

Array Programming in Whiley

David J. Pearce

School of Engineering and Computer Science, Victoria University of Wellington, New Zealand
djp@ecs.vuw.ac.nz

Abstract

Arrays are a fundamental mechanism for developing and reasoning about programs. Using them, one can easily encode a range of important algorithms from various domains, such as for sorting, graph traversal, heap manipulation and more. However, the encoding of such problems in traditional languages is relatively opaque. That is, such programming languages do not allow those properties important for the given problem to be encoded within the language itself and, instead, rely up on programmer-supplied comments.

This paper explores how array-based programming is enhanced by programming languages which support specifications and invariants over arrays. Examples of such systems include Dafny, Why3, Whiley, Spec# and more. For this paper, we choose Whiley as this language provides good support for array-based programming. Whiley is a programming language designed for verification and employs a *verifying compiler* to ensure that programs meet their specifications. A number of features make Whiley particularly suitable for array-based programming, such as type invariants and abstract properties. We explore this through a series of worked examples.

CCS Concepts • **Theory of computation** → **Program reasoning**; • **Software and its engineering** → *Imperative languages*; *Data types and structures*

Keywords Array Programming, Software Verification, Loop Invariants

1. Introduction

Arrays provide a powerful and oft-overlooked primitive for use in programming. Arrays can be used to describe fundamental algorithms from a wide range of areas and disciplines. Sorting algorithms are, of course, one such example. With multi-dimensional arrays and integer types we can easily encode more complex data structures, such as *sets*, *trees* and *graphs*, etc. Even *linked structures*, typically implemented with references, can be implemented using arrays (i.e. where array indices replace references, etc). Unfortunately, modern languages make array-based programming feel very “low level” and, instead, promote the use of linked structures by making `struct` and/or `class` constructs feel more “sophisticated”. The essential argument of this paper is that, with proper programming language support, array-based programming can be made to feel just as sophisticated.

Numerous programming languages have been developed with first-class support for specification in the form of pre- and post-conditions and which typically emphasise the use of automated theorem provers, such as Simplify [1] or Z3 [2], for static checking. Good examples include ESC/Java [3], Spec# [4], Dafny [5], Why3 [6], VeriFast [7], SPARK/Ada [8], and Whiley [9–11]. The development of such tools has proved something of a boon for array-based programming. This is because the specification languages they employ, as well as the underlying automated theorem provers, typically have limited support for dealing with linked data structures [12]. To work-around this, tool developers and practitioners have placed great emphasis on the use of arrays for encoding problems of interest.

Contribution. This paper explores how array-based programming is enhanced by programming languages which support specifications and invariants over arrays. This is achieved through a series of worked examples in the Whiley programming language.

2. Background

In this section, we introduce the Whiley language in the context of software verification through a series of examples. We do not provide an exhaustive examination of the language and the interested reader may find more detailed introductions elsewhere [11].

2.1 Overview

The Whiley programming language has been developed from the ground up to enable compile-time verification of programs [11]. The Whiley Compiler (WyC) attempts to ensure that every function in a program meets its specification. When it succeeds in this endeavour, we know that: 1) all function post-conditions are met (assuming their pre-conditions held on entry); 2) all invocations meet their respective function’s pre-condition; 3) runtime errors such as divide-by-zero, out-of-bounds accesses and null-pointer dereferences are impossible. Note, however, such programs may still loop indefinitely and/or exhaust available resources (e.g. RAM).

2.2 Example 1 — Preconditions and Postconditions

Whiley allows explicit *pre-* and *post-conditions* to be given for functions. For example, the following function accepts a positive integer and returns a natural number:

```
function decrement (int x) -> (int y)
// Parameter x must be greater than zero
requires x > 0
// Return must be greater or equal to zero
ensures y >= 0:
//
return x - 1
```

Here, `decrement()` includes **requires** and **ensures** clauses which correspond (respectively) to its *precondition* and

postcondition. In this context, `y` represents the return value and may be used only within the `ensures` clause. The Whiley compiler statically verifies this function meets its specification (note, integers are unbounded and cannot underflow).

The Whiley compiler reasons about functions by exploring their control-flow paths. As it learns more about the variables encountered, it takes this into account. For example:

```
function max(int x, int y) -> (int z)
// Must return either x or y
ensures x == z || y == z
// Return must be as large as x and y
ensures x <= z && y <= z:
//
  if x > y:
    return x
  else:
    return y
```

Here, multiple `ensures` clauses are given which are conjoined to form the function’s postcondition. We find that allowing multiple `ensures` clauses helps readability, and note that JML [13], Spec# [4] and Dafny [5] also permit this. Furthermore, multiple `requires` clauses are permitted in the same manner.

As an aside, we note that the body of the `max()` above is almost completely determined by its specification. However, in general, this is not usually the case and typically there is scope for significant variation between implementations.

2.3 Example 2 — Data Type Invariants

Type invariants over data can also be explicitly defined:

```
// A natural number is an integer greater-than-or-equal-to zero
type nat is (int n) where n >= 0
// A positive number is an integer greater-than zero
type pos is (int p) where p > 0
```

Here, the `type` declaration includes a `where` clause constraining the permitted values. The declared variable (e.g., `n` or `p`) represents an arbitrary value of the given type. Thus, `nat` defines the type of natural numbers. Likewise, `pos` gives the type of positive integers. Constrained types are helpful for ensuring specifications remain as readable as possible. For example, we can update `decrement()` as follows:

```
function decrement(pos x) -> (nat n):
//
  return x - 1
```

Types in Whiley are more fluid than in typical languages as variables can move seamlessly between them. If two types T_1 and T_2 have the same *underlying* type, then T_1 is a subtype of T_2 iff the constraint on T_1 implies that of T_2 . For example, `pos` above is a subtype of `nat` as the constraint on `pos` implies that of `nat`. As another interesting example, consider the following:

```
type anat is (int x) where x >= 0
type bnat is (int x) where 2*x >= x

function f(anat x) -> bnat:
  return x
```

In this case, we have two alternate (and completely equivalent) definitions for a natural number (we can see that `bnat` is equivalent to `anat` by subtracting `x` from both sides).

2.4 Example 3 — Loop Invariants

Whiley supports loop invariants which are necessary for proving properties about loops. The following illustrates:

```
function sum(int[] xs) -> (int r)
// Every item in xs is greater or equal to zero
requires all { i in 0..|xs| | xs[i] >= 0 }
// Return must be greater or equal to zero
ensures r >= 0:
//
  int s = 0
  int i = 0
  while i < |xs| where s >= 0 && i >= 0:
    s = s + xs[i]
    i = i + 1
  return s
```

Here, a bounded quantifier enforces that `sum()` accepts an array of natural numbers (which could equally have been expressed as type `nat[]`). A key constraint is that summing an array of natural numbers yields a natural number (recall arithmetic is unbounded and does not overflow). The Whiley compiler statically verifies that `sum()` does meet its specification. The loop invariant is necessary to help the compiler generate a sufficiently powerful verification condition to prove the function meets the post condition.

2.5 Example 4 — Properties

Properties are provided as a convenient mechanism for expressing common parts of a specification/invariant:

```
property contains(int[] xs, int x, int n)
where some { k in 0..n | xs[k] == x }
```

Here, the property `contains` captures the notion that an item is contained in an array. Properties provide a mechanism for defining the “language” in which a function’s specification is written.¹

3. Case Studies

We now examine array programming in the context of a language supporting specification and invariants. Whilst our presentation does focus on the use of Whiley, it should be noted that it applies equally to other similar languages (e.g. Dafny [5], Spec# [4], etc).

3.1 Maximum Element

Our first example is the well-known problem of finding the maximum element of an array. This is an oft-used example in the context of verification [14, 15]. The problem definition is simple:

“Given an array, return the largest value contained therein.”

To begin the process, we must first state this more precisely as a specification. To do this, we define the concept of “largest” as follows:

```
property largest(int[] xs, int x, int n)
where all { k in 0..n | xs[k] <= x }
```

This states that no element in the array from zero upto (but not including) `n` is larger than `x` (the reason for including the upper bound `n` will be apparent shortly). This is not itself sufficient to

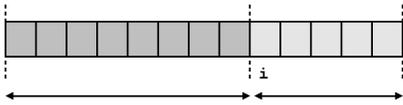
¹One may wonder why properties are needed at all. However, properties have special treatment within the verifier, compared with functions which are “uninterpreted”.

specify our problem but, using `contains()` from above, we can now do so:

```
function max(int[] items) -> (int r)
// Input array must have at least one item
requires |items| > 0
// Item returned must be largest of any in array
ensures largest(items, r, |items|)
// Item returned must be contained in array
ensures contains(items, r, |items|):
```

We can immediately see how the language facilitates a more expressive description of the function, compared with mainstream languages. In particular, the use of properties and quantification over arrays is critical here.

At this point, it remains to implement `max()` and we simply take the “most obvious” approach. Key to this is a loop invariant which, at each point, maintains the largest value seen. This can be viewed diagrammatically as follows:



This represents our implementation operating on a given array which, at any given point, has determined the maximum element (as determined by our specification) for all elements upto (but not including) `i`.

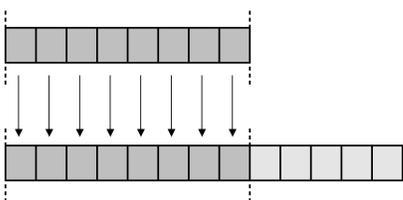
Our implementation of the `max()` function is as follows (where the specification is omitted as it matches above):

```
function max(int[] items) -> (int r)
...
  nat i = 1
  int m = items[0]
  //
  while i < |items|
  where i <= |items|
  where largest(items, m, i)
  where contains(items, m, i):
    if items[i] > m:
      m = items[i]
    i = i + 1
  //
  return m
```

The reason that both `contains()` and `largest()` were defined with an upper bound `n` is now apparent. Doing so allows us to reuse them for both specification and loop invariant. The latter is necessary for the Whiley compiler to statically verify that this function meets its specification.

3.2 Resize

Our next example is a simple function for resizing an array. If the array size is reduced, all items up to that point are unchanged. If the array size increases, the new portion of the array is filled with a default. The following illustrates:



To give the specification for `resize()`, we define the notion of “unchanged” as follows:

```
property unchanged(int[] xs, int n, int[] ys)
// All elements upto n unchanged
where all { j in 0..|xs| |
  j < n ==> xs[j] == ys[j]
}
```

This states that everything up to a given size `n` is the same between `xs` and `ys`. Using this, we can specify `resize()`:

```
function resize(int[] xs, nat n, int v)
-> (int[] ys)
// Returned array is of specified size n
ensures |ys| == n
// All elements retained from old array (up to new n)
ensures unchanged(xs, n, ys)
// All new elements match default value
ensures all { i in |xs|..n | ys[i] == v}:
```

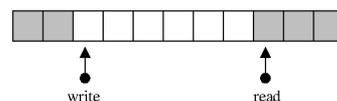
The specification restricts the possible implementations considerably but still leaves open, for example, the order of iteration through the arrays. As before, our implementation takes the straightforward approach:

```
function resize(int[] xs, nat n, int v)
...
  int[] rs = [v; n]
  nat i = 0
  while i < n && i < |xs|
  // Ensure array size remains unchanged
  where |rs| == n
  // All elements up to i match as before
  where unchanged(xs, i, rs)
  // All elements about n match element
  where all { j in |xs|..n | rs[j] == v}:
    //
    rs[i] = xs[i]
    i = i + 1
  //
  return rs
```

One curious aspect of our implementation is the loop invariant clause “`|rs| == n`”. This clause seems unnecessary as no statement in the loop affects the size of `rs`. The need for it is an artefact of the verification system (which is based on Hoare logic). Specifically, within the body of a loop all knowledge about variables modified in that loop is lost, except that given in the loop invariant and condition.

3.3 Cyclic Buffer

A *cyclic buffer* consists of a fixed-sized array with *read* and *write* pointers that “wrap around”. The key is that we should be able to write when the buffer is not full, and read when it is not empty. The following gives a diagrammatic view:



Here, we see a buffer holding five items (marked in gray) where `write` has wrapped around and is below `read`. We can also see that `write` identifies the first *unused* space, whilst `read`

identifies the first *used* space (if one exists). The buffer is *empty* when “read == write” and *full* when there is exactly one unused space (i.e. roughly when “write+1 == read”). We can encode the general concept of a cyclic buffer as follows:

```

type Buffer is {
  int[] data,
  nat read, // read pointer
  nat write // write pointer
}
// Read / write pointers within bounds
where read < |data| && write < |data|

```

Our definition above is not very restrictive and simply requires both pointers are within bounds. We can now easily define the notion of an empty buffer:

```

type EmptyBuffer is (Buffer b)
where b.read == b.write

```

The type EmptyBuffer further constrains Buffer to the case where both pointers are the same. Using this we can give a suitable constructor for Buffer:

```

function Buffer(int n) -> EmptyBuffer
// Cannot create buffer with zero size!
requires n > 0:
//
return {data: [0;n], read: 0, write: 0}

```

Here, “[v;n]” generates an array of size n with each element initialised to v. We see a precondition is necessary on the constructor to ensure “n > 0” as, otherwise, one cannot construct a valid buffer (i.e. since this requires e.g. “0 <= read < n”). We now proceed to define the concept of a “non-full” buffer as follows:

```

type NonFull is (Buffer b)
// Write cannot be immediately behind read
where ((b.write+1) % |b.data|) != b.read

```

This simply states that write cannot be one space behind read (accounting for wrap around). Using this definition, we can give the function for writing into the buffer:

```

function write(NonFull b,int v)->(Buffer r)
// All items unchanged except at (old) write position
ensures unchanged(b.data, b.write, r.data)
// Item v now at old write pos
ensures r.data[b.write] == v
// Read pointer is unchanged
ensures b.read == r.read
// Write pointer has advanced one position
ensures (b.write+1) % |b.data| == r.write:
//
  b.data[b.write] = v
  b.write = (b.write + 1) % |b.data|
return b

```

The post-condition here is perhaps surprisingly involved and, indeed, is longer than the implementation itself! In fact, it provides a fairly exact description of what the function does (though, in general, we find that specifications leave some “wriggle” room for the implementation). The property unchanged() simply states

that everything in the array is unchanged, except for the item being written:

```

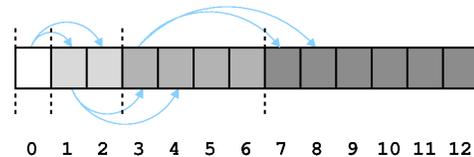
property unchanged(int[] xs,int i,int[] ys)
// Size of arrays remains unchanged
where |xs| == |ys|
// All elements except i are unchanged
where all {
  j in 0..|xs| | (j != i) ==> (xs[j]==ys[j])
}

```

At this point, we leave the remainder of the implementation to the reader’s imagination and simply remark that it follows from the above as expected.

3.4 Binary Heap

The binary heap is a common data-structure for implementing a priority queue. From the perspective of array programming, it is convenient as it implemented using an array with relative ease. Diagrammatically, we can view it as thus:



The key is that, for a given node at index i, the position of its children is calculated as “(2*i)+1” and “(2*i)+2”. Furthermore, the value of each element is larger (i.e. has higher priority) than either of its children. We can formalise this quite nicely as follows:

```

property valid(int[] H, int n, int i,int c)
// For valid child, value is less than parent
where (2*i)+c < n ==> H[(2*i)+c] < H[i]

property validChildren(int[] H, int n)
// Items on left branch are below their parent's item
where all{ i in 0..n | valid(H,n,i,1) }
// Items on right branch are below their parent's item
where all{ i in 0..n | valid(H,n,i,2) }

```

Property validChildren() gives the invariant which holds for any valid heap where n represents the number of items currently in the heap (and for which we will additionally require n <= |H|). Using this property we can define a suitable data type for representing heaps:

```

type Heap is {int[] data, nat length}
// All children within the heap are valid
where validChildren(data,length)
// Never more allocated items than space
where length <= |data|

```

Thus, a Heap defines a region of data within which all items are stored and which limits the heap’s maximum possible size. At any given moment *some*, *all* or *none* of the items in that region may be contained in the heap proper (as determined by length).

For completeness, we now consider one operation for manipulating heaps. Figure 1 gives the insert() function. This puts the item v being inserted into the next available slot, which may break the heap invariant (i.e. because v is larger than its parent). To restore the invariant, the function proceeds to swap v up the tree until it find the correct spot.

```

function insert(Heap h, int v) -> (Heap r)
// Heap cannot be full
requires h.length < |h.data|
// Item v is added to heap
ensures contains(r.data, v, r.length)
// All items in heap remain in heap
ensures containsAll(r.data, h.data)
// All items now in heap were in heap (except perhaps v)
ensures containsEx(h.data, r.data, v) :
//
  nat i = h.length
  // Add v to end of heap
  h.data[i] = v
  // Create (ghost) copy of heap
  Heap oh = h
  //
  int parent = (i-1)/2
  // Swap v up to restore invariant
  while parent > 0 && h.data[parent] < v
  where parent < h.length
  where contains(h.data, v, |h.data|)
  where containsAll(h.data, oh.data)
  where containsEx(oh.data, h.data, v) :
    // perform a swap
    h.data[i] = h.data[parent]
    h.data[parent] = v
    // update indices
    i = parent
    parent = (i-1)/2
  //
  return h

property containsAll(int[] a, int[] b)
// Every item in b is in a
where all { i in 0..|b| |
  contains(a, b[i], |b|)
}
property containsEx(int[] a, int[] b, int v)
// Every item in b is in a, except perhaps v
where all { i in 0..|b| |
  b[i] == v || contains(a, b[i], |b|)
}

```

Figure 1. The function for inserting an item into a heap

One interesting aspect of `insert()` is the use of variable “oh”. This is commonly referred to as a *ghost variable*, since its only purpose is to aid verification [3, 16, 17]. Using this the loop invariant can refer to `h` as it was before the loop. Finally, `insert()` also employs two properties which extend the property `contains()` from before.

4. Discussion

The above case studies highlight how the use of specifications gives life to simple array-based implementations. That is, the ability to express invariants explicitly can offset the “low-level” nature of the array implementation. Such invariants are typically otherwise hidden (though one can optimistically hope they are at least mentioned in the docs). We now provide some further commentary arising from this.

Properties and Type Invariants. Our case studies make extensive use of properties to ensure specifications are as readable as possible. Typically, we also use them as the language for defining the problem. We can observe some reuse of properties between examples. For example, `contains()` is used in the `max()` and binary heap examples. Likewise, a similar notion of `unchanged()` is used in `resize()` and the cyclic buffer. We imagine common properties will be combined into libraries for describing related problems. And, most likely, that a core set of properties will emerge as a primary language for describing array problems (which, ideally, would be incorporated into the standard library).

Type invariants are another mechanism we employ to help characterise our problems. In the cyclic buffer problem, defining different states of the buffer (e.g. `EmptyBuffer`, `NonFull`, etc) proved useful. We expect type invariants to be reused less than properties as, by their nature, they are somehow more “concrete” (i.e. they range over fixed types).

Other Tools. Whilst this paper focuses primarily on Whiley, the argument put forward (i.e. that specifications/invariants enhance array programming) extends to other similar tools. We now consider one such tool, namely Dafny, which is quite comparable to Whiley [5, 16]. Dafny is an imperative language with support for objects and classes without inheritance. Like Whiley, Dafny employs unbound arithmetic, pure functions and provides immutable collection types with value semantics. Dafny also employs *dynamic frames* [18] as a simple mechanism for reasoning about pointer-based programs. Finally, Dafny has been used successfully in many verification challenges [15, 19–22] and for reasoning about pointer-based programs such as the Schorr-Waite algorithm for marking reachable nodes in an object graph [5]

Figure 2 provides a Dafny implementation of the `max()` function from §3.1. Generally speaking, this is largely similar to the Whiley implementation. However, there are some differences between the two. For example, Dafny treats arrays as references which can be `null` and, hence, requires it be established they are *not*. Likewise, `reads` clauses are required to declare any reference variables that may be accessed within a method. Finally, predicate methods (analogous to properties in Whiley) must declare preconditions to ensure, for example, arrays are not null and array accesses are within bounds.

Loop Invariants. The need for loop invariants remains something of a barrier-to-adoption for tools like Whiley and Dafny. Flanagan and Qadeer commented that “*While method specifications also function as useful documentation and may be helpful for code maintenance, loop invariants do not provide comparable benefits*” [23]. Their experiences using ESC/Java lead them believe that “*the burden of specifying loop invariants is substantial*”. Likewise, Beckert *et al.* note that, in general, “*loop invariants are polluted by formulas stating what the loop does not do.*” [24].

Numerous techniques exist for automatically inferring loop invariants and we hope such work will ease the burden [25–30]. Abstract interpretation over some mathematical domain (e.g. linear inequalities [31], polynomials [32] or convex polyhedra [33]) is one common approach and, for example, Dafny employs some techniques here already. This is particularly effective at identifying arithmetic relationships between variables, but not so much for generating quantified invariants (i.e. as needed for array-manipulating programs [28]). Predicate abstraction is another well-studied technique from model checking which has been adopted for inferring loop invariants by several researchers [23, 34–36]. Predicate abstraction works by exploring all permutations of a predefined set of predicates to (hopefully) identify one or more true invariants. The limitation here is the need for a predefined set of predicates, as this restricts what invariants can be inferred. Another common ap-

```

predicate method
largest(xs: array<int>, x: int, n: int)
requires xs != null && n <= xs.Length;
reads xs;{
  (forall k :: 0 <= k < n ==> xs[k] <= x)
}
predicate method
contains(xs: array<int>, x: int, n: int)
requires xs != null && n <= xs.Length;
reads xs;{
  (exists k :: 0 <= k < n && xs[k] == x)
}

method max(items:array<int>) returns (r:int)
// Input array must have at least one item
requires items != null && items.Length > 0;
// Item returned must be largest of any in array
ensures largest(items,r,items.Length);
// Item returned must be contained in array
ensures contains(items,r,items.Length);
{
  r := items[0];
  var i := 1;

  while i < items.Length
  invariant i <= items.Length;
  invariant largest(items,r,i);
  invariant contains(items,r,i);
  {
    if items[i] > r { r := items[i]; }
    i := i + 1;
  }
}

```

Figure 2. Dafny implementation of the `max()` function from §3.1.

proach is to exploit the connection between a loop’s invariant and its postcondition (assuming this is known) [37]. The strategy here is to progressively weaken a loop’s postcondition yielding one or more candidate invariants. For example, Furia and Meyer do this by exhaustively examining all mutations of the postcondition [28].

5. Related Work

Verification. Small array-oriented programs have long been used in the context of verification. Dijkstra’s *Dutch National Flag* problem, although somewhat contrived, is perhaps one of the most widely used for introducing loop invariants [38]. The essential problem is to sort an array of three colours (i.e. those making up the Dutch National Flag) into the correct order. The solution requires only a single loop with a rather nice loop invariant. The introductory books of Broda *et al.* [39], Kourie and Watson [14] and Backhouse [40] all provide worked solutions to this problem to illustrate loop invariants, whilst Gries leaves it as an exercise for the reader [37]. Numerous verified solutions have also been given [22, 41, 42].

The *Correctness-by-Construction (CbC)* approach pioneered by Dijkstra, Hoare and others promotes the development of programs by progressive refinement from specifications [43, 44]. Kourie and Watson applied this to a range of simple array-based programs, such as finding the maximum element of an array, performing a linear search through an array, finding the longest matching array segment, etc [14]. Dony and Le Charlier developed a tool for

teaching students the CbC method [45]. Their stated ambition was to restrict themselves to “*a very simple programming language with simple types (limited to finite domains) and arrays, and to address (relatively) small examples such as searching and sorting algorithms*”. This, they claimed, allowed problems to be specified clearly and formally.

In works on verification, especially those related to teaching, the presentation of loop invariants almost always coincides with the use of array-based examples (which is perhaps not entirely surprising). Back emphasises the value of “*drawing good figures to illustrate the way the algorithm is intended to work*” where his examples primarily correspond with array diagrams similar to those seen in this (and many other) papers [46]. He goes on to claim clear benefit from focusing on “*well understood application domains (mostly array manipulation programs)*” when teaching. In a similar vein, Astrachan argues that “*invariants are especially useful in introductory courses*” but are often avoided because of the mathematical notation involved [47]. Instead, a preference in this context is for the use of invariant diagrams which, again, are typically (though not exclusively) over arrays.

Finally, we observe that a large portion of the problems set in verification challenges are array problems of the kind illustrated here. For example, in VSCOMP’10 three out of five were array problems (with the remainder on linked lists) [20]. Likewise, in COST’11 two out of three were array problems (with the others on trees) [15]. However, in VerifyThis’16 only one out of three was an array problem [22].

Languages. Fortran, for a long time, has been the undeniable king of high-performance computing in part, at least, because of its strong emphasis on arrays. Early success with vectorising compilers for Fortran77 lead to numerous efforts to support parallelisation across different architectures, such as the High-Performance Fortran (HPF) standard [48], Vienna Fortran [49] and Co-Array Fortran [50]. Fortran arrays are declared via the `DIMENSION` attribute which allows one to specify the *rank* (number of dimensions), *shape* (size of each dimension) and *extent* (ranges for each dimension).

Pascal, despite the lack of dynamic arrays (described as a “defect” by Wirth [51]), also has surprisingly expressive support for static arrays. Since array types include the exact range of permitted indices, interesting subtype relationships can be enforced. For example, a value of type “`ARRAY[1..100] of T`” is not permitted to flow into a location of type “`ARRAY[0..99] of T`”. Indeed, any ordinal type can be used to restrict the range of permitted indices. For example, a variable of type “`ARRAY[byte] of T`” has an index for every possible `byte` value.

The Ada programming language provides a similar mechanism to Pascal for defining and accessing bounded arrays, though additionally supports explicit subtypes [52]. The SPARK/Ada subset [8] also provides syntax for expressing specifications and invariants (and support for static checking), some of which migrated into Ada 2012 [53].

The X10 language from IBM was designed from the beginning with an aim to “*include a rich array sub-language that supports dense and sparse distributed multi-dimensional arrays*” for the purposes of high performance computing [54]. However, the language designers struggled to achieve “*acceptable levels of performance for a single unified array sub-language*” [55]. To resolve this, they instead identified a set of core mechanisms which allowed powerful array primitives to be provided via libraries. X10 also provided, like Whyley, support for constrained types [56]. This meant their libraries, for example, could “*use constrained types to enforce that the number of dimensions used in indexing operations ... matches the rank*” [55].

The Java-based language Lime, also from IBM, bears some similarity to X10, though it focused on exploiting GPU and FPGAs as accelerators [57]. Key differences over Java include the introduction of *deeply immutable* arrays and also *bounded* (i.e. fixed-sized) arrays. The latter plays an important role in Lime for expressing inputs and outputs for stream operators. Furthermore, bounded arrays are always accessed by index expressions which are statically guaranteed to be within bounds which, like Pascal, is achieved through ordinal ranges.

Trojahn and Grellck argued that array programming “imposes non-trivial structural constraints on ranks, shapes, and element values of arrays” [58]. They employed dependent types on Single Assignment C (SAC) (a compiled functional array language) programs to ensure such constraints are enforced. For example, that array accesses are within bounds, and that arguments for function invocations have appropriate shape [59]. In doing this, they acknowledge that the system may sometimes be unable to decide correctness and, in such cases, generate runtime checks instead. Finally, we note that it would be interesting to consider the use of other dependently typed languages (e.g. Agda [60], Idris [61], F* [62], etc) for array programming.

6. Conclusion

Through a series of short case studies, we have explored how specifications and invariants enhance array programming. We believe much is gained from expressing specifications and invariants as these are typically otherwise hidden. This is especially true for array programs which often involve subtle and complex invariants. We also find the specification languages used in systems like Whiley and Dafny are ideally suited for describing array problems. Of course, their real power comes from statically verifying programs meet their specifications, and we hope further advances will see them used more widely.

Acknowledgments

The author would like to thank the anonymous reviewers for their helpful feedback on this paper.

References

- [1] D. Detlefs, G. Nelson, and J. B. Saxe. Simplify: A theorem prover for program checking. *Journal of the ACM*, 52(3):365–473, 2005.
- [2] L. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *Proceedings of the conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 337–340, 2008.
- [3] C. Flanagan, K. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for Java. In *Proceedings of the ACM conference on Programming Language Design and Implementation (PLDI)*, pages 234–245, 2002.
- [4] M. Barnett, M. Fähndrich, K. R. M. Leino, P. Müller, W. Schulte, and H. Venter. Specification and verification: the Spec# experience. *Communications of the ACM*, 54(6):81–91, 2011.
- [5] K. Rustan M. Leino. Dafny: An automatic program verifier for functional correctness. In *Proceedings of the Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, volume 6355 of *LNC3*, pages 348–370. Springer-Verlag, 2010.
- [6] J. Filliâtre and A. Paskevich. Why3 — where programs meet provers. In *Proceedings of the European Symposium on Programming (ESOP)*, pages 125–128, 2013.
- [7] B. Jacobs, J. Smans, P. Philippaerts, F. Vogels, W. Penninckx, and F. Piessens. Verifast: A powerful, sound, predictable, fast verifier for C and Java. In *Proceedings of the NASA Formal Methods Symposium*, pages 41–55. Springer-Verlag, 2011.
- [8] J. Barnes. *High Integrity Ada: The SPARK Approach*. Addison Wesley Longman, Inc., Reading, 1997.
- [9] The Whiley Programming Language, <http://whiley.org>.
- [10] D. J. Pearce and L. Groves. Whiley: a platform for research in software verification. In *Proceedings of the Conference on Software Language Engineering (SLE)*, pages 238–248, 2013.
- [11] D. J. Pearce and L. Groves. Designing a verifying compiler: Lessons learned from developing Whiley. *Science of Computer Programming*, pages 191–220, 2015.
- [12] P. Müller, M. Schwerhoff, and A. J. Summers. Viper: A verification infrastructure for permission-based reasoning. In *Proceedings of the Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI)*, pages 41–62, 2016.
- [13] D. R. Cok and J. Kiniry. ESC/Java2: Uniting ESC/Java and JML. In *Proceedings of the Conference on Construction and Analysis of Safe, Secure, and Interoperable Smart Devices (CASSIS)*, pages 108–128, 2005.
- [14] D. G. Kourie and B. W. Watson. *The Correctness-by-Construction Approach to Programming*. Springer, 2012.
- [15] T. Borner, M. Brockschmidt, D. Distefano, G. Ernst, J. Filliâtre, R. Grigore, M. Huisman, V. Klebanov, C. Marché, R. Monahan, W. Mostowski, N. Polikarpova, C. Scheben, G. Schellhorn, B. Tofan, J. Tschannen, and M. Ulbrich. The COST IC0701 verification competition 2011. In *Proc. FoVeOOS*, pages 3–21, 2011.
- [16] K. R. M. Leino. Developing verified programs with Dafny. In *Proceedings of the Conference on Verified Software: Theories, Tools, Experiments (VSTTE)*, pages 82–82, 2012.
- [17] E. Cohen, M. Dahlweid, M. Hillebrand, D. Leinenbach, M. Moskal, T. Santen, W. Schulte, and S. Tobies. VCC: A practical system for verifying concurrent C. In *Proceedings of the Conference on Theorem Proving in Higher Order Logics (TPHOL)*, pages 23–42, 2009.
- [18] I. Kassios. Dynamic frames: Support for framing, dependencies and sharing without restrictions. In *Proceedings of the Symposium on Formal Methods (FM)*, pages 268–283, 2006.
- [19] K.R.M. Leino and R. Monahan. Dafny meets the verification benchmarks challenge. In *Proceedings of the Conference on Verified Software: Theories, Tools, Experiments (VSTTE)*, pages 112–126, 2010.
- [20] V. Klebanov, P. Müller, N. Shankar, G. T. Leavens, V. Wüstholtz, E. Alkassar, R. Arthan, D. Bronish, R. Chapman, E. Cohen, M. Hillebrand, B. Jacobs, K.R.M. Leino, R. Monahan, F. Piessens, N. Polikarpova, T. Ridge, J. Smans, S. Tobies, T. Tuerk, M. Ulbrich, and B. Weiß. The 1st verified software competition: Experience report (VSComp). In *Proceedings of the Symposium on Formal Methods (FM)*, 2011.
- [21] M. Huisman, V. Klebanov, and R. Monahan. Verifythis verification competition 2012 - organizer’s report, 2013.
- [22] M. Huisman, R. Monahan, P. Müller, and E. Poll. Verifythis 2016 — a program verification competition. *International Journal on Software Tools for Technology Transfer (STTT)*, 2016.
- [23] C. Flanagan and S. Qadeer. Predicate abstraction for software verification. In *Proceedings of the ACM symposium on the Principles Of Programming Languages (POPL)*, pages 191–202, 2002.
- [24] B. Beckert, S. Schlager, and P. Schmitt. An improved rule for while loops in deductive program verification. In *Proceedings of the International Conference on Formal Engineering Methods (ICFEM)*, pages 315–329, 2005.
- [25] R. Chadha and D. A. Plaisted. On the mechanical derivation of loop invariants. *Journal of Symbolic Computation*, 15(5 & 6):705–744, 1993.
- [26] A. Ireland, B. Ellis, and T. Ingulfen. Invariant patterns for program reasoning. In *Proceedings of the Mexican International Conference on Artificial Intelligence (MICAI)*, pages 190–201, 2004.
- [27] K. R. M. Leino and F. Logozzo. Loop invariants on demand. In *Proceedings of the Asian Symposium on Programming Languages and Systems (APLAS)*, pages 119–134, 2005.
- [28] C. A. Furia and Bertrand Meyer. Inferring loop invariants using postconditions. *CoRR*, abs/0909.0884, 2009.

- [29] M. Aponte, P. Courtieu, Y. Moy, and M. Sango. Maximal and compositional pattern-based loop invariants. In *Proceedings of the Symposium on Formal Methods (FM)*, pages 37–51, 2012.
- [30] D. Cachera, T. P. Jensen, A. Jobin, and F. Kirchner. Inference of polynomial invariants for imperative programs: A farewell to Gröbner bases. In *Proceedings of the Static Analysis Symposium (SAS)*, pages 58–74, 2012.
- [31] M. Colón, S. Sankaranarayanan, and H. Sipma. Linear invariant generation using non-linear constraint solving. In *Proceedings of Conference on Computer Aided Verification (CAV)*, volume 2725 of *LNCS*, pages 420–432. Springer-Verlag, 2003.
- [32] S. Sankaranarayanan, H. B. Sipma, and Z. Manna. Non-linear loop invariant generation using Grobner bases. In *Proceedings of the ACM symposium on the Principles Of Programming Languages (POPL)*, pages 318–329. ACM Press, 2004.
- [33] R. Bagnara, E. Rodríguez-Carbonell, and E. Zaffanella. Generation of basic semi-algebraic invariants using convex polyhedra. In *Proceedings of the Static Analysis Symposium (SAS)*, volume 3672 of *LNCS*, pages 19–34. Springer-Verlag, 2005.
- [34] S. Lahiri and R. Bryant. Constructing quantified invariants via predicate abstraction. In *Proceedings of the Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI)*, volume 2937 of *LNCS*, pages 267–281. Springer-Verlag, 2004.
- [35] S. Lahiri, T. Ball, and B. Cook. Predicate abstraction via symbolic decision procedures. In *Proceedings of Conference on Computer Aided Verification (CAV)*, volume 3576 of *LNCS*, pages 24–38. Springer-Verlag, 2005.
- [36] S. Srivastava and S. Gulwani. Program verification using templates over predicate abstraction. In *Proceedings of the ACM conference on Programming Language Design and Implementation (PLDI)*, pages 223–234. ACM Press, 2009.
- [37] D. Gries. *The science of programming*. Springer-Verlag, 1981.
- [38] E. W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
- [39] K. Broda, S. Eisenbach, H. Khoshnevisan, and Steven Vickers. *Reasoned Programming*. Prentice Hall, 1994.
- [40] R. Backhouse. *Program Construction*. Wiley, 2003.
- [41] C. Marche, C. Paulin Mohring, and X. Urbain. The KRAKATOA tool for certification of JAVA/JAVACARD programs annotated in JML. *JLAP*, 58(1–2):89–106, 2004.
- [42] C. Hunter, P. Robinson, and P. Strooper. Agent-based distributed software verification. In *Proceedings of the Australasian conference on Computer Science (ACSC)*, pages 159–164, 2005.
- [43] E. W. Dijkstra. A constructive approach to the problem of program correctness. *BIT*, 8:174–186, 1968.
- [44] C. A. R. Hoare. Proof of a program: FIND. *Communications of the ACM*, 14:39–45, 1971.
- [45] I. Dony and B. Le Charlier. A tool for helping teach a programming method. In *Proceedings of the Conference on Innovation and technology in Computer Science Education (ITICSE)*, pages 212–216, 2006.
- [46] R. Back. Invariant based programming: basic approach and teaching experiences. *Formal Aspects of Computing*, 21(3):227–244, 2009.
- [47] O. Astrachan. Pictures as invariants. In *Proceedings of the Technical Symposium on Computer Science Education (SIGCSE)*, pages 112–118. ACM, 1991.
- [48] K. Kennedy, C. Koebel, and H. Zima. The rise and fall of high performance fortran. *Communications of the ACM*, 54(11):74–82, 2011.
- [49] B. Chapman, P. Mehrotra, and H. Zima. Programming in vienna fortran. *Science of Computer Programming*, 1(1):31–50, 1992.
- [50] R. Numrich and J. Reid. Co-array fortran for parallel programming. *ACM Fortran Forum*, 17(2):1–31, 1998.
- [51] N. Wirth. Recollections about the development of pascal. In *Proceedings of the Conference on History of Programming Languages (HOPL)*, pages 333–342. ACM Press, 1993.
- [52] J. Barnes. *Programming in Ada 2012*. Addison Wesley, 2014.
- [53] R. Dewar. Ada2012: Ada with contracts. *Dr. Dobbs's Journal*, 2013.
- [54] P. Charles, C. Grothoff, V. Saraswat, C. Donawa, A. Kielstra, K. Ebcioglu, C. Praun, and V. Sarkar. X10: an object-oriented approach to non-uniform cluster computing. In *Proceedings of the ACM conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, pages 519–538, 2005.
- [55] D. Grove, J. Milthorpe, and O. Tardieu. Supporting array programming in X10. In *Proceedings of the Workshop on Libraries, Languages and Compilers for Array Programming (ARRAY)*, pages 38–43, 2014.
- [56] N. Nystrom, V. Saraswat, J. Palsberg, and C. Grothoff. Constrained types for object-oriented languages. In *Proceedings of the ACM conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, pages 457–474, 2008.
- [57] J. Auerbach, D. Bacon, P. Cheng, and R. Rabbah. Lime: a Java-compatible and synthesizable language for heterogeneous architectures. In *Proceedings of the ACM conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, pages 89–108, 2010.
- [58] K. Trojahner and C. Grelck. Dependently typed array programs don't go wrong. *Journal of Logic and Algebraic Programming*, 78(7):643, 2009.
- [59] C. Grelck and F. Tang. Towards hybrid array types in SAC. In *Proceedings of the Software Engineering Workshop*, pages 129–145, 2014.
- [60] A. Bove, P. Dybjer, and U. Norell. A brief overview of agda - A functional language with dependent types. In *Proceedings of the Conference on Theorem Proving in Higher Order Logics (TPHOL)*, volume 5674 of *LNCS*, pages 73–78, 2009.
- [61] E. Brady. Idris, a general-purpose dependently typed programming language: Design and implementation. *Journal of Functional Programming*, 23(5):552–593, 2013.
- [62] N. Swamy, C. Hritcu, C. Keller, A. Rastogi, A. Delignat-Lavaud, S. Forest, K. Bhargavan, C. Fournet, P. Strub, M. Kohlweiss, J. Zinzindohoue, and S. Béguélin. Dependent types and multi-monadic effects in F. In *Proceedings of the ACM symposium on the Principles Of Programming Languages (POPL)*, pages 256–270, 2016.