(j);
    dispose(p)
```

This is the expected procedure that walks a tree, recursively disposing left and right subtrees and then the root pointer. It uses a representation of tree nodes with left, right and data fields, and the empty tree is represented by nil.

A first attempt at a specification might be something like

$$\{\mathsf{tree}(p) \wedge \mathtt{reach}(p, n)\} \quad \mathtt{DispTree}(p) \, \{\neg \mathtt{allocated}(n)\}$$

assuming that we have defined the predicates that say when p points to a (binary) tree in memory, when n is reachable (following l and r links) from p, and when n is allocated. This spec says that any node n which is in the tree pointed to by p is not allocated on conclusion.

While this specification says part of what we would like to say, it leaves too much unsaid. It does not say what the procedure does to nodes that are not in the tree; we have left out the notorious *frame axioms* [33].

The result is that, while the specification is something that we would expect to be true of the procedure, it is too weak to use at many call sites. For example, consider the first recursive call, $\mathtt{DispTree}(i)$, to dispose the left subtree. If we use the specification (instantiating p by i) as an hypothesis, in the usual way when reasoning about recursive procedures [22], then we have a problem. For, the specification does not rule out the possibility that the procedure call alters the right subtree j, perhaps creating a cycle or even disposing some of its nodes. As a consequence, when we come to the second call $\mathtt{DispTree}(j)$, we will not know that the required $\mathsf{tree}(j)$ part of the precondition will hold. So our reasoning will get stuck.

The moral of this story is that [37]

> *if one does not have some way of representing or inferring frame axioms, then the proofs of even simple programs with procedure calls will not go through.*

The `DispTree` program makes this point especially vivid because of its use of recursion, where the spec and the call sites have to get along:

> *for recursive programs attention to framing is essential if one is to obtain strong enough induction hypotheses.*

The problem does not depend on having low-level operations such as pointer disposal. For example, specifying tree copying leads to similar difficulties.

2.2 An Unfortunate Fix

How can we fix the specification of `DispTree`? Here is my attempt:

$\{\texttt{tree}(p) \wedge \texttt{reach}(p,n) \wedge \neg\texttt{reach}(p,m) \wedge \texttt{allocated}(m) \wedge m.f = m' \wedge \neg\texttt{allocated}(q)\}$
$\texttt{DispTree}(p)$
$\{\neg\texttt{allocated}(n) \wedge \neg\texttt{reach}(p,m) \wedge \texttt{allocated}(m) \wedge m.f = m' \wedge \neg\texttt{allocated}(q)\}$

This says, in addition, that any allocated cell not reachable from p has the same contents in memory and that any previously unallocated cell remains unallocated. The additional clauses are the frame axioms. (I am assuming that m, m', n and q are auxiliary variables, guaranteed not to be altered. The reason why, say, the predicate $\neg\texttt{allocated}(q)$ could conceivably change, even if q is constant, is that the `allocated` predicate refers to a behind-the-scenes heap component. f is used in the spec as an arbitrary field name.)

I *believe* that this specification is strong enough to prove the procedure, but I have never attempted to carry out a proof. It would be complex. But, more importantly, I believe that the specification is badly wrong from a conceptual point of view.

The problem is not that we cannot specify `DispTree` at all, but rather is that final specification makes ugly statements about what is not reachable and what is not allocated that have, really, nothing to do with the program. Programmers *think locally*, and when reasoning about a program they concentrate on the resources that are relevant to its correct operating [37]. The need to state these frame axioms explicitly is violently at odds with programming intuition. So, even if technically alright, I view such a specification as conceptually wrong, a symptom of a problem in program logic.

2.3 The Frame Problem

The frame problem is that, traditionally, an inordinate amount of effort needs to be spent specifying what a program doesn't change, so much so that these frame axioms distract from the main concern – what changes [33]. In the absence

of pointers what doesn't change can be succinctly summarized using modifies clauses, which list the program variables corresponding to locations that can be altered by a program. But, in the presence of pointers of other forms of indirect addressing the relevant locations are not always directly named by program variables, and the idea of modifies clause is then much more difficult to make work. The unhappy consequence is that sound, modular specification methods are lacking for widely-used programming languages such as C and Java.

A full solution to the frame problem would allow us to make a positive statement about what changes, like in our first, faulty, specification, with the frame axioms coming along for free. A partial solution would at least let us represent the frame axioms compactly and intuitively.

The frame problem is extremely irritating. When you see it, you expect that there should be some sort of easy solution. It should be possible for a specification to say just what is relevant, like in our first specification of `DispTree`, and for the rest (the frame axioms) to come along for free. I have often felt that way.

The frame problem has been intensely studied in AI, and there are too many papers to survey here; I mention only one, the extremely clear paper of Reiter [44], which can serve as a good introduction to the problem. Unfortunately, there has been little crossover work applying the techniques there to programs (a notable exception is [9]). Although the frame problem is irritating, it is genuine, and a central problem in modular reasoning. But it is not the whole story, as we shall see in later sections.

The frame problem is stated above in a decidedly negative manner. I prefer to take a more positive perspective [37]:

> *When specifying a program, it should be possible to concentrate exclusively on the information (data, resources, etc) that is relevant to its correct operating. Any information it is independent of should not have to be mentioned.*

3 Separation Logic

The separation logic specification of `DispTree` is just

$$\{\mathsf{tree}(p)\}\ \texttt{DispTree}(p)\ \{\texttt{empty}\}$$

which says that if you have a tree at the beginning then you end up with the empty heap at the end. And the proof is very simple. The crucial part, in the `else` branch, looks like this:

$$\begin{array}{l}
\{p{\mapsto}[l\colon x, r\colon y] * \mathsf{tree}(x) * \mathsf{tree}(y)\} \\
\quad i := p{\to}l;\ j := p{\to}r; \\
\{p{\mapsto}[l\colon i, r\colon j] * \mathsf{tree}(i) * \mathsf{tree}(j)\} \\
\quad \texttt{DispTree}(i); \\
\{p{\mapsto}[l\colon i, r\colon j] * \mathsf{tree}(j)\} \\
\quad \texttt{DispTree}(j); \\
\{p{\mapsto}[l\colon i, r\colon j]\} \\
\quad \texttt{dispose}\ p; \\
\{\texttt{empty}\}
\end{array}$$

After we enter the conditional statement we know that $p{\neq}\mathsf{nil}$, so that p is an allocated node that points to left and right subtrees occupying separate storage. Then the roots of the two subtrees are loaded into i and j. Notice how the proof steps then follow operational intuition. The first recursive call removes the left subtree, the second call removes the right subtree, and the final instruction removes the root pointer p. This verification is carried out using the procedure specification as an assumption, as in the usual treatment of recursive procedures in Hoare logic [22].

I have just given you a proof snippet in what is probably an unfamiliar formalism, so some explanation is in order. To understand separation logic intuitively you should think in terms of *heaplets*, portions of heap, rather than the whole global heap. The separating conjunction $P * Q$ holds of a given heaplet if it can be split into two disjoint heaplets, one of which satisfies P and the other of which satisfies Q. So, the assertion $p{\mapsto}[l\colon i, r\colon j] * \mathsf{tree}(i) * \mathsf{tree}(j)$ describes a portion of heap with a pointer p that points to a record with l and r fields holding values i and j that themselves point to trees. The use of $*$ indicates that there is no overlap between p and i's tree and j's tree.

A question that often comes up is whether a pointer can go from one $*$-conjunct to another. The answer is yes. For instance, $p{\mapsto}[l\colon i, r\colon j]$ describes just a single cell, p, whose contents i and j point across $*$ into other heaplets in $p{\mapsto}[l\colon i, r\colon j] * \mathsf{tree}(i) * \mathsf{tree}(j)$. it helps to use a graphical intuition: take a directed graph, and then draw a line, partitioning it in two. Some of the links in the graph will go over the partition. The p to the left of $\mapsto$ corresponds to the sources of links to targets i and j.

There is a subtle point in the specification of `DispTree` that the reader might have noticed: In order to get the empty heap in the postcondition the precondition must say that "p points to a tree, *and there are no other cells in the given heaplet*". For, if there were other cells then you could not conclude `empty`, those cells that were not originally in the tree would still be around. This "no other cells" aspect is treated implicitly in separation logic. The `tree` predicate satisfies the recursive specification

$$\begin{aligned}\mathsf{tree}(E) \iff (E{=}\mathsf{nil} \wedge \ \ \mathtt{empty}) \\ \vee\ (\exists x, y.\, E{\mapsto}l\colon x, r\colon y * \mathsf{tree}(x) * \mathsf{tree}(y))\end{aligned}$$

where the use of `empty` when $E{=}\mathsf{nil}$ leads, inductively, to $\mathsf{tree}(E)$ not having additional cells.

Finally, there is a crucial interplay between the separating conjunction and a "tight" interpretation of Hoare triples [37,51,50]. A specification $\{P\}C\{Q\}$ means that C will (if it terminates) transform a heaplet satisfying P into one satisfying Q. It does this transformation in an in-place fashion, leaving the global heap surrounding the input heaplet unchanged. This in-place aspect can be seen clearly in the proof steps. For instance, for the first recursive call to `DispTree` the precondition is $p{\mapsto}[l\colon i, r\colon j] * \mathsf{tree}(i) * \mathsf{tree}(j)$, and this does not match up with the overall specification, which would expect only $\mathsf{tree}(i)$. What we do is use the

overall specification to replace $\texttt{tree}(i)$ by $\texttt{empty}$, obtaining $p{\mapsto}[l\colon i, r\colon j] * \texttt{empty} * \texttt{tree}(j)$, and then we can take one further step using the identity $\texttt{empty} * P \leftrightarrow P$.

These intuitions about heaplets and in-place update are codified in an inference rule, the frame rule

$$\frac{\{P\}\ C\ \{Q\}}{\{R * P\}\ C\ \{R * Q\}}\ \text{ModifiesOnly}(C) \cap \text{free}(R) = \emptyset$$

The R here is a frame axiom. The idea of this rule is that if C works on a portion of heap described by P, then it will not alter any additional heaplet described by R. There is also a side condition which has to do with named variables; e.g., the i and j in `DispTree`. (It is an embarassment that the heap is treated more cleanly than simple variables here, and we hope someday to get rid of the variable conditions altogether; see [10].)

As it is a relatively recent development, research on mechanized reasoning with separation logic is just beginning. The Smallfoot static assertion checker discovers proofs of lightweight shape specifications done using the logic [5]. And there are developing applications using interactive proof tools [32,49] and abstract interpretation [17,12,7,20,21].

4 Independence, Interference and Concurrency

Reasoning about concurrency is a subject that has received significant attention, and for good reason. The tremendous number of potential interactions between concurrent processes makes concurrent programs hard to grasp; a successful Program Verifier could provide considerable help to the concurrent programmer.

But, though it has received much attention, the difficulties that the theory meets on even simple examples are not as widely appreciated as perhaps they ought to be. To illustrate, I consider a very simple program: parallel mergesort.

$\{array(a, i, j)\}$
`procedure` $\texttt{ms}(a, i, j)$
`local` $m := (i+j)/2$;
`if` $i < j$ `then`
  $(\texttt{ms}(a, i, m) \parallel \texttt{ms}(a, m+1, j))$;
  $\texttt{merge}(a, i, m+1, j)$;
$\{sorted(a, i, j)\}$

For simplicity this specification just says that the final array is sorted, not that it is a permutation of the initial array.

Now, this program displays a trivial form of concurrency: *disjoint concurrency*. The recursive calls are completely independent, because they act on disjoint array segments. And yet, the program causes immediate difficulties for all of the best known proof methods.

Hoare had provided a beautiful rule for disjoint concurrency [23]

$$\frac{\{P\}C\{Q\} \qquad \{P'\}C'\{Q'\}}{\{P \wedge P'\}C \parallel C'\{Q \wedge Q'\}}$$

where C does not modify any variables free in P', C', Q', and conversely. Unfortunately, using this rule we cannot reason about the parallel calls in mergesort, because Hoare logic treats array-component assignment globally, where an assignment to $a[i]$ is viewed as an assignment to the entire array

$$\{P[(a \mid i\colon E)/a]\}\ a[i]{:=}\,E\ \{P\}$$

In this view the two parallel calls to `ms` are judged to be altering the *same* variable, a. So, the rule does not apply.

Cliff Jones has proposed a powerful approach to reasoning about concurrency, in his rely-guarantee formalism [27] (see also, [34]). For this example, we would add two conditions to the pre/post specification, formalizing the

- **Rely:** No other process touches my array segment $array(a, i, j)$; and
- **Guarantee:** I do not touch any storage outside my segment $array(a, i, j)$.

The Guarantee condition here is something like a frame axiom. The Rely, however, goes beyond the frame issue (one might fancifully consider it a kind of inverse frame axiom).

The point of this example is that it illustrates a breakdown of modularity. The guarantee condition (when formalized) talks about parts of the array not touched by a procedure call. In the worst case, this would have to be extended to other parts of memory than the single array given as a parameter. The issue is not just the cost for individual steps of reasoning, but rather that the rely and guarantee conditions, which are present to deal with subtle issues of interference, complicate the specification itself, even when no interference is present.

I have focussed on rely-guarantee here because is rightly lauded as providing a compositional approach to reasoning about concurrency. My point is that compositionality in program text does not guarantee locality in reasoning about resources such as program state: compositional reasoning can be extremely global. Also, I used a pre/post specification just because it is appropriate to the example, but the same modularity problem I have described here arises as well in temporal logics.

Because it is intuitively about separation, this example can be treated very easily in a concurrent extension of separation logic [39,11]. The crucial part of the proof is the following proof figure for the parallel composition.

$$\begin{array}{ccc} \{array(a,i,m) * \ array(a,m{+}1,j)\} & & \\ \{array(a,i,m)\} & & \{array(a,m{+}1,j)\} \\ \texttt{ms}(a,i,m) & \parallel & \texttt{ms}(a,m{+}1,j) \\ \{sorted(a,i,m)\} & & \{sorted(a,m{+}1,j)\} \\ \{sorted(a,i,m) \ * \ sorted(a,m{+}1,j)\} & & \end{array}$$

The use of the $*$ connective in $array(a,i,m) * \ array(a,m{+}1,j)\}$ implies that the array segments occupy separate memory, and we can then use a proof rule

$$\frac{\{P\}C\{Q\} \qquad \{P'\}C'\{Q'\}}{\{P * P'\}C \parallel C'\{Q * Q'\}}$$

that lets us reason independently about the two processes independently.

This rule is, of course, a descendent of Hoare's rule for disjoint concurrency. There are two reasons why we are able to treat this example where the original rule was not: (i) the assignment $a[i]{:=}e$ is not viewed by separation logic as an assignment to a, but rather to a single cell; (ii) $*$ can be used to describe partitioning of an array that is dynamic, depending on the program state.

My remarks on the rely-guarantee method should be taken in the right spirit: Indeed, they agree with a criticism of it lodged by Jones himself [28]. What he wants, and what I want, is a way to use complex methods where necessary to deal with interference when it is present, but to contain this complexity and default to simpler specification forms for interfaces between components that do not interfere with one another. The desire is to prevent *interference flooding*, where the mere possibility of interference complicates the specification notation, even in situations where there is a great degree of independence.

I do not claim that concurrent separation logic in its current state is the answer. It is good at specifying independence, but struggles with tightly-coupled, interfering processes. In contrast, rely-guarantee is good at describing interference, but is not well oriented to specifications of independent processes. Recently, there have been attempts to marry the advantages of concurrent separation logic and rely/guarantee [41,18]; these are perhaps further steps on the way to modular reasoning about (shared variable) concurrent processes.

5 Information Hiding

Pointers can wreak havoc with data abstraction. It is difficult to keep track of aliases, different copies of the same address, and so it is difficult to know when there are no pointers into the internals of a module. This problem has received attention in the object-oriented types community in work on ownership and confinement [13,3], stemming Hogg's colorful declaration "that objects provide encapsulation is the big lie of object-oriented programming [25]". Further difficulties, beyond confinement, are caused by low-level features such as address arithmetic and storage deallocation.

A good initial challenge which illustrated many issues is a resource management module, that provides primitives for allocating and deallocating resources which are held in a local free list. A client program should not alter the free list, except through the provided primitives; for example, the client should not tie a cycle in the free list. However, it is entirely possible for a client program to hold an alias to an element of the free list, after a deallocation operation is performed.

As an example, suppose that we have written our own memory manager, with operations $\texttt{alloc}(x)$ and $\texttt{free}(x)$ for allocating and deallocating records, where our implementation uses a free list in the usual way. A first attempt at specification might be something like

$$\{\texttt{allocated}(y) \wedge y.f = m \wedge \neg\texttt{allocated}(z)\}$$
$$\texttt{alloc}(x)$$
$$\{\texttt{allocated}(y) \wedge y.f = m \wedge \texttt{allocated}(x) \wedge y \neq x$$
$$\wedge(z \neq x \Rightarrow \neg\texttt{allocated}(z))\}$$

$$\{\texttt{allocated}(y) \wedge y.f = m \wedge \texttt{allocated}(x) \wedge y \neq x \wedge \neg\texttt{allocated}(z)\}$$
$$\texttt{free}(x)$$
$$\{\texttt{allocated}(y) \wedge y.f = m \wedge \neg\texttt{allocated}(x) \wedge y \neq x \wedge \neg\texttt{allocated}(z)\}$$

where, in addition to saying that x is allocated or deallocated, I have included a lot of frame axioms. I admit to some unease, I am not sure I have got the frame axioms exactly right (echoing the discussion from earlier), but there is a further problem I want to show, so let us assume that these are indeed the correct frame axioms. Here, I am again assuming that all variables other than x are auxiliary variables that are guaranteed not to be changed, and that $\{x\}$ is the entire modifies set of the specs (modifies for variables, not heap cells).

The further problem is that this specification does not stop a user of the memory manager from corrupting the free list, breaking the abstraction. For example, a sequence of statements

$$\texttt{alloc}(x)\,;\;\; \texttt{free}(x)\,; x{\rightarrow}r{:=}x$$

might tie a cycle in the free list, if the implementation uses the r field to point to the next record in the free list.

We can get around this problem by adding an invariant to the specifications. To each precondition and postcondition we add a predicate $\texttt{freelist}(free)$ saying that variable $free$ used by the manager points to a linked list without cycles, and where $\neg\texttt{allocated}(n)$ holds for each element in the list.

This fix, though, has come at great cost: we have exposed the invariant describing the ostensibly private storage of the memory management module. To see the cost, suppose a program makes use of n different modules. It would be unfortunate if we had to complicate specifications of user procedures by including descriptions of the internal resources of all modules that might be accessed. A change to a module's internal representation would necessitate altering the specifications of all other procedures that use it.

Stated plainly,

> *information hiding should be the bedrock of modular reasoning, but it is difficult to support soundly*

and this presents a great challenge for research in program logic.

This sort of example has been successfully treated in separation logic [38]. The details are much more involved than the earlier examples, and I will not give the proof here. The basic idea is that the $*$ connective allows the separation of the state owned by a client and the state owned by the manager (the free list). Crucially, since $*$ is a logical connective, the partition it describes can change

over time: in a sense, the logic tracks the right to dereference a cell transfers back and forth between client and module.

6 The Boogie Methodology and Relatives

Many of the issues touched on in this paper have also been approached in work on the "Boogie methodology" [30,36,4], and also in its precursors (see [29]). The basic idea of Boogie it to use certain auxiliary variables, such as ones to describe "ownership" of heap cells, to structure specifications and to constrain who can access what and when. Boogie builds on type systems for ownership [13,16], but uses assertions rather than types. Ownership gives a way to express a form of separation, and frame axioms are avoided by using general invariants which relate the states of auxiliary variables and the program state. The auxiliary variables allow fine control over when certain assertions, such as object invariants, must hold; this has allowed a novel approach to the old and vexing problem of object invariants for re-entrant modules (which allow implicit or explicit recursion).

I discussed an example similar to the first one in this paper with Peter Müller (we discussed copytree rather than disposetree). The early versions of Boogie could not handle that example due to inadequate framing properties, but a later version [31] could. Conversely, the earliest approach to information hiding using separation logic [38] could not handle re-entrant modules, but the later approach of [40] can. As shown in [8], the approach pioneered in [40] can be understood as using quantified predicates in a way that is analogous to the use of polymorphic typing to account for hiding of internal representation types [45,35]. On the other hand, Boogie has "pack" and "unpack" primitives which are intuitively similar to the corresponding primitives for existential types.

I just wanted to mention Boogie, to acknowledge (and point the reader to) the advances it and its relatives have made on difficult problems concerning modular reasoning about object-oriented programs. The exact relationship between Boogie and separation logic is not clear; there are similarities in intuition, but many differences in technique. The reader is referred to [29] for more information on this line of work, including work on ESC/Java and JML that I have not mentioned here.

7 Conclusion

In this paper I wanted to show some difficulties as regards modularity that traditional program logic has on even simple examples, and how it is not impossible to do much better, at least on those examples. In doing this I purposely started from programs rather than specifications; it is a good way to show where formalisms have difficulties. There are many other, more difficult, programs that can serve as challenging test cases.

Although I enjoy starting from programs, I would also love to be able to arrive at the kinds of program I considered by refinement, starting from a simple

specification. I just don't know how to do so. The refinement formalisms that I am aware of (VDM, B, etc) are based on a static form of modularity, where the state that a program component can change is listed in a fixed collection of variables, and the frame properties used are with respect to modifies clauses for these variables. This fixed modularity does not deal well when the partitions between the state used by program components is more dynamic, as is the case in parallel mergesort, in the resource manager example, and typically in systems programs. Of course, this last point should be taken as a challenge. It seems inconceivable that the modularity issues that separation logic and Boogie attempt to address should not show up as well on a design level. Furthermore, there are all sorts of dynamic, interconnected structures other than the program heap, those obtained from networks and message passing being prime examples. One might hope for a design formalism (say, an analogue of B or Z) that goes beyond static modularity, and that has the specific heap modularity of separation logic or Boogie as an instance.

Similar remarks apply to my focus on imperative programs, and shared-variable concurrency. A good problem would be to obtain a reasoning formalism for, say, the pi-calculus or for socket programs that displays the same sort of modularity in its account of channel usage as separation logic or Boogie does for the heap.

All of the examples in this paper have concerned safety properties. Recently, there has been progress on automatic proofs of liveness properties properties of software, using novel applications of abstract interpretation [42,15,6]. The problem of modular, or local, specifications and verifications of liveness properties of concurrent processes looms as an extremely difficult one; see [48,1] for important work in this direction.

Finally, one might question whether modular reasoning methods for software are in general even possible. In temporal logic there have been negative technical results [43], and we should be on the lookout for others. But, there has been considerable progress on modular reasoning about programs and this author, for one, plans to continue searching.

References

1. Amadi, M., Lamport, L.: Composing specifications. ACM TOPLAS 15(1), 73–132 (1993)
2. Ball, T., Cook, B., Levin, V., Rajamani, S.K.: SLAM and Static Driver Verifier: Technology Transfer of Formal Methods inside Microsoft. In: Boiten, E.A., Derrick, J., Smith, G.P. (eds.) IFM 2004. LNCS, vol. 2999, pp. 1–20. Springer, Heidelberg (2004)
3. Banerjee, A., Naumann, D.A.: Ownership confinement ensures representation independence for object-oriented programs. J.ACM (to appear, 2005)
4. Barnett, M., DeLine, R., Fahndrich, M., Leino, K.R.M., Schulte, W.: Verification of object-oriented programs with invariants. Journal of Object Technology 3(6), 27–56 (2004)
5. Berdine, J., Calcagno, C., O'Hearn, P.W.: Smallfoot: Automatic modular assertion checking with separation logic. In: 4th FMCO, pp. 115–137 (2006)

6. Berdine, J., Chawdhary, A., Cook, B., Distefano, D., O'Hearn, P.W.: Variance analyses from invariance analyses. In: 34th POPL, pp. 211–224 (2007)
7. Berdine, J., Cook, B., Distefano, D., O'Hearn, P.W., Wies, T., Yang, H.: Shape Analysis for Composite Data Structures. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 178–192. Springer, Heidelberg (2007)
8. Biering, B., Birkedal, L., Torp-Smith, N.: BI-hyperdoctrines, higher-order separation logic, and abstraction. ACM TOPLAS (to appear, 2007)
9. Borgida, A., Mylopoulos, J., Reiter, R.: On the frame problem in procedure specifications. IEEE Transactions of Software Engineering 21, 809–838 (1995)
10. Bornat, R., Calcagno, C., Yang, H.: Variables as resources in separation logic. In: 19th MFPS (2005)
11. Brookes, S.D.: A semantics for concurrent separation logic. In: Gardner, P., Yoshida, N. (eds.) CONCUR 2004. LNCS, vol. 3170, pp. 227–270. Springer, Heidelberg (2004)
12. Calcagno, C., Distefano, D., O'Hearn, P.W., Yang, H.: Beyond reachability: Shape abstraction in the presence of pointer arithmetic. In: Yi, K. (ed.) SAS 2006. LNCS, vol. 4134, pp. 182–203. Springer, Heidelberg (2006)
13. Clarke, D., Noble, J., Potter, J.: Simple ownership types for object containment. In: Knudsen, J.L. (ed.) ECOOP 2001. LNCS, vol. 2072, pp. 53–76. Springer, Heidelberg (2001)
14. Cok, D., Kiniry, J.: ESC/Java2: Uniting ESC/Java and JML. In: CASSIS, pp. 108–128 (2004)
15. Cook, B., Podelski, A., Rybalchenko, A.: Termination proofs for systems code. In: 13th PLDI (2006)
16. Dietl, W., Müller, P.: Universes: Lightweight ownership for JML. Journal of Object Technology (JOT) (to appear, 2005)
17. Distefano, D., O'Hearn, P., Yang, H.: A local shape analysis based on separation logic. In: 12th TACAS, pp. 287–302 (2006)
18. Feng, X., Ferreira, R., Shao, Z.: On the Relationship Between Concurrent Separation Logic and Assume-Guarantee Reasoning. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 173–188. Springer, Heidelberg (2007)
19. Flanagan, C., Leino, K.R.M., Lillibridge, M., Nelson, G., Saxe, J.B., Stata, R.: Extended static checking for Java. In: 9th PLDI (2002)
20. Gotsman, A., Berdine, J., Cook, B., Sagiv, M.: Thread-modular shape analysis. In: PLDI (to appear, 2007)
21. Guo, B., Vachharajani, N., August, D.: Shape analysis with inductive recursion synthesis. In: PLDI (to appear, 2007)
22. Hoare, C.A.R.: Procedures and parameters: An axiomatic approach. In: Engler, E. (ed.) Symposium on the Semantics of Algebraic Languages. Lecture Notes in Math. vol. 188, pp. 102–116. Springer, Heidelberg (1971)
23. Hoare, C.A.R.: Towards a theory of parallel programming. In: Hoare, Perrot (eds.) Operating Systems Techniques, Academic Press, London (1972)
24. Hoare, C.A.R.: The verifying compiler: A grand challenge for computing research. J. ACM 50(1), 63–69 (2003)
25. Hogg, J.: Islands: aliasing protection in object-oriented languages. In: 6th OOPSLA (1991)
26. Isthiaq, S., O'Hearn, P.W.: BI as an assertion language for mutable data structures. In: 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, London, January 2001, pp. 36–49 (2001)

27. Jones, C.B.: Specification and design of (parallel) programs. In: IFIP Conference (1983)
28. Jones, C.B.: Wanted: A compositional approach to concurrency. In: McIver, A., Morgan, C. (eds.) Programming Methodology, pp. 1–15. Springer, Heidelberg (2003)
29. Leavens, G.T., Leino, K.R.M., Müller, P.: Specification and verification challenges for sequential object-oriented programs. In: Formal Aspects of Computing (to appear, 2007)
30. Leino, K.R.M., Müller, P.: Object Invariants in Dynamic Contexts. In: Odersky, M. (ed.) ECOOP 2004. LNCS, vol. 3086, pp. 491–515. Springer, Heidelberg (2004)
31. Leino, K.R.M., Müller, P.: A Verification Methodology for Model Fields. In: Sestoft, P. (ed.) ESOP 2006 and ETAPS 2006. LNCS, vol. 3924, pp. 115–130. Springer, Heidelberg (2006)
32. Marti, N., Affeldt, R., Yonezawa, A.: Verification of the heap manager of an operating system using separation logic. In: Proceedings of the 3rd SPACE Workshop, Charleston (2006)
33. McCarthy, J., Hayes, P.: Some philosophical problems from the standpoint of artificial intelligence. In: Machine Intelligence, vol. 4, pp. 463–502 (1969)
34. Misra, J., Chandy, K.M.: Proofs of networks of processes. IEEE Trans. Software Eng. 7(4), 417–426 (1981)
35. Mitchell, J.C., Plotkin, G.D.: Abstract types have existential types. ACM Trans. Programming Languages and Systems 10(3), 470–502 (1988)
36. Naumann, D.A., Barnett, M.: Towards imperative modules: Reasoning about invariants and sharing of mutable state. In: 19th LICS, pp. 313–323 (2004)
37. O'Hearn, P., Reynolds, J., Yang, H.: Local reasoning about programs that alter data structures. In: Proceedings of 15th Annual Conference of the European Association for Computer Science Logic. LNCS, pp. 1–19. Springer, Heidelberg (2001)
38. O'Hearn, P.W., Yang, H., Reynolds, J.C.: Separation and information hiding. In: 31st POPL, pp. 268–280 (2004)
39. O'Hearn, P.W.: Resources, Concurrency and Local Reasoning. In: Gardner, P., Yoshida, N. (eds.) CONCUR 2004. LNCS, vol. 3170, pp. 49–67. Springer, Heidelberg (2004)
40. Parkinson, M., Bierman, G.: Separation logic and abstraction. In: Proceedings of POPL (2005)
41. Parkinson, M., Vafeiadis, V.: A Marriage of Rely/Guarantee and Separation Logic. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR. LNCS, vol. 4703, pp. 256–271. Springer, Heidelberg (2007)
42. Podelski, A., Rybalchenko, A.: Transition invariants. In: 19th LICS (2004)
43. Rabinovich, A.: On compositionality and its limitations. ACM TOCL 8(1), 73–132 (2007)
44. Reiter, R.: The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression. In: Lifschitz, V. (ed.) Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy, pp. 359–380. Academic Press, London (1991)
45. Reynolds, J.C.: Types, abstraction and parametric polymorphism. In: Proceedings of IFIP (1983)
46. Reynolds, J.C.: Intuitionistic reasoning about shared mutable data structure. In: Davies, J., Roscoe, B., Woodcock, J. (eds.) Millennial Perspectives in Computer Science, Houndsmill, Hampshire, Palgrave, pp. 303–321 (2000)

47. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: 17th LICS, pp. 55–74 (2002)
48. Stark, E.W.: A proof technique for rely/guarantee properties. In: Maheshwari, S.N. (ed.) FSTTCS 1985. LNCS, vol. 206, pp. 369–391. Springer, Heidelberg (1985)
49. Tuch, H., Klein, G., Norrish, M.: Types, bytes, and separation logic. In: 34th POPL (2007)
50. Yang, H.: Local Reasoning for Stateful Programs. Ph.D. thesis, University of Illinois, Urbana-Champaign (2001)
51. Yang, H., O'Hearn, P.W.: A Semantic Basis for Local Reasoning. In: Nielsen, M., Engberg, U. (eds.) ETAPS 2002 and FOSSACS 2002. LNCS, vol. 2303, Springer, Heidelberg (2002)

A Discussion on Peter O'Hearn's Presentation

Willem-Paul de Roever

You presented two examples explaining why concurrency and pointer manipulation are complicated. Actually, you needed only one example, the first one, because if you look at its scheme, it is part of the concurrent garbage collector of Dijkstra, Lamport and Scholten. And if you have concurrently with this a so-called mutator program, which changes the links, this marking strategy, which you give, is wrong. This was the famous error of these three persons in 1976. This is what you just took, that is the famous error.

Peter O'Hearn

No, the point I am making is that the specifications will be far more complicated. I was not trying to point out an error in a program. Perhaps I misunderstood what you said.

Willem-Paul de Roever: No, you used two examples...

Peter O'Hearn: Yes.

Willem-Paul de Roever: ...you could have used one example.

Peter O'Hearn: Oh yes, I could use a specially prepared example...

Willem-Paul de Roever: No, no, this one!

Peter O'Hearn:

I could have used a parallel-disposed tree, and that would have shown both of my points. But the other point, I wanted to make is, even with our pointers, even with simpler examples than concurrency, we still have more complex specifications than we would like.

Willem-Paul de Roever: OK.

Peter O'Hearn: Oh, but your point is taken.

Willem-Paul de Roever: Yes, OK.

Peter Schmitt

I want to come back to the frame problem. A method that turned out to work pretty well for us is to use this modifies or assignable clause from JML. So, we specify: It is only these elements, only these expressions that might change. But this [example] needs an additional twist, because here, you refer to an unbounded number of elements. But we have these star notations or reachable notations, and, including this, it turned out that it worked pretty well.

Peter O'Hearn

I believe that it works. Many things have been tried that worked pretty well. But the problem here is that the things that are changed are not named by a fixed, finite number of program identifiers, and so the simple approach to modifies clauses would not work for an example like this. There might be more subtle approaches, but I would repeat one thing: that the frame problem was set down 35 years ago or so. And I think, there is still no general solution, and so... I would love to see any approaches people have to solve these problems apart from separation logic!

Peter Schmitt

My point is, we do not need a general solution, we need a solution here in programming logic context, and you are right, we need a more subtle mechanism. So we need a way to describe: We want to address all locations that are reachable from some node by all the L- and R-operations. But the machinery is there to do this.

Peter O'Hearn

Is there to do what? Is there to specify the frame axioms? The problem is not to have to write the frame axioms, yes?

Peter Schmitt

Oh, what you have to do is to specify in the assignable clauses all the elements that get at most changed. That is what you have to do.

Egon Börger, University of Pisa

I think, you have a problem with the approach. Let us look at what people do in mathematics. They never would complicate a definition for the reason that during the attempt to prove something by induction, you need a stronger hypothesis at the inductive step. You see? So, separating definitions-that means specifications and what you need for the proofs-I think this is really crucial for being able to do challenging proofs of really relevant properties. Now, I know that this is heretical in this community for many of my colleagues, but if you mix up these two things, you will always be in full trouble.

Peter O'Hearn

I understand what you are saying. To repeat what he is saying: We might want a simpler spec for use, and a more complicated spec to have a strong enough induction hypothesis. But the problem is, the simple spec won't be usable at very many call sites, even not just these two call sites. So, I agree with what you say, but I don't agree that this impacts this problem.

Rustan Leino

I wondered what you thought of the abstraction dependencies or the data groups or the Boogie methodology for handling the frame problem.

Peter O'Hearn

I do not know the extent to which it solves it. I have seen some approaches based on type systems [that] I think definitely do now solve it.

Rustan Leino

None of which I said uses type systems. Each one of the three has been used in verification. We use the Boogie methodology right now, for example, in Spec#.

Peter O'Hearn

Yes, I am hoping that someone can explain that methodology to me, and we can go for a few examples to see if it handles, for instance, examples like dispose tree. Simple ones like that.

Wolfgang Paul

I am slightly confused, as often in my life. I have a very simple question. I have seen you had some simple programs and you have presented certain proofs, and although they were not completely trivial, I failed to see why those proofs were complicated. There are certain things you want to prove. You have to prove that things that change, change in the right way, and on the way, you have to prove that certain things do not change in undesired ways. It is completely normal.

You observe this, write down the right tools, and everything you wrote down was very nice and beautiful and the right things, but I would not call it complicated! So, as a consequence, if we call things complicated that for Russian scientists certainly are not, then we are limiting the things which we can do.

Peter O'Hearn

I only showed you the proofs that were in the formalism that at least partially solve those problems.

David Naumann

To put a slightly different twist on it... I think Wolfgang is pointing out that in large systems, there is all sorts of complexity in the interactions. And the ghost variable approach that we have been talking about a little this morning, and Rustan was just alluding to as well, is somewhat ad-hoc, but one can look for reasoning patterns particular to situations and try to formalize the dependencies there and be able to make sense of the footprints of various predicates and reason about their interferences or risks thereof or absence thereof. By contrast, separation logic has complete separation of the footprint of predicates expressed by a logical connective, and then the footprint per se does not exist. That is sort of the interpretation of these triples. It's gorgeous for small algorithms, but it poses the questions of: Will this scale to more complex interactions and overlaps between resources? And will there be further connectives needed for such?

Peter O'Hearn

I don't know. I mean, I was struck by you using the "star of heaps", because what is behind separation logic, the reason it works so well on the small examples, is the local way that programs operate, which guarantees that many frame axioms are simply true. And you might be able to, in a traditional logic, make use of that same observation. So, it is just that separation logic gives you a convenient way to do it, but there is a deeper reason for why it is working well.

Wolfgang Paul

Let me insist! Maybe this is controversial and I do not want to hurt anybody, but: I always hear "reasoning patterns". If I want to teach mathematics to somebody, I do not teach him reasoning patterns, I teach him how to find proofs. When I teach a class of retarded children in lower classes to prepare for a school examination, then I teach them reasoning patterns, but among scientists, for heaven's sake! I do not try to identify reasoning patterns; I think that is harmful!

Peter O'Hearn: To identify reasoning patterns?

Wolfgang Paul

Reasoning patterns might be good if I want to automate things, and then I say: With what reasoning pattern can I automatically kill the following problem? Then, it is great, but just for finding proofs... what's the problem? In mathematics, it is not done this way, and this is mathematics!

Peter O'Hearn

No, mathematics also, we are allowed to be scientists, and we can insist that we do not like a solution because it is complicated, and we can try for a simpler solution. And so, that is why I am insisting.