Chapter 2

Dependent Types for Distributed Arrays

Wouter Swierstra[1] and Thorsten Altenkirch[1]
Category: Research Paper

Abstract: Locality-aware algorithms over distributed arrays can be very difficult to write. Yet such algorithms are becoming more and more important as desktop machines boast more and more processors. We show how a dependently-typed programming language can help develop such algorithms by hosting a domain-specific embedded language that ensures every well-typed program will only ever access local data. Such static guarantees can help catch programming errors early on in the development cycle and maximise the potential speedup that multicore machines offer. At the same time, the functional specification of effects we provide facilitates the testing of and reasoning about algorithms that manipulate distributed arrays.

2.1 INTRODUCTION

Computer processors are not becoming significantly faster. To satisfy the demand for more and more computational power, manufacturers are now assembling computers with multiple microprocessors. It is hard to exaggerate the impact this will have on software development: tomorrow's programming languages must embrace parallel programming on multicore machines.

Researchers have proposed several new languages to maximise the potential speedup that multicore processors offer [2, 6, 7, 8, 12, 18]. Although all these languages are different, they share the central notion of a *distributed array*, where the elements of an array may be distributed over separate processors or even over separate machines. To write efficient code, programmers must ensure that processors only access *local* parts of a distributed array—it is much faster to access data stored locally than remote data on another core.

[1]University of Nottingham, {wss,txa}@cs.nott.ac.uk

When writing such locality-aware algorithms it is all too easy to make subtle mistakes. Programming languages such as X10 require all arrays operations to be local [8]. Any attempt to access non-local data results in an exception. To preclude such errors, Grothoff *et al.* have designed a type system, based on a dependently-typed lambda calculus, for a small core language resembling X10 that is specifically designed to guarantee that programs only access local parts of a distributed array [9]. Their proposed system is rather intricate and consists of a substantial number of complicated type rules that keep track of locality information.

In this paper, we explore an alternative avenue of research. Designing and implementing a type system from scratch is a lot of work. New type systems typically require extensive proofs of various meta-theoretical results. Instead, we show how to tailor a powerful type system to enforce certain properties—resulting in a *domain-specific embedded type system*. We immediately inherit all the desirable properties of our dependently-typed host type system, such as subject reduction, decidable type checking, and principle typing. Functional programmers have studied domain-specific embedded languages for years [11]; the time is ripe to take these ideas one step further.

In previous work [20], we described a pure specification of several parts of the IO monad, the interface between pure functional languages such as Haskell [16] and the 'real world.' By providing functional, executable specifications we can test, debug, and reason about impure programs as if they were pure. When we release the final version of our code, we can replace our pure specifications with their impure, more efficient, counterparts. In the presence of dependent types, we will show how our specifications can provide even stronger static guarantees about our programs. To this end, we make several novel contributions:

- We begin by giving a pure specification of arrays (Section 2.3). Due to our rich type system, the specification is *total*: there is no way to access unallocated memory; there are no 'array index out of bounds' exceptions. As a result, these specifications can not only be used to *program* with, but also facilitate *formal proofs* about array algorithms.
- Distributed arrays pose more of a challenge (Section 2.4). Not only do we attend to locality constraints, but we must also accommodate place-shifting operators. The pure specification we present is, once again, executable and total: it can be interpreted both as a domain-specific embedded language for writing algorithms on distributed arrays and as an executable denotational model for specifying and proving properties of such algorithms.
- Finally, we demonstrate how programmers may write their own locality-aware control structures. We implement a simple distributed algorithm using these control structures, and conclude by discussing the limitations of our approach and directions for further research (Section 2.5).

Throughout this paper, we will use the dependently-typed programming language Agda [1, 15] as a vehicle of explanation. In fact using lhs2TeX [13], the

sources of this paper generate an Agda program that can be compiled and executed.[1] We will briefly introduce the syntax by means of several examples, as it may be unfamiliar to many readers.

2.2 AN OVERVIEW OF AGDA

Data types in Agda can be defined using a similar syntax to that for Generalized Algebraic Data Types, or GADTs, in Haskell [17]. For example, consider the following definition of the natural numbers.

```
data Nat : ⋆ where
  Zero : Nat
  Succ : Nat → Nat
```

There is one important difference with Haskell. We must explicitly state the *kind* of the data type that we are introducing; in particular, the declaration $Nat : \star$ states that *Nat* is a base type.

We can define functions by pattern matching and recursion, just as in any other functional language. To define addition of natural numbers, for instance, we could write:

```
_+_ : Nat → Nat → Nat
Zero   + m = m
Succ n + m = Succ (n + m)
```

Note that Agda uses underscores to denote the positions of arguments when defining new operators.

Polymorphic lists are slightly more interesting than natural numbers:

```
data List (a : ⋆) : ⋆ where
  Nil : List a
  Cons : a → List a → List a
```

To uniformly parameterise a data type, we can write additional arguments to the left of the copula. In this case, we add $(a : \star)$ to our data type declaration to state that lists are type *constructors*, parameterised over a type variable a of kind $\star$.

Just as we defined addition for natural numbers, we can define an operator that appends one list to another:

```
append : (a : ⋆) → List a → List a → List a
append a Nil         ys = ys
append a (Cons x xs) ys = Cons x (append a xs ys)
```

The *append* function is polymorphic. In Agda, such polymorphism can be introduced via the *dependent function space*, written $(x : a) \to y$, where the variable x may occur in the type y. This particular example of the dependent function space is not terribly interesting: it corresponds to parametric polymorphism. Later we will encounter more interesting examples, where *types* depend on *values*.

One drawback of using the dependent function space for such parametric poly-

[1] The resulting code is available from the first author's website.

morphism, is that we must explicitly instantiate polymorphic functions. For example, the recursive call to *append* in the *Cons* case takes a type as its first argument. Fortunately, Agda allows us to mark certain arguments as *implicit*. Using implicit arguments, we could also define *append* as in any other functional language:

$$
\begin{aligned}
&\mathit{append} : \{a : \star\} \to \mathit{List}\ a \to \mathit{List}\ a \to \mathit{List}\ a \\
&\mathit{append}\ \mathit{Nil} \qquad\quad \mathit{ys} = \mathit{ys} \\
&\mathit{append}\ (\mathit{Cons}\ x\ \mathit{xs})\ \mathit{ys} = \mathit{Cons}\ x\ (\mathit{append}\ \mathit{xs}\ \mathit{ys})
\end{aligned}
$$

Arguments enclosed in curly brackets, such as $\{a : \star\}$, are implicit: we do not write a to the left of the equals sign and do not pass a type argument when we make a recursive call. The Agda type checker will automatically instantiate this function whenever we call it, much in the same way as type variables are automatically instantiated in Haskell.

Besides polymorphic data types, Agda also supports *indexed families*. Like Haskell's GADTs, indexed families allow a data type's constructors to have different codomains. Indexed families, however, are more general as they also capture data types that are indexed by *values* instead of types. For example, we can define the family of finite types:

$$
\begin{aligned}
&\textbf{data}\ \mathit{Fin} : \mathit{Nat} \to \star\ \textbf{where} \\
&\quad \mathit{Fz} : \{n : \mathit{Nat}\} \to \mathit{Fin}\ (\mathit{Succ}\ n) \\
&\quad \mathit{Fs} : \{n : \mathit{Nat}\} \to \mathit{Fin}\ n \to \mathit{Fin}\ (\mathit{Succ}\ n)
\end{aligned}
$$

The type $\mathit{Fin}\ n$ corresponds to a finite type with n distinct values. For example, $\mathit{Fin}\ 1$ is isomorphic to the unit type; $\mathit{Fin}\ 2$ is isomorphic to *Bool*. Note that the argument n is left implicit in both the constructors of *Fin*. From the types of these constructors, it is easy to see that $\mathit{Fin}\ 0$ is uninhabited. For every n, the *Fs* constructor embeds $\mathit{Fin}\ n$ into $\mathit{Fin}\ (\mathit{Succ}\ n)$; the *Fz* constructor, on the other hand, adds a single new element to $\mathit{Fin}\ (\mathit{Succ}\ n)$ that was not in $\mathit{Fin}\ n$. By induction it is easy to see that $\mathit{Fin}\ n$ does indeed have n elements.

Agda has many other features, such as records and a module system, that we will hardly use in this paper. Although there are a few more concepts we will need, we will discuss them as they pop up in later sections.

2.3 MUTABLE ARRAYS

With this brief Agda tutorial under our belt, we can start our specification of mutable arrays. We will specify three different operations on arrays: the creation of new arrays; reading from an array; and updating a value stored in an array. Before we can define the behaviour of these operations, however, we need to introduce several data types to describe the layout and contents of memory. Using these data types, we can proceed by defining an *IO* type that captures the syntax of array operations. Finally, we will define a *run* function that describes how the array operations affect the heap, assigning semantics to our syntax. This semantics can be used to simulate and reason about computations on mutable arrays in a pure functional language. When compiled, however, these operations should be

replaced by their more efficient, low-level counterparts.

To keep things simple, we will only work with flat arrays storing natural numbers. This is, of course, a drastic oversimplification. The techniques we present here, however, can be adapted to cover multidimensional arrays that may store different types of data (Section 2.5).

To avoid confusion between numbers denoting the size of an array and the data stored in an array, we introduce the *Data* type synonym. Throughout the rest of this paper, we will use *Data* to refer to the data stored in arrays; the *Nat* type will always refer to the size of an array.

```
Data : ⋆
Data = Nat
```

Using the *Fin* type, we can give a functional specification of arrays of a fixed size by mapping every index to the corresponding value.

```
Array : Nat → ⋆
Array n = Fin n → Data
```

How should we represent the heap? We need to be a bit careful—as the heap will store arrays of different sizes, its type should explicitly state how many arrays it stores and how large each array is. To accomplish this, we begin by introducing a data type representing the *shape* of the heap:

```
Shape : ⋆
Shape = List Nat
```

The *Shape* of the heap is simply a list of natural numbers, representing the size of the arrays stored in memory.

We can now define a *Heap* data type that is indexed by a *Shape*. The *Empty* constructor corresponds to an empty heap; the *Alloc* constructor adds an array of size *n* to any heap of shape *ns* to build a larger heap with the layout *Cons n ns*.

```
data Heap : Shape → ⋆ where
  Empty : Heap Nil
  Alloc : {n : Nat} → {ns : Shape} →
    Array n → Heap ns → Heap (Cons n ns)
```

Finally, we will want to model references, denoting locations in the heap. A value of type *Loc n ns* corresponds to a reference to an array of size *n* in a heap with shape *ns*. The *Loc* data type shares a great deal of structure with the *Fin* type. Every non-empty heap has a *Top* reference; we can weaken any existing reference to denote the same location in a larger heap using the *Pop* constructor.

```
data Loc : Nat → Shape → ⋆ where
  Top : {n : Nat} → {ns : Shape} → Loc n (Cons n ns)
  Pop : forall {n k ns} → Loc n ns → Loc n (Cons k ns)
```

Note that in the type signature of the *Pop* constructor, we omit the types of three implicit arguments and quantify over them using the **forall** keyword. When we use the **forall**-notation, the types of *n*, *k*, and *ns* are inferred from the rest of the signature by the Agda type checker. Alternatively, we could also have written the

more verbose:

```
Pop : {n : Nat} → {k : Nat} → {ns : Shape} →
  Loc n ns → Loc n (Cons k ns)
```

We will occasionally use the **forall**-notation to make large type signatures somewhat more legible.

With these data types in place, we can define a data type capturing the syntax of the permissible operations on arrays. Crucially, the *IO* type is indexed by *two* shapes: a value of type *IO a ns ms* denotes a computation that takes a heap of shape *ns* to a heap of shape *ms* and returns a result of type *a*. This pattern of indexing operations by an initial and final 'state' is a common pattern in dependently-typed programming [14].

```
data IO (a : ⋆) : Shape → Shape → ⋆ where
  Return : {ns : Shape} → a → IO a ns ns
  Write : forall {n ns ms} →
    Loc n ns → Fin n → Data → IO a ns ms → IO a ns ms
  Read : forall {n ns ms} →
    Loc n ns → Fin n → (Data → IO a ns ms) → IO a ns ms
  New : forall {ns ms} →
    (n : Nat) → (Loc n (Cons n ns) → IO a (Cons n ns) ms) →
    IO a ns ms
```

The *IO* type has four constructors. The *Return* constructor returns a pure value of type *a* without modifying the heap. The *Write* constructor takes four arguments: the location of an array of size *n*; an index in that array; the value to write at that index; and the rest of the computation. Similarly, reading from an array requires a reference to an array and an index. Instead of requiring the data to be written, however, the last argument of the *Read* constructor may refer to data that has been read. Finally, the *New* constructor actually changes the size of the heap. Given a number *n*, it allocates an array of size *n* on the heap; the second argument of *New* may then use this fresh reference to continue the computation in a larger heap.

The *IO* data type is a *parameterised monad* [3]—that is, a monad with *return* and bind operators that satisfy certain coherence conditions with respect to the *Shape* indices.

```
return : {a : ⋆} → {ns : Shape} → a → IO a ns ns
return x = Return x

_>>=_ : forall {a b ns ms ks} →
  IO a ns ms → (a → IO b ms ks) → IO b ns ks
Return x >>= f       = f x
Write a i x wr >>= f = Write a i x (wr >>= f)
Read a i rd >>= f    = Read a i (λx → rd x >>= f)
New n io >>= f       = New n (λa → io a >>= f)
```

The return of the *IO* data type lifts a pure value into a computation that can run on a heap of any size. Furthermore, *return* does not modify the shape of the heap. The bind operator, $\gg\!=$, can be used to compose monadic computations. To

sequence two computations, the heap resulting from the first computation must be a suitable starting point for the second computation. This condition is enforced by the type of the bind operator:

To actually program using these array operations, we need to introduce smart constructors. For example, we could define the *readArray* function as follows:

```
readArray : forall {n ns} → Loc n ns → Fin n → IO Data ns ns
readArray a i = Read a i Return
```

There is a slight problem with this definition. As we allocate new memory, the size of the heap changes; correspondingly, we must explicitly modify any existing pointers to denote locations in a larger heap. We can achieve this by revising the above definition slightly, applying the *inj* function to weaken references:

```
inj : forall {ms ns n} → Loc n ns → Loc n (append ms ns)
inj {Nil}        i = i
inj {Cons k ks} i = Pop (inj i)
```

For the purpose of this paper, however, we will ignore this technicality. The first definition will suffice for the examples we cover. For a more comprehensive discussion, we refer to the first author's forthcoming thesis [19].

Denotational model We have described the syntax of array computations using the *IO* data type, but we have not specified how these computations *behave*. Recall that we can model arrays as functions from indices to natural numbers:

```
Array : Nat → ⋆
Array n = Fin n → Data
```

Before specifying the behaviour of *IO* computations, we define several auxiliary functions to update an array and lookup a value stored in an array.

```
lookup : forall {n ns} → Loc n ns → Fin n → Heap ns → Data
lookup Top      i (Alloc a _) = a i
lookup (Pop k) i (Alloc _ h) = lookup k i h
```

The *lookup* function takes a reference to an array l, an index i in the array at location l, and a heap, and returns the value stored in the array at index i. It dereferences l, resulting in a function of type $Fin\ n \to Data$; the value stored at index i is the result of applying this function to i.

Next, we define a pair of functions to update the contents of an array.

```
updateArray : {n : Nat} → Fin n → Data → Array n → Array n
updateArray i d a = λj → if i ≡ j then d else a j

updateHeap : forall {n ns} →
  Loc n ns → Fin n → Data → Heap ns → Heap ns
updateHeap Top      i x (Alloc a h) = Alloc (updateArray i x a) h
updateHeap (Pop k) i x (Alloc a h) = Alloc a (updateHeap k i x h)
```

The *updateArray* function overwrites the data stored at a single index. The function *updateHeap* updates a single index of an array stored in the heap. It proceeds by dereferencing the location on the heap where the desired array is stored and

updates it accordingly, leaving the rest of the heap unchanged.

We now have all the pieces in place to assign semantics to *IO* computations. The *run* function below takes a computation of type *IO a ns ms* and an initial heap of shape *ns* as arguments, and returns a pair consisting of the result of the computation and the final heap of shape *ms*.

```
data Pair (a : ⋆) (b : ⋆) : ⋆ where
  pair : a → b → Pair a b

run : forall {a ns ms} → IO a ns ms → Heap ns → Pair a (Heap ms)
run (Return x) h       = pair x h
run (Read a i rd) h    = run (rd (lookup a i h)) h
run (Write a i x wr) h = run wr (updateHeap a i x h)
run (New n io) h       = run (io Top) (Alloc (λ i → Zero) h)
```

The *Return* constructor simply pairs the result and heap; in the *Read* case, we lookup the data from the heap and recurse with the same heap; for the *Write* constructor, we recurse with an appropriately modified heap; finally, when a new array is created, we extend the heap with a new array that stores *Zero* at every index, and continue recursively. Note that, by convention, the *Top* constructor always refers to the most recently created reference. Our smart constructors will add additional *Pop* constructors when new memory is allocated.

We refer to this specification as a denotational model. As Agda is a programming language based on type theory, may also view it as a constructive set theory. In that sense, the *run* function constitutes a denotational semantics of mutable arrays. By implementing these semantics in Agda, we build an executable denotational model in Agda's underlying type theory.

Example Using our smart constructors and the monad operators, we can now define functions that manipulate arrays. For example, the swap function exchanges the value stored at two indices:

```
swap : forall {n ns} → Loc n ns → Fin n → Fin n → IO () ns ns
swap a i j = readArray a i >>= λ vali →
             readArray a j >>= λ valj →
             writeArray a i valj >>
             writeArray a j vali
```

In a dependently-typed programming language such as Agda, we can prove properties of our code. For example, we may want to show that swapping the contents of any two array indices twice, leaves the heap intact :

```
swapProp : forall {n ns} →
  (l : Loc n ns) → (i : Fin n) → (j : Fin n) → (h : Heap ns) →
  (h ≡ snd (run (swap l i j >> swap l i j) h))
```

The proof requires a lemma about how *updateHeap* and *lookupHeap* interact and is not terribly interesting in itself. The fact that we can formalise such properties and have our proof verified by a computer is much more exciting.

2.4 DISTRIBUTED ARRAYS

Arrays are usually represented by a continuous block of memory. *Distributed arrays*, however, can be distributed over different *places*—where every place may correspond to a different core on a multiprocessor machine, a different machine on the same network, or any other configuration of interconnected computers.

We begin by determining the type of places, where data is stored and code is executed. Obviously, we do not want to fix the type of all possible places prematurely: you may want to execute the same program in different environments. Yet regardless of the exact number of places, there are certain operations you will always want to perform, such as iterating over all places, or checking when two places are equal.

We therefore choose to abstract over the *number* of places in the module we will define in the coming section. Agda allows modules to be paramaterised:

module *DistrArray* (*placeCount* : *Nat*) **where**

When we import the *DistrArray* module, we are obliged to choose the number of places. Typically, there will be one place for every available processor. From this number, we can define a data type corresponding to the available places:

Place : $\star$
Place = *Fin placeCount*

The key idea underlying our model of locality-aware algorithms is to index computations by the place where they are executed. The new type declaration for the *IO* monad corresponding to operations on *distributed* arrays will become:

data *DIO* (*a* : $\star$) : *Shape* $\rightarrow$ *Place* $\rightarrow$ *Shape* $\rightarrow$ $\star$ **where**

You may want to think of a value of type *DIO a ns p ms* as a computation that can be executed at place *p* and will take a heap of shape *ns* to a heap of shape *ms*, yielding a final value of type *a*.

We strive to ensure that any well-typed program written in the *DIO* monad will never access data that is not local. The specification of distributed arrays now poses a twofold problem: we want to ensure that the array manipulations from the previous section are 'locality-aware,' that is, we must somehow restrict the array indices that can be accessed from a certain place; furthermore, X10 facilitates several *place-shifting* operations that change the place where certain chunks of code are executed. As we shall see in the rest of this section, both these issues can be resolved quite naturally.

Regions, Points, and Distributed Arrays Before we define the *DIO* monad, we need to introduce several new concepts. In what follows, we will try to stick closely to X10's terminology for distributed arrays. Every array is said to have a *region* associated with it. A region is a set of valid index points. A *distribution* specifies a place for every index point in a region.

Once again, we will only treat flat arrays storing natural numbers and defer any discussion about how to deal with more complicated data structures for the

moment. In this simple case, a region merely determines the size of the array.

$$\begin{aligned}&\mathit{Region} : \star\\&\mathit{Region} = \mathit{Nat}\end{aligned}$$

As we have seen in the previous section, we can model array indices using the *Fin* data type:

$$\begin{aligned}&\mathit{Point} : \mathit{Region} \to \star\\&\mathit{Point}\ n = \mathit{Fin}\ n\end{aligned}$$

To model distributed arrays, we now need to consider the distribution that specifies *where* this data is stored. In line with existing work [9], we assume the existence of a fixed distribution. Agda's **postulate** expression allows us to assume the existence of a distribution, without providing its definition.

$$\begin{aligned}&\textbf{postulate}\\&\quad \mathit{distr} : \textbf{forall}\ \{n\ ns\} \to \mathit{Loc}\ n\ ns \to \mathit{Point}\ n \to \mathit{Place}\end{aligned}$$

Although we have implemented several of X10's combinators for defining distributions, we do not have the space to cover them here.

Now that we have all the required auxiliary data types, we proceed by defining the *DIO* monad. As it is a bit more complex than the data types we have seen so far, we will discuss every constructor individually.

The *Return* constructor is analogous to one we have seen previously for the *IO* monad: it lifts any pure value into the *DIO* monad.

$$\mathit{Return} : \{p : \mathit{Place}\} \to \{ns : \mathit{Shape}\} \to a \to \mathit{DIO}\ a\ ns\ p\ ns$$

The *Read* and Write operations are more interesting. Although they correspond closely to the operations we have seen in the previous section, their type now keeps track of the place where they are executed. Any read or write operation to point *pt* of an array *l* can *only* be executed at the place specified by the distribution. This invariant is enforced by the types of our constructors:

$$\begin{aligned}&\mathit{Read} : \textbf{forall}\ \{n\ ns\ ms\} \to\\&\quad (l : \mathit{Loc}\ n\ ns) \to (pt : \mathit{Point}\ n) \to\\&\quad (\mathit{Data} \to \mathit{DIO}\ a\ ns\ (\mathit{distr}\ l\ pt)\ ms) \to\\&\quad \mathit{DIO}\ a\ ns\ (\mathit{distr}\ l\ pt)\ ms\\&\mathit{Write} : \textbf{forall}\ \{n\ ns\ ms\} \to\\&\quad (l : \mathit{Loc}\ n\ ns) \to (pt : \mathit{Point}\ n) \to \mathit{Data} \to\\&\quad \mathit{DIO}\ a\ ns\ (\mathit{distr}\ l\ pt)\ ms \to\\&\quad \mathit{DIO}\ a\ ns\ (\mathit{distr}\ l\ pt)\ ms\end{aligned}$$

In contrast to *Read* and Write, new arrays can be allocated at any place.

$$\begin{aligned}&\mathit{New} : \textbf{forall}\ \{p\ ns\ ms\} \to\\&\quad (n : \mathit{Nat}) \to\\&\quad (\mathit{Loc}\ n\ (\mathit{Cons}\ n\ ns) \to \mathit{DIO}\ a\ (\mathit{Cons}\ n\ ns)\ p\ ms) \to\\&\quad \mathit{DIO}\ a\ ns\ p\ ms\end{aligned}$$

Finally, we add a constructor for a place-shifting operator. Using this *At* operator lets us execute a computation at another place.

$$
\begin{aligned}
&\mathit{At} : \textbf{forall}\ \{p\ ns\ ms\ ps\} \rightarrow \\
&\quad (q : \mathit{Place}) \rightarrow \mathit{DIO}\ ()\ ns\ q\ ms \rightarrow \mathit{DIO}\ a\ ms\ p\ ps \rightarrow \mathit{DIO}\ a\ ns\ p\ ps
\end{aligned}
$$

Note that we will discard the result of the computation that is executed at another place. We therefore require this computation to return an element of the unit type.

We can add our smart constructors for each these operations, as we have done in the previous section. We can also show that *DIO* is indeed a parameterised monad. We have omitted the definitions of the *return* and bind operators for the sake of brevity:

$$
\begin{aligned}
&\mathit{return} : \textbf{forall}\ \{ns\ a\ p\} \rightarrow a \rightarrow \mathit{DIO}\ a\ ns\ p\ ns \\
&\_ \gg\!= \_ : \textbf{forall}\ \{ns\ ms\ ks\ a\ b\ p\} \rightarrow \\
&\quad \mathit{DIO}\ A\ ns\ p\ ms \rightarrow (A \rightarrow \mathit{DIO}\ B\ ms\ p\ ks) \rightarrow \mathit{DIO}\ B\ ns\ p\ ks
\end{aligned}
$$

It is worth noting that the bind operator $\gg\!=$ can only be used to sequence operations at the same place.

Denotational model To run a computation in the *DIO* monad, we follow the *run* function defined in the previous section closely. Our new *run* function, however, must be locality-aware. Therefore, we parameterise the *run* function explicitly by the place where the computation is executed.

$$
\begin{aligned}
&\mathit{run} : \textbf{forall}\ \{a\ ns\ ms\} \rightarrow \\
&\quad (p : \mathit{Place}) \rightarrow \mathit{DIO}\ a\ ns\ p\ ms \rightarrow \mathit{Heap}\ ns \rightarrow \mathit{Pair}\ a\ (\mathit{Heap}\ ms) \\
&\mathit{run}\ p\ (\mathit{Return}\ x)\ h = \mathit{pair}\ x\ h \\
&\mathit{run}\ .(\mathit{distr}\ l\ i)\ (\mathit{Read}\ l\ i\ rd)\ h = \mathit{run}\ (\mathit{distr}\ l\ i)\ (rd\ (\mathit{lookup}\ l\ i\ h))\ h \\
&\mathit{run}\ .(\mathit{distr}\ l\ i)\ (\mathit{Write}\ l\ i\ x\ wr)\ h = \textbf{let}\ h' = \mathit{updateHeap}\ l\ i\ x\ h \\
&\qquad\qquad\qquad\qquad\qquad\qquad\quad \textbf{in}\ \mathit{run}\ (\mathit{distr}\ l\ i)\ wr\ h' \\
&\mathit{run}\ p\ (\mathit{New}\ n\ io)\ h = \mathit{run}\ p\ (io\ \mathit{Top})\ (\mathit{Alloc}\ (\lambda i \rightarrow \mathit{Zero})\ h) \\
&\mathit{run}\ p\ (\mathit{At}\ q\ io1\ io2)\ h = \mathit{run}\ p\ io2\ (\mathit{snd}\ (\mathit{run}\ q\ io1\ h))
\end{aligned}
$$

Now we can see that the *Read* and *Write* operations may not be executed at *any* place. Recall that the *Read* and *Write* constructors both return computations at the place *distr l i*. When we pattern match on a *Read* or *Write*, we know exactly what the place argument of the *run* function must be. Correspondingly, we do not pattern match on the place argument—we know that the place can only be *distr l i*. Agda's syntax allows us to prefix expressions by a single period, provided we know that there is only one possible value an argument may take. This may be unfamiliar to many functional programmers who are used to thinking of patterns being built-up from variables and constructors: *distr l i* is an expression, not a pattern! The situation is somewhat similar to pattern matching on GADTs in Haskell, which introduces equalities between *types*. The *DIO* monad, however, is indexed by values. As a result, pattern matching in the presence of dependent types may introduce equalities between *values*.

The other difference with respect to the previous *run* function, is the new case for the *At* constructor. In that case, we sequence the two computations *io1* and *io2*. To do so, we first execute the *io1* at *q*, but discard its result; we continue executing the second computation *io2* with the heap resulting from the execution

of *io1* at the location p. Conform to previous proposals [10], we have assumed that *io1* and *io2* are performed synchronously—executing *io1* before continuing with the rest of the computation. Using techniques to model concurrency that we have presented previously [20], we believe we could give a more refined treatment of the X10's *globally asynchronous/locally synchronous* semantics and provide specifications for X10's clocks, *finish*, and *force* constructs.

Locality-aware combinators Using the place-shifting operator *at*, we can define several locality-aware control structures. With our first-class distribution and definition of *Place*, we believe there is no need to define more primitive operations.

The distributed map, for example, applies a function to all the elements of a distributed array at the place where they are stored. We define it in terms of an auxiliary function, *for*, that iterates over all the indices of an array:

$$
\begin{aligned}
&\mathit{for} : \textbf{forall}\ \{n\ ns\ p\} \to (\mathit{Point}\ n \to \mathit{DIO}\ ()\ ns\ p\ ns) \to \mathit{DIO}\ ()\ ns\ p\ ns\\
&\mathit{for}\ \{\mathit{Succ}\ k\}\ \mathit{dio} = \mathit{dio}\ \mathit{Fz} \gg (\mathit{for}\ \{k\}\ (\mathit{dio}\ .\ \mathit{Fs}))\\
&\mathit{for}\ \{\mathit{Zero}\}\quad \mathit{dio} = \mathit{return}\ ()
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{dmap} : \textbf{forall}\ \{n\ ns\ p\} \to (\mathit{Data} \to \mathit{Data}) \to \mathit{Loc}\ n\ ns \to \mathit{DIO}\ ()\ ns\ p\ ns\\
&\mathit{dmap}\ f\ l = \mathit{for}\ (\lambda i \to \mathit{at}\ (\mathit{distr}\ l\ i)\ (\mathit{readArray}\ l\ i \gg\!= \lambda x \to\\
&\qquad\qquad\qquad\qquad \mathit{writeArray}\ l\ i\ (f\ x)))
\end{aligned}
$$

Besides *dmap*, we implement two other combinators: *forallplaces* and *ateach*. The *forallplaces* operation executes its argument computation at all available places. We define it using the *for* function to iterate over all places. The *ateach* function, on the other hand, is a generalisation of the distributed map operation. It iterates over an array, executing its argument operation once for every index of the array, at the place where that index is stored.

$$
\begin{aligned}
&\mathit{forallplaces} : \textbf{forall}\ \{p\ ns\} \to\\
&\quad ((q : \mathit{Place}) \to \mathit{DIO}\ ()\ ns\ q\ ns) \to \mathit{DIO}\ ()\ ns\ p\ ns\\
&\mathit{forallplaces}\ io = \mathit{for}\ (\lambda i \to \mathit{at}\ i\ (io\ i))
\end{aligned}
$$

$$
\begin{aligned}
&\mathit{ateach} : \textbf{forall}\ \{n\ ns\ p\} \to\\
&\quad (l : \mathit{Loc}\ n\ ns) \to ((pt : \mathit{Point}\ n) \to \mathit{DIO}\ ()\ ns\ (\mathit{distr}\ l\ pt)\ ns) \to\\
&\quad \mathit{DIO}\ ()\ ns\ p\ ns\\
&\mathit{ateach}\ l\ io = \mathit{for}\ (\lambda i \to \mathit{at}\ (\mathit{distr}\ l\ i)\ (io\ i))
\end{aligned}
$$

Example We will now show how to write a simple algorithm that sums all the elements of a distributed array. To do so efficiently, we first locally sum all the values at every place. To compute the total sum of all the elements of the array, we add together all these local sums. In what follows, we will need the following auxiliary function, *increment*:

$$
\begin{aligned}
&\mathit{increment} : \textbf{forall}\ \{n\ ns\ p\} \to\\
&\quad (l : \mathit{Loc}\ n\ ns) \to (i : \mathit{Fin}\ n) \to \mathit{Nat} \to (\mathit{distr}\ l\ i \equiv p) \to \mathit{DIO}\ ()\ ns\ p\ ns\\
&\mathit{increment}\ l\ i\ x\ \mathit{Refl} = \mathit{readArray}\ l\ i \gg\!= \lambda y \to \mathit{writeArray}\ l\ i\ (x + y)
\end{aligned}
$$

Note that *increment* is a bit more general than strictly necessary. We could return a computation at *distr l i*, but instead we choose to be a little more general: *increment* can be executed at any place, as long as we have a proof that this place is equal to *distr l i*. The $\equiv$-type is inhabited by single constructor *Refl*.

We can use the *increment* function to define a simple sequential *sum* function:

```
sum : forall {n ns p} → Loc n ns → Loc 1 ns → DIO () ns p ns
sum l out = ateach l (λ i → readArray l i >>= λ n →
                              at (distr out Fz) (increment out Fz n Refl))
```

The *sum* function takes an array as its argument, together with a reference to a single-celled array, *out*. It reads every element of the array, and increments *out* accordingly.

Finally, we can use both these functions to define a parallel sum:

```
psum : forall {n ns} →
  (l : Loc n ns) → (localSums : Loc placeCount ns) →
  ((i : Place) → distr localSums i ≡ i) →
  (out : Loc 1 ns) → DIO Nat ns (distr out Fz) ns
psum l localSums locDistr out =
  ateach l (λ i → (readArray l i >>= λ n →
                     increment localSums (distr l i) n (locDistr (distr l i))))
  >> sum localSums out
  >> readArray out Fz
```

The *psum* function takes four arguments: the array *l* whose elements you would like to sum; an array *localSums* that will store the intermediate sums; an assumption regarding the distribution of this array; and finally, the single-celled array to which we write the result. For every index *i* of the array *l*, we read the value stored at index *i*, and increment the corresponding local sum. We then add together the local sums using our previous sequential *sum* function, and return the final result. We use our assumption about the distribution of the *localSums* array when calling the increment function. Without this assumption, we would have to use the place-shifting operation *at* to update a (potentially) non-local array index.

There are several interesting issues that these examples highlight. First of all, as our *at* function only works on computations returning a unit type, the results of intermediate computations must be collected in intermediate arrays.

More importantly, however, whenever we want to rely on properties of the global distribution, we need to make explicit assumptions in the form of proof arguments. This is rather unfortunate: it would be interesting to research how a specific distribution can be associated with an array when it is created. This would hopefully allow for a more fine-grained treatment of distributions and eliminate the need for explicit proof arguments.

2.5 DISCUSSION

Using a dependently-typed host language, we have shown how to implement a domain-specific library for distributed arrays, together with an embedded type system that guarantees all array access operations are both safe and local. In contrast to existing work [10], we have not designed a specific set of type rules; instead, we have shown how equivalent properties can be enforced by a general-purpose language with dependent types. We have provided semantics for our library in the form of a total, functional specification. Although our semantics may not take the form of deduction rules, they are no less precise or concise. Besides these functional specifications are both executable and amenable to computer-aided formal verification. More generally, we hope that this approach can be extended to other domains: a dependently-typed language accommodates domain specific libraries with their own embedded type systems.

Having said this, there are clearly several serious limitations of this work as it stands. We have had to make several simplifying assumptions. First and foremost, we have assumed that every array only stores natural numbers, disallowing more complex structures such as multi-dimensional arrays. This can be easily fixed by defining a more elaborate *Shape* data type. In its most general form, we could choose our *Shape* data type as a list of types; a heap then corresponds to a list of values of the right type.[2] We decided to restrict ourself to this more simple case for the purpose of presentation. We believe that there is no fundamental obstacle preventing us from incorporating the rich region calculus offered by X10 in the same fashion.

Furthermore, our pure model is rather naive. It would be interesting to explore a more refined model, where every place maintains its own heap. As our example in the previous section illustrated, assuming the presence of a global distribution does not scale well. Decorating every array with a distribution upon its creation should help provide locality-information when it is needed.

We have not discussed how code in the *IO* or *DIO* monad is actually compiled. At the moment, Agda can only be compiled to Haskell. Agda does provide several pragmas to customise how Agda functions are translated to their Haskell counterparts. The ongoing effort to support data parallelism in Haskell [4, 5] may therefore provide us with a most welcome foothold.

There are many features of X10 that we have not discussed here at all. Most notably, we have refrained from modelling many of X10's constructs that enable asynchronous communication between locations, even though we would like to do so in the future.

Finally, we should emphasise that we need to explore larger examples to acquire a better understanding of how this approach scales. At the moment, we cannot predict how efficient the resulting code will be; we do not know how difficult it will be to reason about large, realistic distributed algorithms. Unfortunately,

[2] There are some technical details involving 'size problems' that are beyond the scope of this paper. The standard technique of introducing a universe, closed under natural numbers and arrays, should resolve these issues.

we do not have the space to explore such examples further in this paper. Despite these many limitations, however, we believe this paper provides an important first stepping-stone for such further work.

Acknowledgements We wish to express our gratitude to Jens Palsberg for our interesting discussions; to Ulf Norell for his fantastic new incarnation of Agda; and to Mauro Jaskelioff, Nicolas Oury, Liyang HU, and the anonymous reviewers for their helpful comments on a draft version of this paper.

REFERENCES

[1] Agda. `http://www.cs.chalmers.se/~ulfn/Agda`.

[2] Eric Allen, David Chase, Victor Luchangco, Jan-Willem Maessen, Sukyoung Ryu, Guy L. Steele Jr., and Sam Tobin-Hochstadt. The Fortress language specification. Technical report, Sun Microsystems, Inc., 2005.

[3] Robert Atkey. Parameterised notions of computation. In *Proceedings of the Workshop on Mathematically Structured Functional Programming*, 2006.

[4] Manuel M.T. Chakravarty, Gabriele Keller, Roman Lechtchinsky, and Wolf Pfannenstiel. Nepal – Nested Data-Parallelism in Haskell. In *Euro-Par 2001: Parallel Processing, 7th International Euro-Par Conference*, volume LNCS 2150, 2001.

[5] Manuel M.T. Chakravarty, Roman Leshchinskiy, Simon Peyton Jones, Gabriele Keller, and Simon Marlow. Data Parallel Haskell: a status report. *Proceedings of the 2007 Workshop on Declarative Aspects of Multicore Programming*, 2007.

[6] Brad Chamberlain, Steve Deitz, Mary Beth Hribar, and Wayne Wong. Chapel. Technical report, Cray Inc., 2005.

[7] Bradford L. Chamberlain, Sung-Eun Choi, E. Christopher Lewis, Calvin Lin, Lawrence Snyder, and Derrick Weathersby. ZPL: A machine independent programming language for parallel computers. *Software Engineering*, 26(3), 2000.

[8] Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun, and Vivek Sarkar. X10: an object-oriented approach to non-uniform cluster computing. In *OOPSLA '05*, 2005.

[9] Christian Grothoff, Jens Palsberg, and Vijay Saraswat. Safe arrays via regions and dependent types. Submitted for publication.

[10] Christian Grothoff, Jens Palsberg, and Vijay Saraswat. A type system for distributed arrays. Unpublished draft.

[11] Paul Hudak. Building domain-specific embedded languages. *ACM Computing Surveys*, 28, 1996.

[12] Ben Liblit and Alexander Aiken. Type systems for distributed data structures. In *POPL '00: Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 199–213, 2000.

[13] Andres Löh. lhs2tex. `http://people.cs.uu.nl/andres/lhs2tex/`.

[14] James McKinna and Joel Wright. A type-correct, stack-safe, provably correct, expression compiler in Epigram. To appear in the Journal of Functional Programming.

[15] Ulf Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Chalmers University of Technology, 2007.

[16] Simon Peyton Jones, editor. *Haskell 98 Language and Libraries – The Revised Report*. Cambridge University Press, 2003.

[17] Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Geoffrey Washburn. Simple unification-based type inference for GADTs. In *ICFP '06: Proceedings of the Eleventh ACM SIGPLAN International Conference on Functional Programming*, 2006.

[18] Sven-Bodo Scholz. Single Assignment C — efficient support for high-level array operations in a functional setting. *Journal of Functional Programming*, 13(6):1005–1059, 2003.

[19] Wouter Swierstra. *A Functional Specification of Effects*. PhD thesis, University of Nottingham, 2008.

[20] Wouter Swierstra and Thorsten Altenkirch. Beauty in the beast: a functional semantics of the awkward squad. In *Proceedings of the ACM SIGPLAN Haskell Workshop*, 2007.