

A Lambda Calculus for Transfinite Arrays

Unifying Arrays and Streams

ARTJOMS ŠINKAROVŠ, Heriot-Watt University

SVEN-BODO SCHOLZ, Heriot-Watt University

We propose a design for a functional language that natively supports infinite arrays. We use ordinal numbers to introduce the notion of infinity in shapes and indices. By doing so, we obtain a calculus that naturally extends existing array calculi and, at the same time, allows for recursive specifications as they are found in stream- and list-based settings. Furthermore, the main language construct that can be thought of as an n -fold *cons* operator gives rise to expressing transfinite recursion in data, something that lists or streams usually do not support. This makes it possible to treat the proposed calculus as a unifying theory of arrays, lists and streams. We give an operational semantics of the proposed language, discuss design choices that we have made, and demonstrate its expressibility with several examples. We also demonstrate that the proposed formalism preserves a number of well-known universal equalities from array/list/stream theories, and discuss implementation-related challenges.

CCS Concepts: • **Theory of computation** → **Operational semantics**;

Additional Key Words and Phrases: ordinals, arrays, semantics, functional languages

ACM Reference format:

Artjoms Šinkarovš and Sven-Bodo Scholz. 2017. A Lambda Calculus for Transfinite Arrays. *Proc. ACM Program. Lang.* 1, 1, Article 1 (January 2017), 30 pages.
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Conceptually, lists and streams are different objects. Lists are finite inductive objects that can be characterised as the smallest fixpoint: $\text{Lst } A = \mu X. 1 + A \times X$, and streams are infinite co-inductive objects that are characterised as the greatest fixpoint: $\text{Str } A = \nu X. A \times X$.

Despite these conceptual differences between lists and streams, it has been proven useful to enable programmers to specify functions that can operate on both forms equally well. In particular languages that allow for the construction of cyclic structures can support a list type $[A]$ as the greatest fix point $\nu X. 1 + A \times X$ without requiring extra implementation effort. With this construction, any function that operates on lists inherently is applicable to streams as well.

A similar unification of streams and arrays is less straight-forward. The main obstacle to such a unification lies in the fact that array computations usually make heavy use of random access selections, while stream computations are expressed in a step-wise fashion on a temporarily available window of elements. This difference has led to two distinct programming styles: stream processing [Hinze 2010; Stephens 1997; Thies et al. 2002] and array programming [Grelck and Scholz 2006; IBM 1994; Svensson and Svenningsson 2014]. If we want to apply some array-based program to a stream, it typically requires the given program to be massively rewritten.

The key towards a unification of arrays and streams, at least on a conceptual level, becomes evident when looking at arrays as index-value mappings. We can model arrays of element type A as a family of types:

$$[A]_n = \text{Fin}(n) \rightarrow A \quad n : \text{Nat}$$

2017. 2475-1421/2017/1-ART1 \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

where $\text{Fin}(n)$ denotes the set $\{0, \dots, n-1\}$.

With this in mind, we can observe the following correspondence for streams:

$$\text{Str } A \simeq A^\omega \simeq \text{Nat} \rightarrow A \simeq [A]_\omega$$

Streams are isomorphic to infinite sequences, and A^ω is an exponential object that can be seen as a mapping of positions in that sequence to its values. Such an object is nothing but an array of infinite length. Consequently, a unification of arrays and streams can be achieved by extension of our type family for arrays to:

$$[A]_\alpha = \text{Fin}(\alpha) \rightarrow A \quad \alpha : \text{Nat} + 1$$

where the right injection of the sum type contains ω and the definition of Fin is extended by $\text{Fin}(\omega) = \text{Nat}$. While this conceptually unites arrays and streams in the same way as the type $[A]$ unites lists and streams, we identify two main challenges that we address in this paper.

The first challenge arises from the fact that algebraic properties on finite structures often are lost when switching to the infinite setting. As an example consider some classical list properties: value-related properties such as $\text{map } f \circ \text{map } g = \text{map } (f \circ g)$ hold for lists and streams alike but properties that relate to the structure of lists such as $\text{drop } (\text{len } a) (a ++ b) = b$ typically only hold for (finite) lists; for streams, they break. While this loss of properties might be deemed acceptable in the context of list programming, in the context of array programming such structural properties play a very important role. Sophisticated array calculi have evolved around such properties such as Mullin's ψ -calculus [Mullin and Thibault 1994; Mullin 1988], Nial [Glasgow and Jenkins 1988] and the many APL-inspired array languages [Bernecky and Berry 1993; Breed et al. 1972; Hui and Iverson 1998]. Losing the generality of such properties for the sake of including streams would constitute an unacceptable loss. We tackle this issue by extending our type families for arrays further. We introduce the notion of *Transfinite Arrays* as we expand our type indices to countable ordinals:

$$[A]_\alpha = \text{Fin}(\alpha) \rightarrow A \quad \alpha : \text{Ord}$$

With this extension, we can resurrect most algebraic array properties for the infinite case.

The second challenge arises from the observation that transfinite arrays imply the existence of transfinite streaming, a concept that rarely considered in stream processing. We discuss what implications this extension has on classical streaming problems such as filtering and we propose solutions on how to deal with it.

The individual contributions of this paper are as follows:

- (1) We define an applied λ -calculus on finite arrays, its operational semantics and a type system for array operations. The calculus is a generic core language that implicitly supports several array calculi as well as compilation to highly efficient parallel code.
- (2) We expand the λ -calculus to support infinite arrays and show that the use of ordinals as indices enables a wide range of array-algebraic laws to carry over from the finite case to the infinite case.
- (3) We show that the proposed calculus also maintains many streaming properties even in the context of transfinite streaming.
- (4) We show that the proposed calculus inherently supports transfinite recursion. Several examples are contrasted to traditional list-based solutions.
- (5) We provide and describe a prototypical implementation¹. It demonstrates the viability of our semantics and it shows how the strict and finite fragment of the language can be mapped

¹The implementation is provided in the anonymous supplementary materials.

into high-performance code. We also provide a brief discussion on the opportunities and challenges involved when compiling the full language capabilities into efficient code.

We start with a description of the finite array calculus and naive extensions for infinite arrays in Section 2, before presenting the ordinal-based approach and its potential in Sections 3–5. Section 6 presents our prototypical implementation. Related work is discussed in Section 7; we conclude in Section 8.

2 EXTENDING ARRAYS TO INFINITY

We define an idealised, data-parallel array language, based on an applied λ -calculus that we call λ_α . The key aspect of λ_α is built-in support for shape- and rank-polymorphic array operations, similar to what is available in APL [Iverson 1962], J [Jsoftware 2016], or SAC [Grelck and Scholz 2006].

In the array programming community, it is well-known [Falster and Jenkins 1999; Jenkins and Mullin 1991] that basic design choices made in a language have an impact on the array algebras to which the language adheres. While we believe that our proposed approach is applicable within various array algebras, we chose one concrete setting for the context of this paper. We follow the *design decisions* of the functional array language SAC, which are compatible with many array languages, and which were taken directly from K.E. Iverson’s design of APL.

DD 1 *All expressions in λ_α are arrays.* Each array has a shape which defines how components within arrays can be selected.

DD 2 *Scalar expressions, such as constants or functions, are 0-dimensional objects with empty shape.* Note that this maintains the property that all arrays consist of as many elements as the product of their shape, since the product of an empty shape is defined through the neutral element of multiplication, *i.e.* the number 1.

DD 3 *Arrays are rectangular — the index space of every array forms a hyper-rectangle.* This allows the shape of an array to be defined by a single vector containing the element count for each axis of the given array.

DD 4 *Nested arrays that cater for inhomogeneous nesting are not supported. Homogeneously nested array expressions are considered isomorphic with non-nested higher-dimensional arrays.* Inhomogeneous nesting, in principle, can be supported by adding dual constructs for enclosing and disclosing an entire array into a singleton, and vice versa. **DD 2** implies that functions and function application can be used for this purpose.

DD 5 *λ_α supports infinitely many distinct empty arrays that differ only in their shapes.* In the definition of array calculi, the choice whether there is only one empty array or several has consequences on the universal equalities that hold. While a single empty array benefits value-focussed equalities, structural equalities require knowledge of array shapes, even when those arrays are empty. In this work, we assume an infinite number of empty arrays; any array with at least one shape element being 0 is empty. Empty arrays with different shape are considered distinct. For example, the empty arrays of shape [3, 0] and [0] are different arrays.

Further we describe the syntax and informal semantics of the language in Section 2.1 and we present types for the main array constructs in Section 2.2. Readers who feel more comfortable when explanation of the language starts with types can immediately refer to Section 2.2.

2.1 Syntax Definition and Informal Semantics of λ_α

We define the syntax of λ_α in Fig. 1. Its core is an untyped, applied λ -calculus. Besides scalar constants, variables, abstractions and applications, we introduce conditionals, a recursive let operator and some basic functions on the constants, including arithmetic operations such as +, -, *, /, a

148				
149	$c ::= 0, 1, \dots,$	(numbers)	\sim	
150	$true, false$	(booleans)	$reduce\ e\ e\ e$	(reduction)
151				
152	$e ::= c$	(constants)	$imap\ s$	$\begin{cases} g_1 : & e_1, \\ & \dots \\ g_n : & e_n \end{cases}$ (index map)
153	x	(variables)		
154	$\lambda x. e$	(abstractions)		
155	$e\ e$	(applications)		
156	$if\ e\ then\ e\ else\ e$	(conditionals)	$s ::= e$	(scalar imap)
157	$letrec\ x = e\ in\ e$	(recursive let)	$e e$	(generic imap)
158	$e + e, \dots$	(built-in binary)	$g ::= (e <= x < e)$	(index set)
159	$[e, \dots, e]$	(array constructor)	$-(x)$	(full index set)
160	$e.e$	(selections)		
161	$ e $	(shape operation)		
162	\sim			

Fig. 1. The syntax of λ_α

remainder operation denoted as %, and comparisons $<$, $<=$, $=$, etc. The actual support for arrays as envisioned by the aforementioned design principles is provided through five further constructs: array construction, selection, shape operation, *reduce* and *imap* combinators.

All arrays in λ_α are immutable. Arrays can be constructed by using potentially nested sequences of scalars in square brackets. For example, $[1, 2, 3, 4]$ denotes a four-element vector, while $[[1, 2], [3, 4]]$ denotes a two-by-two-element matrix. We require any such nesting to be homogeneous, for adherence to DD 4. For example, the term $[[1, 2], [3]]$ is irreducible, so does not constitute a value.

The dual of array construction is a built-in operation for element selection, denoted by a dot symbol, used as an infix binary operator between an array to select from, and a valid index into that array. A valid index is a vector containing as many elements as the array has dimensions; otherwise it is undefined.

$$[1, 2, 3, 4].[0] = 1 \quad [[1, 2], [3, 4]].[1, 1] = 4 \quad [[1, 2], [3, 4]].[1] = \perp$$

The third array-specific addition to λ_α is the primitive *shape* operation, denoted by enclosing vertical bars. It is applicable to arbitrary expressions, as demanded by DD 1, and it returns the shape of its argument as a vector, leveraging DD 3. For our running examples, we obtain: $|[1, 2, 3, 4]| = [4]$ and $|[[1, 2], [3, 4]]| = [2, 2]$. DD 5 and DD 2 imply that we have:

$$|[]| = [0] \quad |[[]]| = [1, 0] \quad |true| = [] \quad |42| = [] \quad |\lambda x.x| = []$$

λ_α includes a *reduce* combinator which in essence, it is a variant of *foldl*, extended to allow for multi-dimensional arrays instead of lists. *reduce* takes three arguments: the binary function, the neutral element and the array to reduce. For example, we have:

$$reduce\ (+)\ 0\ [[1, 2], [3, 4]] = (((0 + 1) + 2) + 3) + 4$$

assuming row-major traversal order. This allows for shape-polymorphic reductions such as:

$$sum \equiv \lambda a. \mathbf{reduce}\ (\lambda x. \lambda y. x + y)\ 0\ a \ ; \ \text{also works for scalars and empty arrays}$$

The final, and most elaborate, language construct is the *imap* (index map) construct. It bears some similarity to the classical map operation, but instead of mapping a function over the elements of an array, it constructs an array by mapping a function over all legal indices into the index space

denoted by a given shape expression². Added flexibility is obtained by supporting a piecewise definition of the function to be mapped. Syntactically, the *imap*-construct starts out with the keyword *imap*, followed by a description of the result shape (rule *s* in Fig. 1). The shape description is followed by a curly bracket that precedes the definition of the mapping function. This function can be defined piecewise by providing a set of index-range expression pairs. We demand that the set of index ranges constitutes a partitioning of the overall index space defined through the result shape expression, *i.e.* their union covers the entire index space and the index ranges are mutually disjoint. We refer to such index ranges as *generators* (rule *g* in Fig. 1), and we call a pair of a generator and its subsequent expression a *partition*. Each generator defines an index set and a variable (denoted by *x* in rule *g* in Fig. 1) which serves as the formal parameter of the function to be mapped over the index set. Generators can be defined in two ways: by means of two expressions which must evaluate to vectors of the same shape, constituting the lower and upper bounds of the index set, or by using the underscore notation which is syntactic sugar for the following expansion rule:

$$(\mathbf{imap} \ s \ \{ _ (iv) \ \dots \}) \equiv (\mathbf{imap} \ s \ \{ \underbrace{[0, \dots, 0]}_n \ \leq iv \ < \ s : \ \dots \})$$

assuming that $|s| = [n]$. The variable name of a generator can be referred to in the expression of the corresponding partition.

The \leq and $<$ operators in the generators can be seen as element-by-element array counterparts of the corresponding scalar operators which, jointly, specify sets of constraints on the indices described by the generators. As the index-bounds are vectors, we have:

$$v_1 \leq v_2 \implies |v_1|.0 = |v_2|.0 \wedge \forall 0 \leq i < |v_1|.0 : v_1.[i] \leq v_2.[i]$$

In the rest of the paper, we use the same element-wise extensions for scalar operators, denoting the non-scalar versions with dot on top: $c = a + b \implies c.i = a.i + b.i$. This often helps to simplify the notation³.

As an example of an *imap*, consider an element-wise increment of an array *a* of shape $[n]$. While a classical *map*-based definition can be expressed as *map* $(\lambda x.x + 1) \ a$, using *imap*, the same operation can be defined as:

```
imap [n] { [0] <= iv < [n]: a.iv + 1
```

Having mapping functions from indices to values rather than values to values adds to the flexibility of the construct. Arrays can be constructed from shape expressions without requiring an array of the same shape available:

```
imap [3,3] { [0,0] <= iv < [3,3]: iv.[0]*3 + iv.[1]
```

defines a 2-dimensional array $[[0, 1, 2], [3, 4, 5], [6, 7, 8]]$. Structural manipulations can be defined conveniently as well. Consider a *reverse* function, defined as follows:

```
reverse  $\equiv \lambda a.$  imap |a| { [0] <= iv < |a|: a.(|a| - iv - 1)
```

In order to express this with *map*, one needs to construct an intermediate array, where indices of *a* appear as values. Note also that the explicit shape of the *imap* construct makes it possible to define shape-polymorphic functions in a way similar to our definition of *reverse*. An element-wise increment for arbitrarily shaped arrays can be defined as:

```
increment  $\equiv \lambda a.$  imap |a| { \_(iv): a.iv + 1 ; also works for scalars & empty arrays
```

²For readers familiar with Haskell: the *imap* defined here derives the index space from a shape expression. It does not require an argument array of that shape.

³A formal definition of the extended operator is: $(\oplus) \equiv \lambda a.\lambda b.$ *imap* |a| { _(iv) : a.iv \oplus b.iv where $\oplus \in \{+, -, \dots\}$.

246 **DD 4** allows *imap* to be used for expressing operations in terms of n -dimensional sub-structures.
 247 All that is required for this is that the expressions on the right hand side of all partitions evaluate to
 248 non-scalar values. For example, matrices can be constructed from vectors. Consider the following
 249 expression:

```
250 imap [n] { [0] <= iv < [n]: [1,2,3,4] ; non-scalar partitions (incorrect attempt)
```

252 Its shape is $[n, 4]$; however, this shape no longer can be computed without knowing the shape of
 253 at least one element. If the overall result array is empty, its shape determination is a non-trivial
 254 problem. To avoid this situation, we require the programmer to specify the result shape by means
 255 of two shape expressions separated by a vertical bar: see the rule (generic *imap*) in Fig. 1. We refer
 256 to these two shape expressions as the *frame shape* which specifies the overall index range of the
 257 *imap* construct as well as the *cell shape* which defines the shape of all expressions at any given
 258 index. The concatenation of those two shapes is the overall shape of the resulting array. For more
 259 discussions related to the concepts of frame and cell shapes, see [Bernecky 1987, 1993; Bernecky
 260 and Iverson 1980]. The above *imap* expression therefore needs to be written as:

```
262 imap [n]||[4] { [0] <= iv < [n]: [1,2,3,4] ; non-scalar partitions (correct)
```

263 to be a legitimate expression of λ_α . The (scalar *imap*) case in Fig. 1, which we use predominantly in
 264 the paper, can be seen as syntactic sugar for the generic version, with the second expression being
 265 an empty vector.
 266

267 2.2 Towards a Type System for λ_α

268 We will present an outline of a type system here so that a reader could develop a better understanding
 269 of the essence of the array calculus that λ_α provides. For the sake of readability, we have taken
 270 some small liberties, like omitting definitions of standard arithmetic operations as well as standard
 271 non-array constructs.
 272

273 We use dependent types to specify array operations. First we define the types we will use as well
 274 as well-formedness criteria for array types.

NAT	BOOL	FIN	FUN
$\frac{}{\text{Nat} : \text{Type}}$	$\frac{}{\text{Bool} : \text{Type}}$	$\frac{n : \text{Nat}}{\text{Fin}(n) : \text{Type}}$	$\frac{A : \text{Type} \quad B : \text{Type}}{A \rightarrow B : \text{Type}}$
ARRAY			
$T : \text{Type}$	$T \notin \{\text{Array}\}$	$d : \text{Nat}$	$s : \text{Fin}(d) \rightarrow \text{Nat}$
$v : \left(\prod_i i : \text{Fin}(d). \text{Fin}(s\ i) \right) \rightarrow T$			
$\text{Array}(T, d, s, v) : \text{Type}$			

284 Nat is a type for natural numbers, Bool is a type for booleans, $\text{Fin}(n)$ is a type for numbers from 0
 285 to $n - 1$. Function types are standard. An array type is a quadruple, where the first element is a type
 286 of the base element. We prohibit T to be of array types, as according to **DD 4**, nested arrays are not
 287 supported. The second element of the tuple is the dimensionality of an array. We do not support
 288 nested arrays, but we support multi-dimensional arrays, so instead of having a type $[[A]_m]_n$ we
 289 have a type $[A]_{\langle n, m \rangle}$. Such a shape vector $\langle n, m \rangle$ is a third component of the tuple and it is modeled
 290 as a function from positions into vector components, e.g. $\{0 \rightarrow n, 1 \rightarrow m\}$ in our example. The last
 291 component of the tuple is a function type that maps an index vector type to a value type T . For
 292 each dimension $i : \text{Fin}(d)$ the corresponding index component has to be within the given shape,
 293 i.e. it has to be of type $\text{Fin}(s\ i)$.

294

The definitions of Nat and Fin are standard:

$$\begin{array}{c}
\text{NAT}_0 \\
\hline
0 : \text{Nat} \\
\text{NAT}_s \\
\hline
n : \text{Nat} \\
S n : \text{Nat} \\
\text{FIN-0} \\
\hline
n : \text{Nat} \\
\bar{0} : \text{Fin}(S n) \\
\text{FIN-S} \\
\hline
n : \text{Nat} \quad k : \text{Fin}(n) \\
\bar{S} k : \text{Fin}(S n)
\end{array}$$

We use \bar{x} notation to denote conversion from Nat to $\text{Fin}(x + 1)$:

$$x : \text{Nat} \implies \bar{x} : \text{Fin}(x + 1)$$

We use standard context $\Gamma ::= \cdot \mid \Gamma, x : A$, where $A : \text{Type}$. All the numbers in the language are natural numbers and the shape operation for any array of shape s returns a one-dimensional vector of Nats, with the content s :

$$\begin{array}{c}
\text{CONST} \\
\hline
\Gamma \vdash c : \text{Nat} \\
\text{SHAPE} \\
\hline
\Gamma \vdash a : \text{Array}(T, d, s, v) \\
\Gamma \vdash |a| : \text{Array}(\text{Nat}, 1, \lambda_d, \lambda\phi.s(\phi \bar{0}))
\end{array}$$

To construct a one-dimensional array using the bracket notation $[e_0, \dots, e_{n-1}]$ we ensure that all the elements have the same type, the shape vector of such an array is $\langle n \rangle$ — a single-element vector containing n . The value function of such an array is $\{\langle 0 \rangle \mapsto e_0, \langle 1 \rangle \mapsto e_1, \dots\}$ and we use a meta operator `packvec` to construct it.

$$\begin{array}{c}
\text{1D-ARR} \\
\hline
\forall 0 \leq i < n. \Gamma \vdash e_i : T \quad T \notin \{\text{Array}\} \\
\Gamma \vdash [e_0, \dots, e_{n-1}] : \text{Array}(T, 1, \lambda_n, \text{packvec } e_0 \dots e_{n-1}) \\
\text{packvec } e_0 \dots e_{n-1} = \\
\lambda\phi. \text{if } \phi \bar{0} = \bar{0} \text{ then } e_0 \\
\quad \text{else if } \phi \bar{0} = \bar{1} \text{ then } e_1 \\
\dots
\end{array}$$

To construct a $(d + 1)$ -dimensional array using n d -dimensional arrays we expect all the arrays to have the same dimensionality d and the same shape. Therefore we require d to be the same and we require s to be the same. By the latter we mean extensional equality. As s will be of a type $\text{Fin}(d) \rightarrow \text{Nat}$, such a check is decidable. Finally, we use the `pack` meta operator to create a value function for the resulting array.

$$\begin{array}{c}
\text{ND-ARR} \\
\hline
\forall 0 \leq i < n. \Gamma \vdash e_i : \text{Array}(T, d, s, v_i) \\
s_a \equiv \lambda i. \text{if } i = \bar{0} \text{ then } n \text{ else } s(i - \bar{1}) \\
\Gamma \vdash [e_0, \dots, e_{n-1}] : \text{Array}(T, d + 1, s_a, \text{pack } v_0 \dots v_{n-1}) \\
\text{pack } v_0 \dots v_{n-1} = \\
\lambda\phi. \text{if } \phi \bar{0} = \bar{0} \text{ then } \\
\quad v_0(\lambda i. \phi(i + \bar{1})) \\
\quad \text{else if } \phi \bar{0} = \bar{1} \text{ then } \\
\quad v_1(\lambda i. \phi(i + \bar{1})) \\
\dots
\end{array}$$

When selecting an element from a d -dimensional array, we have to provide an index which shall be a 1-dimensional array of Nats of d elements, where each element is bound by the shape of the array we are selecting from.

$$\begin{array}{c}
\text{SEL} \\
\hline
\Gamma \vdash a : \text{Array}(T, d, s_a, v_a) \\
\Gamma \vdash i : \text{Array}(\text{Nat}, 1, s_i, v_i) \quad \Gamma \vdash s_i \bar{0} = d \quad \forall 0 \leq j < d. \Gamma \vdash (v_i(\lambda_j)) < (s_a \bar{j}) \\
\Gamma \vdash a.i : T
\end{array}$$

The `imap` construct can be seen as a generalisation of the $[e_0, \dots]$ construct, a higher-order function that takes the shape of an array and a set of functions that generate elements for a given range of indices. We demonstrate the typing rule for the scalar `imap`, and we avoid the construction

of the value function of the resulting array, as such a construction is reflected in our semantics.

IMAP-SCAL

$$\frac{\begin{array}{l} \Gamma \vdash s : \text{Array}(\text{Nat}, 1, s_s, v_s) \\ \forall 1 \leq i \leq n. \Gamma \vdash l_i : \text{Array}(\text{Nat}, 1, s_s, _) \\ \forall 1 \leq i \leq n. \Gamma \vdash u_i : \text{Array}(\text{Nat}, 1, s_s, _) \\ \forall 1 \leq i \leq n. \Gamma, iv_i : \text{Array}(\text{Nat}, 1, s_s, _) \vdash e_i : T \quad T \notin \{\text{Array}\} \end{array}}{\Gamma \vdash \text{imap } s \left\{ \begin{array}{l} l_1 \leq iv_1 < u_1 : e_1, \\ \dots \\ l_n \leq iv_n < u_n : e_n, \end{array} \right. : \text{Array}(T, (s_s \bar{0}), \lambda i.v_s (\lambda _ . i), _)}$$

The rest of the typing rules for applications, abstractions, letrec and conditionals are standard, therefore we omit them here.

The type system presented here imposes a distinction between natural numbers and arrays of natural numbers of an empty shape. While this helps keeping the presentation reasonably compact this distinction is undesirable for λ_α from a pragmatical perspective. As most array calculi do, we want to consider scalars to be 0-dimensional arrays with empty shape. Amongst other benefits, this allows the function $\lambda a.\text{imap } |a| \{ _ (iv) : a.iv + 1 \}$ to be applied to regular arrays and scalars alike.

In the above type system, we can create an array of an empty shape: $\text{Array}(\text{Nat}, 0, \text{efq } \lambda \phi.5)$, where $\text{efq} : \text{Fin}(0) \rightarrow \text{Nat}$ (a function from empty type to Nat). The object of such a type will be isomorphic to $5 : \text{Nat}$, but not identical. This means that we will have to introduce explicit coercions not only between numbers of type Fin and Nat, but also between any non-array type T and an empty array of type T .

For the price of further type constructions, some of these equalities can be regained as shown in [Elsman and Dybdal 2014; Slepak et al. 2014; Trojahner and Grellck 2009]. Since this paper is mainly concerned with the calculus itself and its properties, we omit such elaboration. Instead, we assume that $T \equiv \text{Array}(T, 0, _ _)$, for non-array types T , and numbers of type Nat and Fin can be used interchangeably.

2.3 Formal Semantics of λ_α

In this section, we offer a brief overview of the semantics. A complete semantics can be found in [Anonymous-1 2018].

In λ_α , evaluated arrays are pairs of shape and element tuples. A shape tuple consists of numbers, and an element tuple consists of numbers, booleans or functions closures. We denote pairs and tuples, as well as element selection and concatenation on them, using the following notation:

$$\vec{a} = \langle a_1, \dots, a_n \rangle \implies \vec{a}_i = a_i \quad \langle a_1, \dots, a_n \rangle ++ \langle b_1, \dots, b_m \rangle = \langle a_1, \dots, a_n, b_1, \dots, b_m \rangle$$

To denote the product of a tuple of numbers, we use the following notation:

$$\vec{s} = \langle s_1, \dots, s_n \rangle \implies \otimes \vec{s} = s_n \cdot \dots \cdot s_1 \cdot 1$$

When a tuple is empty, its product is one. An array is rectangular, so its shape vector specifies the extent of each axis. The number of elements of each array is finite. Element vectors contain all the elements in a linearised form. While the reader can assume row-major order, formally, it suffices that a fixed linearisation function $F_{\vec{s}}$ exists which, given a shape vector $\vec{s} = \langle s_1, \dots, s_n \rangle$, is a bijection between indices $\{\langle 0, \dots, 0 \rangle, \dots, \langle s_1 - 1, \dots, s_n - 1 \rangle\}$ and offsets of the element vector: $\{1, \dots, \otimes \vec{s}\}$. Consider, as an example, the array $[[1, 2], [3, 4]]$, with F being row-major order. This array is evaluated into the shape-tuple element-tuple pair $\langle \langle 2, 2 \rangle, \langle 1, 2, 3, 4 \rangle \rangle$. Scalar constants are arrays with empty shapes. We have 5 evaluating to $\langle \langle \rangle, \langle 5 \rangle \rangle$. The same holds for booleans and function closures: *true* evaluates to $\langle \langle \rangle, \langle \text{true} \rangle \rangle$ and $\lambda x.e$ evaluates to $\langle \langle \rangle, \langle [\lambda x.e, \rho] \rangle \rangle$.

393 F is an invariant to the presented semantics. In finite cases, the usual choices of F are row-major
 394 order or column-major order. In infinite cases, this might be not the best option, and one could
 395 consider space-filling curves instead. F is only relevant for two operations; the creation of array
 396 values and the selection of elements from it. Selections relate the indices of the index vectors to the
 397 axes of the arrays following the order of nesting and starting with the index 0 on each level. We
 398 have: $[[1, 2], [3, 4]] [1, 0] = 3$, Assuming F is row-major, $F_{(2,2)}(\langle 1, 0 \rangle)$ equals 2 which, when used
 399 as index into $\langle\langle 2, 2 \rangle, \langle 1, 2, 3, 4 \rangle\rangle$ returns the intended result 3.

400 The inverse of F is denoted as $F_{\vec{s}}^{-1}$ and for every legal offset $\{1, \dots, \otimes \vec{s}\}$ it returns an index vector
 401 for that offset.

402
 403 *Deduction rules.* To define the operational semantics of λ_{α} , we use a *natural semantics*, similar to
 404 the one described in [Kahn 1987]. To make sharing more visible, instead of a single environment
 405 ρ that maps names to values, we introduce a concept of storage; environments map names to
 406 pointers and storage maps pointers to values. Environments are denoted by ρ and are ordered lists
 407 of name-pointer pairs. Storage is denoted by S and consists of an ordered list of pointer-value pairs.

408 Formally, we construct storage and environments as lists of pointer-value and variable-pointer
 409 bindings, respectively, using comma to denote extensions:

$$410 \quad S ::= \emptyset \mid S, p \mapsto v \quad \rho ::= \emptyset \mid \rho, x \mapsto p$$

411
 412 A look-up of a storage or an environment is performed *right to left* and is denoted as $S(p)$ and $\rho(x)$,
 413 respectively. Extensions are denoted with comma. Semantic judgements can take two forms:

$$414 \quad S; \rho \vdash e \Downarrow S'; p \quad S; \rho \vdash e \Downarrow S'; p \Rightarrow v$$

415
 416 where S and ρ are initial storage and environment and e is a λ_{α} expression to be evaluated. The
 417 result of this evaluation ends up in the storage S' and the pointer p points to it. The latter form of a
 418 judgement is a shortcut for: $S; \rho \vdash e \Downarrow S'; p \wedge S'(p) = v$.

419
 420 *Values.* The values in this semantics are constants (including arrays) and λ -closures which contain
 421 the λ term and the environment where this term shall be evaluated:

$$422 \quad \langle\langle \dots \rangle, \langle \dots \rangle\rangle \quad \langle\langle \rangle, \langle\langle \lambda x.e, \rho \rangle\rangle\rangle$$

423
 424 *Meta-operators.* Further in this section we use the following meta-operators:

425
 426 $\mathbf{E}(v)$ Lift the internal representation of a vector or a number into a valid λ_{α} expression. For
 427 example: $\mathbf{E}(5) = 5$, $\mathbf{E}(\langle 1, 2, 3 \rangle) = [1, 2, 3]$, etc.

428
 429 $\langle \vec{s}, _ \rangle$ We use underscore to omit the part of a data structure, when binding names. For example:
 430 $S; \rho \Rightarrow \langle \vec{s}, _ \rangle$ refers to binding the variable \vec{s} to the shape of $S(p)$ which must be a constant.

431 2.4 Core Rules

432 In λ_{α} , the rules for the λ -calculus core, *i.e.* constants, variables, abstractions and applications are
 433 straightforward adaptations of the standard rules for strict functional languages to our notation
 434 with storage and pointers:

435
 436
 437
 438
 439
 440
 441

442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490

$$\begin{array}{c}
 \text{CONST-SCAL} \\
 \frac{c \text{ is scalar}}{S; \rho \vdash c \Downarrow S_1, p \mapsto \langle \langle \rangle, \langle c \rangle \rangle; p} \\
 \\
 \text{VAR} \\
 \frac{x \in \rho \quad \rho(x) \in S}{S; \rho \vdash x \Downarrow S; \rho(x)} \\
 \\
 \text{APP} \\
 \frac{S; \rho \vdash e_1 \Downarrow S_1; p_1 \Rightarrow \langle \langle \rangle, \llbracket \lambda x. e, \rho_1 \rrbracket \rangle \quad S_1; \rho \vdash e_2 \Downarrow S_2; p_2 \quad S_2; \rho_1, x \mapsto p_2 \vdash e \Downarrow S_3; p_3}{S; \rho \vdash e_1 e_2 \Downarrow S_3; p_3} \\
 \\
 \text{ABS} \\
 \frac{S; \rho \vdash \lambda x. e \Downarrow S, p \mapsto \langle \langle \rangle, \llbracket \lambda x. e, \rho \rrbracket \rangle; p}{}
 \end{array}$$

As an illustration, consider the evaluation of $(\lambda x.x)$ 42:

$$\begin{array}{lll}
 \emptyset; \emptyset & (\lambda x.x) \ 42 & \text{ABS} \\
 S_1 = p_1 \mapsto \langle \langle \rangle, \llbracket \lambda x.x, \emptyset \rrbracket \rangle; \emptyset & p_1 \ 42 & \text{CONST-SCAL} \\
 S_2 = S_1, p_2 \mapsto \langle \langle \rangle, \langle 42 \rangle \rangle; \emptyset & p_1 \ p_2 & \text{APP} \\
 S_2; x \mapsto p_2 & x & \text{VAR} \\
 S_2; \emptyset & p_2 & \square
 \end{array}$$

We start with an empty storage and an empty environment. The outer application demands that the APP-rule be used. It enforces three computations: the evaluation of the function, the evaluation of the argument and the evaluation of the function body with an appropriately expanded environment. The function is evaluated by the ABS-rule which adds a closure $p_1 \mapsto \langle \langle \rangle, \llbracket \lambda x.x, \emptyset \rrbracket \rangle$ to the storage and returns the pointer p_1 to it. The argument is evaluated by the CONST-SCAL-rule which adds $p_2 \mapsto \langle \langle \rangle, \langle 42 \rangle \rangle$ to the storage and returns p_2 . Finally, the APP-rule demands the evaluation of the body of the function with an environment $\rho_1 = x \mapsto p_2$. The body being just the variable x , the VAR-rule gives us $S_2; p_2$ as final result.

The rules for array constructors and array selections are rather straightforward as well. Both these constructs are strict:

$$\begin{array}{c}
 \text{IMM-ARRAY} \\
 \frac{P = \langle p_1, \dots, p_n \rangle \quad n \geq 1 \quad \forall_{i=1}^n S_i; \rho \vdash c_i \Downarrow S_{i+1}; p_i \quad \text{AllSameShape}(S_{n+1}, P) \quad S' = S_{n+1}, p_o \mapsto \langle \langle 1 \rangle, \langle n \rangle \rangle, p_i \mapsto S_{n+1}(p_1) \quad S', \rho \vdash \text{imap}_{p_1} p_o | p_i \{ \langle i-1 \rangle \mapsto p_i \mid i \in \{1, \dots, n\} \} \Downarrow S''; p}{S_1; \rho \vdash [c_1, \dots, c_n] \Downarrow S''; p} \\
 \\
 \text{IMM-ARRAY-EMPTY} \\
 \frac{}{S; \rho \vdash [] \Downarrow S, p \mapsto \langle \langle 0 \rangle, \langle \rangle \rangle; p} \\
 \\
 \text{SEL-STRICT} \\
 \frac{S; \rho \vdash i \Downarrow S_1; p_i \Rightarrow \langle \langle d \rangle, \vec{i} \rangle \quad S_1; \rho \vdash a \Downarrow S_2; p_a \Rightarrow \langle \vec{s}, \vec{a} \rangle \quad k = F_{\vec{s}}(\vec{i})}{S; \rho \vdash a.i \Downarrow S_3, p \mapsto \langle \langle \rangle, \langle \vec{a}_k \rangle \rangle; p}
 \end{array}$$

Empty arrays are put into the storage with shape $[0]$ (IMM-ARRAY-EMPTY-rule). Non-empty arrays (IMM-ARRAY-rule) evaluate all the components and ensure that they are all of the same finite shape. Subsequently, we assemble evaluated components into the resulting array value ensuring that the flattening adheres to F . This is achieved by using an auxiliary term imap_1 . It takes the form $\text{imap}_1 p_o | p_i \{ \vec{i}^1 \mapsto p_{\vec{i}^1}, \dots, \vec{i}^n \mapsto p_{\vec{i}^n} \}$ where p_o and p_i are pointers to frame and cell shapes, and the set $\{ \vec{i}^1 \mapsto p_{\vec{i}^1}, \dots, \vec{i}^n \mapsto p_{\vec{i}^n} \}$ contains pairs of frame-shape indices and value pointers for all

$$\begin{array}{c}
491 \quad \text{IMAP-STRICT} \\
492 \quad S; \rho \vdash e_{\text{out}} \Downarrow S_1; p_{\text{out}} \Rightarrow \langle \langle d_o \rangle, s_{\text{out}} \rangle \quad S_1; \rho \vdash e_{\text{in}} \Downarrow S_2; p_{\text{in}} \Rightarrow \langle \langle d_i \rangle, s_{\text{in}} \rangle \\
493 \quad \hat{S}_1 = S_2 \quad \bigvee_{i=1}^n \hat{S}_i; \rho \vdash g_i \Downarrow \hat{S}_{i+1}; p_{g_i} \Rightarrow \bar{g}_i \quad \text{FormsPartition}(s_{\text{out}}, \{\bar{g}_1, \dots, \bar{g}_n\}) \\
494 \\
495 \quad \bar{S}_1 = \hat{S}_{n+1} \quad \forall (i, \vec{i}) \in \text{Enumerate}(s_{\text{out}}) \exists k : \left\{ \begin{array}{l} \vec{i} \in \bar{g}_k \wedge \bar{g}_k = \text{Gen}(x_k, _ _) \\ \bar{S}_i, p \mapsto \langle \langle d_o \rangle, \vec{i} \rangle; \rho, x_k \mapsto p \vdash e_k \Downarrow \bar{S}'_i; p_{\vec{i}} \\ \bar{S}'_i; \rho, x \mapsto p_{\vec{i}} \vdash |x| \Downarrow \bar{S}_{i+1}; p'_{\vec{i}} \Rightarrow \langle \langle d_i \rangle, s_{\text{in}} \rangle \end{array} \right. \\
496 \\
497 \quad \frac{\bar{S}_{\otimes s_{\text{out}}+1}, \rho \vdash \text{imap}_1 p_{\text{out}} | p_{\text{in}} \{ \vec{i} \mapsto p_{\vec{i}} \mid (_ _, \vec{i}) \in \text{Enumerate}(s_{\text{out}}) \} \Downarrow S'; p}{S; \rho \vdash \text{imap } e_{\text{out}} | e_{\text{in}} \left\{ \begin{array}{l} g_1 : e_1, \\ \dots \quad \Downarrow S'; p \\ g_n : e_n \end{array} \right.} \\
498 \\
499 \\
500 \quad \text{GEN} \\
501 \quad S; \rho \vdash e_1 \Downarrow S_1; p_1 \Rightarrow \langle \langle n \rangle, \vec{l} \rangle \quad S_1; \rho \vdash e_2 \Downarrow S_2; p_2 \Rightarrow \langle \langle n \rangle, \vec{u} \rangle \\
502 \\
503 \quad \frac{}{S; \rho \vdash (e_1 \leq x < e_2) \Downarrow S, p \mapsto \text{Gen}(x, \vec{l}, \vec{u}); p} \\
504 \\
505 \\
506 \\
507
\end{array}$$

legal indices into the frame shape. The formal definition of the deduction rule for imap_1 is provided in [Anonymous-1 2018, Sec 2.1.1].

The rule for selection (SEL-STRICT-rule) first evaluates the array we are selecting from, and the index vector specifying the array index we wish to select. Then, we compute the offset into the data vector by applying F to the index vector. Finally, we get the scalar value at the corresponding index. When applying F , we implicitly check that:

- the index is within bounds $1 \leq k \leq \otimes \vec{s}$, as $F_{\vec{s}}$ is undefined outside the index space bounded by \vec{s} ; and
- the index vector and the shape vector are of the same length, which means that selections evaluate scalars and not array sub-regions.


Imap. In order to keep the imap rule reasonably concise, we introduce two separate rules, a rule GEN for evaluating the generator bounds, and the main rule for imap , the IMAP-STRICT-Rule. The GEN-rule introduces auxiliary values $\text{Gen}(x, \vec{l}, \vec{u})$ which are triplets that keep a variable name, lower bound and upper bound of a generator together. These auxiliary values are references only by the rule for imap .

Evaluation of an imap happens in three steps. First, we compute shapes and generators, making sure that generators form a partition of s_{out} (FormsPartition is responsible for this). Secondly, for every valid index defined by the frame shape (Enumerate generates a set of offset-index-vector pairs), we find a generator that includes the given index (denoted $\vec{i} \in \bar{g}_k$). We evaluate the generator expression e_k , binding the generator variable x_k to the corresponding index value and check that the result has the same shape as p_{in} . Finally, we combine evaluated expressions for every index of the frame shape into imap_1 for further extraction of scalar values.

All missing rules, including built-in operations, conditionals and recursion through the *letrec*-construct are straightforward adaptations of the standard rules. They can be found in [Anonymous-1 2018]. Formal definitions of helper functions, such as AllSameShape, will also be found there.

2.5 Infinite Arrays

In order to support infinite arrays, we introduce the notion of infinity in λ_α , and we allow infinities to appear in shape components. Syntactically, this can be achieved by adding a symbol for infinity, as shown in Fig. 2. For disambiguation, we refer to the extended version of λ_α as λ_α^∞ . Adding ∞ has

Fig. 2. The syntax of λ_α^∞

several implications. First of all, our built-in arithmetic needs to be extended. We treat infinity in the usual way, applying the model commonly known as a Riemann sphere. That is:

$$z + \infty = \infty \qquad z \times \infty = \infty \qquad \frac{z}{\infty} = 0 \qquad \frac{z}{0} = \infty$$

The following operations are undefined:

$$\infty + \infty \qquad \infty - \infty \qquad \infty \times 0 \qquad \frac{0}{0} \qquad \frac{\infty}{\infty}$$

While these additions to the semantics are trivial, allowing infinity to appear in shapes has a more profound impact on our semantics. Our rule for *imap*-constructs (IMAP-STRICT) forces the evaluation of all elements. If our result shape contains infinity, this can no longer be done. As we want to maintain a strict evaluation regime for function applications in general, we turn our *imap*-construct into a lazy data-structure which does not immediately compute its elements, but only does so when individual elements are being inspected. For this purpose, we extend our set of allowed values of our semantics with an *imap*-closure:

$$\left[\left[\text{imap } p_{\text{out}} | p_{\text{in}} \left\{ \begin{array}{l} \bar{g}_1 : e_1, \\ \dots \\ \bar{g}_n : e_n \end{array} \right. , \rho \right] \right]$$

The *imap* closure contains pointers to frame and element shapes (p_{out} and p_{in} correspondingly), the list of partitions, where generators have been evaluated and the environment in which the *imap* shall be evaluated. The overall idea is to update, in place, this closure whenever individual elements are computed. With this extension, we can now replace our strict *imap*-rule by a lazy variant:

$$\begin{array}{c}
 \text{IMAP-LAZY} \\
 S; \rho \vdash e_{\text{out}} \Downarrow S_1; p_{\text{out}} \Rightarrow \langle \langle _ \rangle, s_{\text{out}} \rangle \quad S_1; \rho \vdash e_{\text{in}} \Downarrow S_2; p_{\text{in}} \Rightarrow \langle \langle _ \rangle, _ \rangle \\
 \hat{S}_1 = S_2 \quad \forall_{i=1}^n \hat{S}_i; \rho \vdash g_i \Downarrow \hat{S}_{i+1}; p_{g_i} \Rightarrow \bar{g}_i \quad \text{FormsPartition}(s_{\text{out}}, \{\bar{g}_1, \dots, \bar{g}_n\},) \\
 \hline
 S; \rho \vdash \text{imap } e_{\text{out}} | e_{\text{in}} \left\{ \begin{array}{l} g_1 : e_1, \\ \dots \\ g_n : e_n \end{array} \right. \Downarrow \hat{S}_{n+1}, p \mapsto \left[\left[\text{imap } p_{\text{out}} | p_{\text{in}} \left\{ \begin{array}{l} \bar{g}_1 : e_1, \\ \dots \\ \bar{g}_n : e_n \end{array} \right. ; \rho \right] \right]
 \end{array}$$

We can see that the new rule for *imap*-constructs, in essence, performs a subset of what the strict rule from the previous section does. It still forces the result shapes, it still computes the boundaries of the generators, and it checks the validity of the overall generator set. Once these computations have been done, further element computation is delayed and an *imap*-closure is created instead.

The actual computation of elements is triggered upon element selection. Consequently, we need a second selection rule which can deal with *imap* closures in the array argument position:

SEL-LAZY-IMAP

$$\begin{array}{c}
 589 \\
 590 \\
 591 \\
 592 \\
 593 \\
 594 \\
 595 \\
 596 \\
 597 \\
 598 \\
 599 \\
 600 \\
 601 \\
 602 \\
 603 \\
 604 \\
 605 \\
 606 \\
 607 \\
 608 \\
 609 \\
 610 \\
 611 \\
 612 \\
 613 \\
 614 \\
 615 \\
 616 \\
 617 \\
 618 \\
 619 \\
 620 \\
 621 \\
 622 \\
 623 \\
 624 \\
 625 \\
 626 \\
 627 \\
 628 \\
 629 \\
 630 \\
 631 \\
 632 \\
 633 \\
 634 \\
 635 \\
 636 \\
 637
 \end{array}$$

$$\begin{array}{c}
 S; \rho \vdash i \Downarrow S_1; p_i \Rightarrow \langle \langle _ \rangle, \vec{v} \rangle \quad S_1; \rho \vdash a \Downarrow S_2; p_a \Rightarrow \left[\text{imap } p_{\text{out}} | p_{\text{in}} \begin{array}{l} \vec{g}_1 \quad e_1 \\ \dots \\ \vec{g}_n \quad e_n \end{array}, \rho' \right] \\
 S_2(p_{\text{out}}) = \langle \langle m \rangle, _ \rangle \quad (\vec{i}, \vec{j}) = \text{Split}(m, \vec{v}) \\
 \exists k : \vec{i} \in \vec{g}_k \quad \vec{g}_k = \text{Gen}(x_k, _ , _) \quad S_2, p \mapsto \mathbf{E}(\vec{i}); \rho', x_k \mapsto p \vdash e_k \Downarrow S_3; p_{\vec{i}} \\
 S_3; \rho', x \mapsto p_{\vec{i}} \vdash x.\mathbf{E}(\vec{j}) \Downarrow S_4; p \quad S_5 = \text{UpdateIMap}(S_4, p_a, \vec{i}, p_{\vec{i}}) \\
 \hline
 S; \rho \vdash a.i \Downarrow S_5; p
 \end{array}$$

Selections into *imap*-closures happen at indices that are of the same length as the concatenation of the *imap* frame and cell shapes. This means that the index the *imap*-closure is being selected from has to be split into frame and cell sub-indices: \vec{i} and \vec{j} correspondingly. Given that \vec{g}_k contains \vec{i} , we evaluate e_k with x_k being bound to \vec{i} . As this value may be non-scalar, we evaluate a selection into it at \vec{j} . Finally, the evaluated generator expression is saved within the *imap* closure. This step is performed by the helper function `UpdateIMap`, which splits the k -th partition into a single-element partition containing \vec{i} with the computed value $p_{\vec{i}}$, and further partitions covering the remaining indices of \vec{g}_k with the expression e_k . For more details see [Anonymous-1 2018, Sec. 2.1.1].

With this, we can define and use infinite arrays in an overall strict setting. Let us consider the definitions of the infinite array of natural numbers in λ_α^∞ on the left and Haskell-like definition on the right:

$$\text{nats} \equiv \mathbf{imap} \ [\infty] \ \{ _(\text{iv}) : \text{iv}.[0] \quad \text{nats} = 0 : \mathbf{map} \ (+1) \ \text{nats}$$

Both versions define an object that delivers the value n when being selected at any index n . Both definitions provide a data structure whose computation unfolds in a lazy fashion. The main difference is that the Haskell-like specification introduces dependencies between the elements of the list. Arguably, for a large number of practical implementations, whenever an element n is selected, the entire spine of the list, up to the n -th element, has to be in place. In the λ_α^∞ case, the specification explicitly states how to compute the element at any position: the underscore in the *imap* is similar to the λ -binder. Therefore, we encode less dependencies, which means that space-efficient implementation of *imap* closures can be derived with less analysis. For example, we can envision representing *imap* closures as a hashmap.

The above comparison demonstrates important difference between a data-parallel programming style and a list-based, inherently recursive programming style. This observation leads us to the question whether similar recursive definitions are possible in λ_α^∞ at all?

2.6 Recursive Definitions

It turns out that the lazy *imap*, together with the *letrec* construct, allows for recursive definitions of arrays. A recursive definition of the natural numbers, including 0, can be defined in λ_α^∞ by:

$$\mathbf{letrec} \ \text{nats} = \mathbf{imap} \ [\infty] \ \{ [0] \Leftarrow \text{iv} < [1] : 0, \\
 [1] \Leftarrow \text{iv} < [\infty] : \text{nats} . (\text{iv} \dot{-} [1]) + 1 \ \mathbf{in} \ \text{nats}$$

The interesting question here is whether the semantics defined thus far ensures that all elements of the array `nats` are actually being inserted into one and the same *imap*-closure. For this to happen, we need the environment of the *imap*-closure to map `nats` to itself, and we need the selection within the body of the *imap* to modify the closure from which it is selecting. While the latter is given

through the SEL-LAZY-IMAP-rule, the former is achieved through the rule for letrec-constructs. For λ_α , we have:

$$\text{LETREC} \quad \frac{S_1 = S, p \mapsto \perp \quad \rho_1 = \rho, x \mapsto p \quad S_1; \rho_1 \vdash e_1 \Downarrow S_2; p_2 \quad S_3 = S_2[p_2/p] \quad S_3; \rho, x \mapsto p_2 \vdash e_2 \Downarrow S_4; p_r}{S; \rho \vdash \text{letrec } x = e_1 \text{ in } e_2 \Downarrow S_4; p_r}$$

where $S[p_2/p]$ denotes substitution of the $x \mapsto p$ bindings inside of the enclosed environments with $x \mapsto p_2$, where x is any legal variable name. This substitution is key for creating the circular reference in the *imap*-closure from the example above.

In conclusion, the above recursive specification denotes an array with the same elements as the data-parallel specification from the previous section. In contrast to data-parallel version, this specification behaves much more like the recursive, Haskell-like version; the computation of individual elements can no longer happen directly. Since there is an encoded dependency between an element and its predecessor, the first access to an element at index n , in this variant, will trigger the computation of all elements from 0 up to n . The implementation of the UpdateIMap operation on *imap*-closures determines how these numbers are stored in memory and, consequently, how efficiently they can be accessed.

The availability of direct indexes makes it possible to encode an arbitrary order for the recursion. Consider the following example:

```
letrec a = imap [10] { [9] <= iv < [10]: 9,
                      [0] <= iv < [9]: a.(iv+[1])-1 in a
```

Selection of the 9th element can be evaluated in one step. In case of lists, the selection request always starts at the beginning of the list. Hence, to obtain the same performance, some optimisation of the list case is required.

2.7 List Primitives in the Array Setting

We have enabled two features that are inherent with lists, but that are usually not supported in an array setting: recursively defined data-structures and infinite arrays. All that is required to achieve this is a recursion-aware, lazy semantics of the *imap*-construct and the inclusion of an explicit notion of infinity. With these extensions, the key primitives for lists, *head*, *tail*, and *cons* can be defined as

```
head ≡ λa.a.[0]
tail ≡ λa.imap |a|÷[1] { _(iv): a.([1]+iv)
cons ≡ λa.λb.imap [1]+|b| { [0] <= iv < [1]: a,
                          [1] <= iv < [1]+|b|: b.(iv÷[1])
```

More complex list-like functions can be defined on top of these. An example is concatenation:

```
letrec (++) = λa.λb.if |a|. [0] = 0 then b
                else cons (head a) ((tail a) ++ b) in (++)
```

In case a is infinite, however, the above definition of concatenation is unsatisfying. The strict nature of λ_α will force *tail a* forever as $|a|. [0] = 0$ never yields *true*. The way to avoid this is to shift the case distinction into the lazy *imap* construct:

```
(++) ≡ λa.λb.imap |a|+|b| { [0] <= iv < |a|: a.iv,
                          |a| <= iv < |a|+|b|: b.(iv÷|a|)
```

As we have seen earlier, λ_α enables the typical constructions of recursive definitions of infinite vectors well-known from the realm of lists such as list of ones, natural numbers or fibonacci sequence.

687 Having a unified interface for arrays and lists enables programmers to switch the algorithmic
 688 definitions of individual arrays from recursive to data-parallel styles without modifying any of the
 689 code that operates on them.

690 However, such a unification comes at a price: we have to support a lazy version of the *imap*-
 691 construct. As a consequence, we conceptually lose the advantage of $O(1)$ access. Despite λ_α offering
 692 many opportunities for compiler optimisations like pre-allocating arrays and potentially enforcing
 693 strictness on finite, non-recursive *imaps*, one may wonder at this point how much λ_α differs from a
 694 lazy array interface in a lazy, list-based language such as Haskell?

695 3 TRANSFINITE ARRAYS

697 We now investigate to what extent λ_α^∞ adheres to the key properties of array programming – array
 698 algebras and array equalities.

699 3.1 Algebraic Properties

701 Array-based operations offer a number of beneficial algebraic properties. Typically, these properties
 702 manifest themselves as universally valid equalities which, once established, improve our thinking
 703 about algorithms and their implementations, and give rise to high-level program transformations.
 704 We define equality between two non-scalar arrays a and b as

$$705 \quad a == b \iff |a| = |b| \wedge \forall iv < |a| : a.iv = b.iv$$

706 that is, we demand equality of the shapes and equality of all elements. The demand for equality of
 707 shapes recursively implies equality in dimensionality and the extensional character of this definition
 708 through the use of array selections ensures that we can reason about equality on infinite arrays as
 709 well.
 710

711 Arrays give rise to many algebras such as Theory of Arrays [More 1973], Mathematics of
 712 Arrays [Mullin 1988], and Array Algebras [Glasgow and Jenkins 1988]. Most of the developed
 713 algebras differ only slightly, and the set of equalities that are ultimately valid depends on some
 714 fundamental choices, such as the ones we made in the beginning of the previous section. At the
 715 core of these equalities is the ability to change the shape of arrays in a systematic way without
 716 losing any of their data.

717 An equality from [Falster and Jenkins 1999] that plays a key role in consistent shape manipulations
 718 is:

$$719 \quad \mathit{reshape} \ |a| \ (\mathit{flatten} \ a) == a \quad (1)$$

720 where *flatten* maps an array recursively into a vector by *concatenating* its sub-arrays in a row-major
 721 fashion and *reshape* performs the dual operation of bringing a row-major linearisation back into
 722 multi-dimensional form. These operations can be defined in λ_α^∞ as

$$723 \quad \begin{aligned} 724 \quad \mathit{flatten} &\equiv \lambda a. \mathbf{imap} \ [\mathit{count} \ a] \ \{ _ (iv) : a.(o2i \ iv \ .[0] \ |a|) \\ 725 \quad \mathit{reshape} &\equiv \lambda \mathit{shp} . \lambda a. \mathbf{imap} \ \mathit{shp} \ \{ _ (iv) : (\mathit{flatten} \ a).[i2o \ iv \ \mathit{shp}] \end{aligned}$$

726 where *count* returns the product of all shape components and *o2i* and *i2o* translate offsets into
 727 indices and vice versa, respectively. These operations effectively implement conversions from
 728 mixed-radix systems into natural numbers using multiplications and additions and back using
 729 division and remainder operations.

730 The above equality states that any array a can be brought into flattened form and, subsequently
 731 be brought back to its original shape. For arrays of finite shape s , this follows directly from the fact
 732 that $o2i \ (i2o \ iv \ s) = iv$ for all legitimate index vectors iv into the shape s .

733 If we want Eq. 1 to hold for all arrays in λ_α^∞ , we need to show that the above equality also holds
 734 for arrays with infinite axes. Consider an array of shape $s = [2, \infty]$. For any legal index vector $[1, n]$
 735

736 into the shape s , we obtain:

$$\begin{aligned}
 737 \quad & o2i (i2o [1, n] [2, \infty]) [2, \infty] = o2i (\infty \cdot 1 + n) [2, \infty] \\
 738 \quad & = o2i \infty [2, \infty] \\
 739 \quad & = [\infty / \infty, \infty \% \infty] \\
 740 \quad & = [\infty / \infty, \infty \% \infty]
 \end{aligned}$$

741 which is not defined. We can also observe that all indices $[1, n]$ are effectively mapped into the
 742 same offset: ∞ which is not a legitimate index into any array in λ_α^∞ . This reflects the intuition that
 743 the concatenation of two infinite vectors effectively loses access to the second vector.

744 The inability to concatenate infinite arrays also makes the following equality fail:

$$745 \quad \text{drop } |a| (a ++ b) == b \quad (2)$$

746 where a and b are vectors and $\text{drop } s x$ removes first s elements from the left. The reason is exactly
 747 the same: given that $|a| = [\infty]$ and b is of finite shape $[n]$, the shape of the concatenation is
 748 $[\infty + n] = [\infty]$, and drop of $|a|$ results in an empty vector.

749 Clearly, λ_α^∞ as presented so far is not strong enough to maintain universal equalities such as Eq. 1
 750 or 2. Instead, we have to find a way that enables us to represent sequences of infinite sequences
 751 that can be distinguished from each other.

752 3.2 Ordinals

753 When numbers are treated in terms of cardinality, they describe the number of elements in a set.
 754 Addition of two cardinal numbers a and b is defined as a size of a union of sets of a and b elements.
 755 This notion also makes it possible to operate with infinite numbers, where the number of elements
 756 in an infinite set is defined via bijections. As a result, differently constructed infinite sets may end
 757 up having the same number of elements. For example, if there exists a bijection from $\mathbb{N} \times \mathbb{N}$ into \mathbb{N} ,
 758 the cardinality of both sets is the same.

759 When studying arrays, treating their shapes and indices using cardinal numbers is an over-
 760 simplification, because arrays have richer structure. Arrays are collections of ordered elements,
 761 where the order is established by the indices. Ordinal numbers, as introduced by G. Cantor in 1883,
 762 serve exactly this purpose — to “label” positions of objects within an ordered collection. When
 763 collections are finite, cardinals and ordinals can be used interchangeably, as we can always count
 764 the labels. Infinite collections are quite different in that regard: despite being of the same size, there
 765 can be many non-isomorphic well-orderings of an infinite collection. For example, consider two
 766 infinite arrays of shapes $[\infty, 2]$ and $[2, \infty]$. Both of these have infinitely many elements, but they
 767 differ in their structure. From a row major perspective, the former is an infinite sequence of pairs,
 768 whereas the latter are two infinite sequences of scalars. Ordinals give a formal way of describing
 769 such different well-orderings.

770 First let us try to develop an intuition for the concept of ordinal numbers and then we give a
 771 formal definition. Consider an ordered sequence of natural numbers: $0 < 1 < 2 < \dots$. These are
 772 the first ordinals. Then, we introduce a number called ω that represents the limit of the above
 773 sequence: $0 < 1 < 2 < \dots < \omega$. Further, we can construct numbers beyond ω by putting a “copy”
 774 of natural numbers “beyond” ω :

$$775 \quad 0 < 1 < 2 < \dots < \omega < \omega + 1 < \omega + 2 < \dots < \omega + \omega$$

776 For the time being, we treat operations such as $\omega + n$ symbolically. The number $\omega + \omega$ which can
 777 be also denoted as $\omega \cdot 2$ is the second limit ordinal that limits any number of the form $\omega + n$, $n \in \mathbb{N}$.
 778 Such a procedure of constructing limit ordinals out of already constructed smaller ordinals can be
 779 applied recursively. Consider a sequence of $\omega \cdot n$ numbers and its limit:

$$780 \quad 0 < \omega < \omega \cdot 2 < \omega \cdot 3 < \dots < (\omega \cdot \omega = \omega^2)$$

and we can carry on further to ω^n , ω^ω , etc. Note though that in the interval from ω^2 to ω^3 we have infinitely many limit ordinals of the form:

$$\omega^2 < \omega^2 + \omega < \omega^2 + \omega \cdot 2 < \dots < \omega^3$$

and between any two of these we have a “copy” of the natural numbers:

$$\omega^2 + \omega < \omega^2 + \omega + 1 < \dots < \omega^2 + \omega \cdot 2$$

3.2.1 Formal definitions. A totally ordered set $\langle A, < \rangle$ is said to be well ordered if and only if every nonempty subset of A has a least element [Ciesielski 1997]. Given a well-ordered set $\langle X, < \rangle$ and $a \in X$, $X_a \stackrel{\text{def}}{=} \{x \in X \mid x < a\}$. An ordinal is a well-ordered set $\langle X, < \rangle$, such that: $\forall a \in X : a = X_a$. As a consequence, if $\langle X, < \rangle$ is an ordinal then $<$ is equivalent to \in . Given a well-ordered set $A = \langle X, < \rangle$ we define an ordinal that this set is isomorphic to as $Ord(A, <)$. Given an ordinal α , its successor is defined to be $\alpha \cup \{\alpha\}$. The minimal ordinal is \emptyset which is denoted with 0. The next few ordinals are:

$$\begin{aligned} 1 &= \{0\} &= \{\emptyset\} \\ 2 &= \{0, 1\} &= \{\emptyset, \{\emptyset\}\} \\ 3 &= \{0, 1, 2\} &= \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\} \\ &\dots & \end{aligned}$$

A limit ordinal is an ordinal that is greater than zero that is not a successor. The set of natural numbers $\{0, 1, 2, 3, \dots\}$ is the smallest limit ordinal that is denoted ω . We use *islim* x to denote that x is a limit ordinal.

3.2.2 Arithmetic on Ordinals.

Addition. Ordinal addition is defined as $\alpha + \beta = Ord(A, <_A)$, where $A = \{0\} \times \alpha \cup \{1\} \times \beta$ and $<_A$ is the lexicographic ordering on A . Ordinal addition is associative but *not* commutative. As an example consider expressions $2 + \omega$ and $\omega + 2$. The former can be seen as follows: $0 < 1 < 0' < 1' < \dots$, which after relabeling is isomorphic to ω . However, the latter can be seen as: $0 < 1 < \dots < 0' < 1'$, which has the largest element $1'$, whereas ω does not. Therefore $2 + \omega = \omega < \omega + 2$. We have used $0'$, $1'$ to indicate the right hand side argument of the addition.

Subtraction. Ordinal subtraction can be defined in two ways, as partial inverse of the addition on the left and on the right. For left subtraction, which will be used by default throughout this paper unless otherwise specified, $\alpha - \beta$ is defined when $\beta \leq \alpha$, as: $\exists \xi : \beta + \xi = \alpha$. As ordinal addition is left-cancelative ($\alpha + \beta = \alpha + \gamma \implies \beta = \gamma$), left subtraction always exists and it is unique.

Right subtraction is a bit harder to define as:

- it is not unique: $1 + \omega = 2 + \omega$ but $1 \neq 2$; therefore $\omega -_R \omega$ can be any number that is less than ω : $\{0, 1, 2, \dots\}$.
- even if $\beta < \alpha$, the difference $\alpha - \beta$ might not exist. For example: $42 < \omega$; however, $\omega -_R 42$ does not exist as $\nexists \xi : \xi + 42 = \omega$.

Despite those difficulties, right subtraction can be useful at times and can be defined for $\alpha -_R \beta$:

$$\min\{\xi : \xi + \beta = \alpha\}$$

Multiplication. Ordinal multiplication $\alpha \cdot \beta = Ord(A, <_A)$ where $A = \alpha \times \beta$ and $<_A$ is the lexicographic ordering on A . Multiplication is associative and left-distributive to addition:

$$\alpha \cdot (\beta + \gamma) = (\alpha \cdot \beta) + (\alpha \cdot \gamma)$$

However, multiplication is not commutative and is not distributive on the right: $2 \cdot \omega = \omega < \omega \cdot 2$ and $(\omega + 1) \cdot \omega = \omega \cdot \omega < \omega \cdot \omega + \omega$.

834 *Exponentiation.* Exponentiation can be defined using transfinite recursion: $\alpha^0 = 1, \alpha^{\beta+1} = \alpha^\beta \cdot \alpha$
 835 and for limit ordinals λ : $\alpha^\lambda = \bigcup_{0 < \xi < \lambda} \alpha^\xi$.
 836

837 *ϵ -ordinals.* Using ordinal operations we can construct the following hierarchy of ordinals:
 838 $0, 1, \dots, \omega, \omega + 1, \dots, \omega \cdot 2, \omega \cdot 2 + 1, \dots, \omega^2, \dots, \omega^3, \dots, \omega^\omega, \dots$. The smallest ordinal for which $\alpha = \omega^\alpha$
 839 is called ϵ_0 . It can also be seen as a limit of the following $\omega^\omega, \omega^{\omega^\omega}, \dots, \omega^{\omega^{\dots}}$.
 840

841 **3.2.3 Cantor Normal Form.** For every ordinal $\alpha < \epsilon_0$ there are unique $n, p < \omega, \alpha_1 > \alpha_2 > \dots >$
 842 α_n and $x_1, \dots, x_n \in \omega \setminus \{0\}$ such that $\alpha > \alpha_1$ and $\alpha = \omega^{\alpha_1} \cdot x_1 + \dots + \omega^{\alpha_n} \cdot x_n + p$. Cantor Normal
 843 Form makes provides a standardized way of writing ordinals. It uniquely represents each ordinal
 844 as a finite sum of ordinal powers, and can be seen as an ω based polynomial. This can be used as a
 845 basis for an efficient implementation of ordinals and their operations.

846 3.3 λ_ω : Adding Ordinals to λ_α

847 The key contribution of this paper is the introduction of λ_ω , a variant of λ_α , which use ordinals
 848 as shapes and indices of arrays and which reestablishes global equalities in the context of infinite
 849 arrays.
 850

851 Before revisiting the equalities, we look at the changes to λ_α that are required to support
 852 transfinite arrays. Syntactically, to introduce ordinals in the language, we make two minor additions
 853 to λ_α . Firstly, we add ordinals⁴ as scalar constants. Secondly, we add a built-in operation, *islim*,
 854 which takes one argument and returns *true* if the argument is a limit ordinal and *false* otherwise.
 855 For example: *islim* ω reduces to *true* and *islim* $(\omega + 21)$ reduces to *false*.

λ_α with ordinals	extends λ_α				
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="padding: 5px;">$e ::= \dots$</td> <td style="padding: 5px;"> <i>islim</i> (limit ordinal predicate)</td> </tr> <tr> <td style="padding: 5px;">$c ::= \dots$</td> <td style="padding: 5px;"> $\omega, \omega + 1, \dots$ (ordinals)</td> </tr> </table>		$e ::= \dots$	<i>islim</i> (limit ordinal predicate)	$c ::= \dots$	$\omega, \omega + 1, \dots$ (ordinals)
$e ::= \dots$	<i>islim</i> (limit ordinal predicate)				
$c ::= \dots$	$\omega, \omega + 1, \dots$ (ordinals)				

862 Fig. 3. The syntax of λ_ω .

863 Semantically, it turns out that all core rules can be kept unmodified apart from the aspect that all
 864 helper functions, arithmetic, and relational operations now need to be able to deal with ordinals
 865 instead of natural numbers. In particular, the semantic for lazy *imaps* as developed for λ_α^∞ can be
 866 used unaltered, provided that all helper functions involved such as for splitting generators *etc.* are
 867 expanded to cope with ordinals.
 868

871 3.4 Array Equalities Revisited

872 With the support of Ordinals in λ_ω , we can now revisit our equalities Eq. 1 and 2. Let us first look
 873 at the counter example for Eq. 1: from Section 3.1: With an array shape $s = [2, \omega]$ and a legal index
 874 vector into $s [1, n]$, we now obtain:
 875

$$\begin{aligned}
 876 \quad o2i \ (i2o \ [1, n] \ [2, \omega]) \ [2, \omega]) &= o2i \ (\omega + n) \ [2, \omega] \\
 877 &= [(\omega + n) / \omega, (\omega + n) \% \omega] \\
 878 &= [1, n] \\
 879
 \end{aligned}$$

880 ⁴Technically, we support ordinal values only up to ω^ω , as ordinals are constructed using the constant ω and $+$, $-$, $*$, $/$ and
 881 $\%$ operations (no built-in ordinal exponentiation).
 882

883 The crucial difference to the situation from λ_α^∞ in Section 3.1 here is the ability to divide $(\omega + n)$
 884 by ω and to obtain a remainder, denoted by $\%$, of that division as well. By means of induction over
 885 the length of the shape and index vectors this equality can be proven to hold for arbitrary shapes
 886 in λ_ω , and, based on this proof, Eq. 1 can be shown as well.

887 In the same way as the arithmetic on ordinals is key to the proof of Eq. 1, it also enables the
 888 proof of Eq. 2 for arbitrary ordinal-shaped vectors⁵ a and b , with the definition of $++$ from the
 889 previous section and $drop$ being defined as:

890 $drop \equiv \lambda s. \lambda a. \mathbf{imap} \ |a| \dot{-} s \ \{ \ [0] \ \leq \ iv \ < \ |a| \dot{-} s : \ a.(s+iv)$

891 After inlining $++$ and $drop$, the left hand side of Eq. 2 can be rewritten as:

892 $\mathbf{letrec} \ ab = \mathbf{imap} \ |a| \dot{+} |b| \ \{ \ [0] \ \leq \ jv \ < \ |a| : \ a.jv ,$
 893 $\quad \quad \quad |a| \ \leq \ jv \ < \ |a| \dot{+} |b| : \ b.(jv \dot{-} |a|) \ \mathbf{in}$
 894 $\mathbf{imap} \ |ab| \dot{-} |a| \ \{ \ [0] \ \leq \ iv \ < \ |ab| \dot{-} |a| : \ ab.(|a|+iv)$

896 Consider the shape of the goal expression of the $letrec$. According to the semantics of the shape of
 897 an $imap$, we get: $|ab| \dot{-} |a|$. The shape of ab is $|a| \dot{+} |b|$. According to ordinal arithmetic: $(|a| \dot{+} |b|) \dot{-} |a|$
 898 is $|b|$. Therefore the shapes of right-hand and left-hand sides of the goal expressions are the same.

899 Let us rewrite the last $imap$ as:

900 $\mathbf{imap} \ |b| \ \{ \ [0] \ \leq \ iv \ < \ |b| : \ ab.(|a|+iv)$

901 Consider now selections into ab . All the selections into ab will happen at indices that are greater
 902 than a . This is because all the legal iv in the $imap$ are from the range $[0]$ to $|b|$.

903 According to the semantics of selections into $imaps$, $ab.(|a|+iv)$ will select from the second
 904 partition of the $imap$ that defines ab , and will evaluate to: $b.((|a|+iv) \dot{-} |a|)$. According to ordinal
 905 arithmetic, $(|a|+iv) \dot{-} |a|$ is identical to iv , therefore we can rewrite the previous $imap$ as:

906 $\mathbf{imap} \ |b| \ \{ \ [0] \ \leq \ iv \ < \ |b| : \ b.iv$

907 As it can be seen, this is an identity $imap$, which is extensionally equivalent to b .

910 4 EXAMPLES

911 *Transfinite tail.* As explained in Section 3.3, the shift from natural numbers to ordinals as indices
 912 in λ_ω implies corresponding extensions of the built-in arithmetic operations. As these operations
 913 lose key properties, such as commutativity, once arguments exceed the range of natural numbers,
 914 we need to ensure that function definitions for finite arrays extend correctly to the transfinite case.

915 As an example, consider the definition of $tail$ from the previous section:

916 $tail \equiv \lambda a. \mathbf{imap} \ |a| \dot{-} [1] \ \{ \ _ (iv) : \ a.([1] \dot{+} iv)$

917 For the case of finite vectors, we can see that a vector shortened by one element is returned, where
 918 the first element is missing and all elements have been shifted to the left by one element.

919 Let us assume we apply $tail$ to an array a with $|a| = [\omega]$. The arithmetic on ordinals gives us a
 920 return shape of $[\omega] \dot{-} [1] = [\omega]$. That is, the tail of an infinite array is the same size as the array
 921 itself, which matches our common intuition when applying $tail$ to infinite lists. The elements of that
 922 infinite list are those of a , shifted by one element to the right, which, again, matches our expected
 923 interpretation for lists.

924 Now, assume we have $|a| = [\omega + 42]$, which means that $(tail \ a).[\omega]$ should be a valid expression.
 925 For the result shape of $tail \ a$, we obtain $[\omega + 42] \dot{-} [1] = [\omega + 42]$. A selection $(tail \ a).[\omega]$ evaluates
 926 to $a.([1] \dot{+} [\omega]) = a.[\omega]$. This means that the above definition of the $tail$ shifts all the elements
 927 at indices smaller than $[\omega]$ one left, and leaves all the other unmodified. While this may seem

928 ⁵Eq. 2 can be generalised and shown to hold in the multi-dimensional case, provided that $++$ and $drop$ operate over the
 929 same axis.

counter-intuitive at first, it actually only means that *tail* can be applied infinitely often but will never be able to reach “beyond” the first limit.

Finally, observe that the body of the *imap*-construct in the definition of *tail* uses $[1]+iv$ is an index expression, not $iv+[1]$. In the latter case, the tail function would behave differently beyond $[\omega]$: it would attempt to shift elements to the left. However, this would make the overall definition faulty. Consider again the case when $|a| = [\omega + 42]$: the shape of the result would be $|a|$, which would mean that it would be possible to index at position $[\omega + 41]$, triggering evaluation of $a.([\omega + 41]+[1])$ and consequently, producing an *index error*, or out-of-bounds access into a .

Zip. Let us now define *zip* of two vectors that produces a vector of tuples. Consider a Haskell definition of *zip* function first:

```
zip (a : as) (b : bs) = (a, b) : zip as bs
zip _         _       = []
```

The result is computed lazily, and the length of the resulting list is a minimum of the lengths of the arguments. Like concatenation, a literal translation into λ_ω is possible, but it has the same drawbacks, *i.e.* it is restricted to arrays whose shape has no components bigger than ω .

A better version of *zip* that can be applied to arbitrary transfinite arrays looks as follows:

```
zip  $\equiv$   $\lambda a . \lambda b . \mathbf{imap}$  (min |a| |b|)|[2] {_(iv): [a.iv, b.iv]}
```

Here, we use a constant array in the body of the *imap*. This forces evaluation of both arguments, even if only one of them is selected. This can be changed by replacing the constant array with an *imap*:

```
zip  $\equiv$   $\lambda a . \lambda b . \mathbf{imap}$  (min |a| |b|)|[2] {_(iv):  $\mathbf{imap}$  [2] { [0] <= jv < [1] a.iv ,
[1] <= jv < [2] b.iv
```

which can be fused in a single *imap* as follows:

```
zip  $\equiv$   $\lambda a . \lambda b . \mathbf{letrec}$  s = (min |a| |b|).[0]  $\mathbf{in}$ 
 $\mathbf{imap}$  [s, 2] { [0, 0] <= iv < [s, 1]: a.[iv].[0] },
[0, 1] <= iv < [s, 2]: b.[iv].[0]}
```

Data Layout and Transpose. A typical transformations in stream programming is changing the granularity of a stream and joining multiple streams. In λ_ω , these transformations can be expressed by manipulating the shape of an infinite array. Consider changing the granularity of a stream a of shape $[\omega]$ into a stream of pairs:

```
 $\mathbf{imap}$  (|a|/[2])|[2] {_(iv): [a.[2*iv].[0], a.[2*iv].[0]+1]}
```

or we can express the same code in a more generic fashion:

```
( $\lambda n . \mathbf{reshape}$  ((|a|/[n])++[n]) a) 2
```

This code can operate on the streams of transfinite length, as well. If we envision compiling such a program into machine code, the infinite dimension of an array can be seen as a time-loop, and the operations at the inner dimension seen as a stream-transforming function. Such granularity changes are often essential for making good use of (parallel) hardware resources, *e.g.* FPGAs.

Transposing a stream makes it possible to introduce synchronisation. Consider transforming a stream a of shape $[2, \omega]$ into a stream of pairs (shape $[\omega, 2]$):

```
 $\mathbf{imap}$  [ $\omega$ ]|[2] {_(iv): [a.[iv].[0], 0], a.[iv].[0], 1]}
```

Conceptually, an array of shape $[2, \omega]$ represents two infinite streams that reside in the same data structure. An operation on such a data structure can progress independently on each stream, unless some dependencies on the outer index are introduced. A transpose, as above, makes it possible to introduce such a dependency, ensuring that the operations on both streams are synchronized. A

key to achieving this is the ability to re-enumerate infinite structures, and ordinal-based infinite arrays make this possible.

Ackermann function. The true power of multidimensional infinite arrays manifests itself in definitions of non-primitive-recursive sequences as data. Consider the Ackermann function, defined as a multi-dimensional stream:

```

letrec a = imap [ω, ω] {_(iv): letrec m = iv.[0] in
                                letrec n = iv.[1] in
                                if m = 0 then n + 1
                                else if m > 0 and n = 0 then a.[m-1, 1]
                                else a.[m-1, a.[m,n-1]] in a

```

Such a treatment of multi-dimensional infinite structures enables simple transliteration of recursive relations *as data*. Achieving similar recursive definitions when using cons-lists is possible, but they have a subtle difference. Consider a Haskell definition of the Ackermann function in data:

```

a = [[ if m == 0 then n+1
      else if m > 0 then a !! (m-1) !! 1
      else a !! (m-1) !! (a !! m !! (n-1))
     | n <- [0..]]
  | m <- [0..]]

```

We use two [0..] generators for explicit indexing, even though at runtime, all necessary elements of the list will be present. The lack of explicit indexes forces one to use extra objects to encode the correct dependencies, essentially implementing *imap* in Haskell. Conceptually, these generators constitute two further locally recursive data structures. Whether they can be always can be optimised away is not clear. Avoiding these structures in an algorithmic specification can be a major challenge.

Game of Life. As a final example, consider Conway's Game of Life which describes an evolution of cells on a plane. The most interesting aspect of this example is the fact that we can encode it in λ_ω in such a way that the shape of the plane is never specified. This means that the program can operate with infinite planes, *e.g.* of shape $[\omega, \omega]$, as well as finite 2d planes with no changes to source code.

First we introduce a few generic helper functions:

```

(or) ≡ λa.λb.if a then a else b
(and) ≡ λa.λb.if a then b else a
any ≡ λa.reduce or false a
gen ≡ λs.λv.imap s {_(iv): v
↖ ≡ λv.λa.imap |a| {_(iv): if any (iv+v >= |a|) then 0 else a.(iv+v)
↘ ≡ λv.λa.imap |a| {_(iv): if any (iv < v) then 0 else a.(iv-v)

```

or and and encode logical conjunction and disjunction, respectively. any folds an array of boolean expressions with the disjunction, and gen defines an array of shape *s* whose values are all identical to *v*. More interesting are the functions \swarrow and \searrow . Given a vector *v* and an array *a*, they shift all elements of *a* towards decreasing indices or increasing indices by *v* elements, respectively. Missing elements are treated as the value 0.

Now, we define a single step of the 2-dimensional Game of Life in APL style⁶: two-dimensional array *a* by:

```

gol_step ≡ λa.
letrec F = [↖ [1,1], ↖ [1,0], ↖ [0,1], λ x. ↖ [1,0] (↘ [0,1] x),
           ↘ [0,1], ↘ [1,0], ↘ [1,1], λ x. ↘ [1,0] (↖ [0,1] x)]
in letrec
    c = (reduce (λf.λg.λx.f x + g x) (λx.gen |a| 0) F) a
in

```

⁶See this video by John Scholes for more details: <https://youtu.be/a9xAKttWgP4>

```

1030     imap |a| { _(iv): if (c.iv = 2 and a.iv = 1) or (c.iv = 3)
1031                 then 1
1032                 else 0

```

1033 We assume an encoding of a live cell in a to be 1, and a dead cell to be 0. The array F contains
 1034 partial applications of the two shift functions to two-element vectors so that shifts into all possible
 1035 directions are present. The actual counting of live cells is performed by a function which folds F
 1036 with the function $\lambda f.\lambda g.\lambda x.f\ x + g\ x$. This produces c , an array of the same shape as a , holding the
 1037 numbers of live cells surrounding each position. Defining the shift operations \swarrow and \searrow to insert 0
 1038 ensures that all cells beyond the shape of a are assumed to be dead.

1039 The definition of the result array is, therefore, a straightforward *imap*, implementing the rules of
 1040 birth, survival and death of the Game of Life.

1041

1042 5 TRANSFINITE ARRAYS VS. STREAMS

1043 Streams have attracted a lot of attention due to the many algebraic properties they expose. [Hinze
 1044 2010] provides a nice collection of examples, many of which are based on the observation that
 1045 streams form an applicative functor. Transfinite arrays are applicative functors as well, not only for
 1046 arrays of shape $[\omega]$, but also for any given shape *shp*. With definitions:

```

1047     pure  $\equiv \lambda x.$  imap shp {_(iv): x
1048     ( $\diamond$ )  $\equiv \lambda a.\lambda b.$  imap shp {_(iv): a.iv b.iv

```

1049 we obtain for arbitrary arrays u, v, w , and x of shape *shp*:

$$\begin{aligned}
 1051 \quad & (\text{pure } \lambda x.x) \diamond u == u & (\text{pure } (\lambda f.\lambda g.\lambda x.f\ (g\ x))) \diamond u \diamond v \diamond w == u \diamond (v \diamond w) \\
 1052 \quad & (\text{pure } f) \diamond (\text{pure } x) == \text{pure } (f\ x) & u \diamond (\text{pure } x) == (\text{pure } (\lambda f.f\ x)) \diamond u
 \end{aligned}$$

1054 This shows that arbitrarily shaped arrays of finite size have this property, as also shown by [Gib-
 1055 boms 2017], and that these properties can be expanded into ordinal-shaped arrays. Classical streams
 1056 are a special instance of these, *i.e.* arrays of shape $[\omega]$.

1057 For stream operations that insert or delete elements, it is less obvious whether these can be
 1058 easily extended into ordinal-shaped arrays other than shape $[\omega]$. As an example, let us consider the
 1059 function *filter*, which takes a predicate p and a vector v and returns a vector that contains only
 1060 those elements x of v that satisfy $(p\ x)$. A direct definition of *filter* can be given as:

```

1062     filter  $\equiv \lambda p.$   $\lambda v.$  if (p v.[0]) then v.[0] ++ filter p (tail v)
1063                 else filter p (tail v)

```

1064 This definition, in principle, is applicable to arrays of any ordinal shape, but the use of *tail* in the
 1065 recursive calls inhibits application beyond ω . Furthermore, the strict semantics of λ_ω inhibits a
 1066 terminating application to any infinite array, including arrays of shape $[\omega]$. For the same reason, a
 1067 definition of *filter* through the built-in *reduce* is restricted to finite arrays.

1068 To achieve possible termination of the above definition of *filter* for transfinite arrays, we would
 1069 need to change to a lazy regime for all function applications in λ_ω and we would need to change
 1070 the semantics of *imap* into a variant where the shape computation can be delayed as well. Even if
 1071 that would be done, we would still end up with an unsatisfying solution. The filtering effect would
 1072 always be restricted to the elements before the first limit ordinal ω . This limitation breaks several
 1073 fundamental properties, like those defined in [Bird 1987], that hold in the finite and stream cases.
 1074 As an example, consider distributivity of *filter* over concatenation:

$$1075 \quad \text{filter } p\ (a\ ++\ b) == (\text{filter } p\ a) ++ (\text{filter } p\ b) \tag{3}$$

1077 This property holds for finite arrays, but fails with the above definition of *filter* in case a is infinite.

1078

To regain this property for transfinite arrays, we need to apply *filter* to all elements of the argument array, not only those before the first limit ordinal ω . When doing this in the context of λ_ω , the necessity to have a strict shape for every object forces us to “guess” the shape of the filtered result in advance. The way we “guess” has an impact on the filter-based equalities that will hold universally.

In this paper we propose a scheme that respects the above equality. For finite arrays *filter* works as usual, and for the infinite ones, we postulate that the result of filtering will be of an infinite-shape:

$$\forall p \forall a : |a| \geq \omega \implies |\text{filter } p \ a| \geq \omega$$

This is further applied to all infinite sequences contained within the given shape as follows:

$$\forall i < |a| : (\exists \text{islim } \alpha : i < \alpha \leq |a|) \implies (\exists k \in \mathbb{N} : p (a.(i+k)) = \text{true})$$

We assume that each infinite sequence contains infinitely many elements for which the predicate holds. Consequently, any limit ordinal component of the shape of the argument is carried over to the result shape and only any potential finite rest undergoes potential shortening. Consider a filter operation, applied to a vector of shape $[\omega \cdot 2]$. Following the above rationale, the shape of the result will be $[\omega \cdot 2]$ as well. This means that the result of applying *filter* to such an expression should allow indexing from $\{0, 1, \dots\}$ as well as from $\{\omega, \omega + 1, \dots\}$ delivering meaningful results.

This decision can lead to non-termination when there are only finitely many elements in the filtered result. For example:

```
filter ( $\lambda x. x > 0$ ) (imap  $[\omega+2]$   $\{_(iv) : 0\}$ )
```

reduces to an array of shape $[\omega]$, which effectively is empty. Any selection into it will lead to a non-terminating recursion.

The overall scheme may be counter-intuitive, but it states that for every index position of the output, the computation of the corresponding value is well-defined.

Assuming the aforementioned behaviour of *filter*, Eq. 3 holds for all transfinite arrays. Another universal equation that holds for all transfinite vectors concerns the interplay of *filter* and *map*:

$$\text{filter } p \ (\text{map } f \ a) == \text{map } f \ (\text{filter } (p \cdot f) \ a)$$

The proposed approach does not only respect the above equalities, but it also behaves similarly to filtering of streams that can be found in languages such as Haskell: *filter* applied to an infinite stream cannot return a finite result.

In principle, the chosen filtering scheme can be defined in λ_ω by using the *islim* predicate within an *imap*. However, the resulting definition is neither concise, nor likely to be runtime efficient. Given the importance of *filter*, we propose an extension of λ_ω . Fig. 4 shows the syntactical extension of λ_ω .




Fig. 4. The syntax of λ_ω with filters.

As *filter* conceptually is an alternative means of constructing arrays, its semantics is similar to that of *imap*. In particular, it constitutes a lazy array constructor, whose elements are being evaluated upon demand created through selections. Technically, this means that we have to introduce a new value to keep *filter*-closures, a rule that builds such a closure from *filter* expression, and we need to define the selection operation that forces evaluation within the filter closure.

We introduce as new value for *filter*-closures:

$$\left[\left[\text{filter } p_f \ p_e \ \begin{cases} \alpha_1 & v_r^1 \ v_i^1 \\ \dots \\ \alpha_n & v_r^n \ v_i^n \end{cases} \right] \right]$$

which contains the pointer to the filtering function p_f , the shape of the argument we are filtering over (p_e) and the list of partitions that consist of a limit ordinal, and a pair of partial result and natural number: v_r and v_i correspondingly.

On every selection at index $[\xi + n]$, where ξ is a limit ordinal or zero, and n is a natural number, we find a ξ partition within the filter closure or add a new one if it is not there. Every partition keeps a vector with a partial result of filtering (v_r), and the index (v_i) with the following property: the element in the array we are filtering over at position $\xi + (v_i - 1)$ is the last element in the v_r , given that $v_r > 0$. This means that if n is within v_r , we return $v_r.[n]$. Otherwise, we extend v_r until its length becomes $n + 1$ using the following procedure: inspect the element in p_e at the position $\xi + v_i$ – if the predicate function evaluates to *true*, append this element to v_r and increase v_i by one, otherwise, increase v_i by one.

A formal description of this procedure can be found in [Anonymous-1 2018, Sec. 2.1.4].

6 IMPLEMENTATION

We implement λ_ω in a system called Heh, which can be found in the anonymous supplementary materials. Heh contains:

- (1) an interpreter for λ_ω covering the full language, and
- (2) a compiler for the strict and finite subset of λ_ω .

The interpreter can be seen as a proof of concept that the proposed semantics is implementable. The implementation is an almost literal translation of the semantic rules provided in the paper into Ocaml code. We carefully implement updates in-place for *imap* and *filter* closures, ensuring that these constructs are evaluated lazily rather than in normal order. All examples provided in the paper can be found in that repository, and run, correctly, in Heh.

Compilation of the finite subset of λ_ω is achieved by translating λ_ω programs into SAC programs and subsequently using the compiler *sac2c* to produce binaries. Multi-core and GPU backends of *sac2c* can be leveraged to execute strict and finite λ_ω programs in parallel on these types of architectures. The Heh implementation comes with more than a 100 unit tests for its internal components.

In the interpreter, ordinals are represented by their Cantor Normal Form. The algorithms for implementing operations on ordinals are based on [Manolios and Vroon 2005]. In the same paper, we also find an in-depth study of the complexities of ordinal operations: comparisons, additions and subtractions have complexities $O(n)$, where n is the minimum of the lengths of both arguments; multiplications have the complexity $O(n \cdot m)$, where m and n are the lengths of the two argument representations.

6.1 Performance considerations

Our compiler for the strict and finite sublanguage of λ_ω shows that this part of the language can be mapped into languages such as SAC, leading to high-performance execution potential on various platforms [Šinkarovs et al. 2013; Wieser et al. 2012]. Whether the full-fledged version of λ_ω can be compiled into high-performance codes as well, mainly relies on the answers to two questions:

- (1) how can we handle finite expressions that are defined by means of recursive *imaps*, and
- (2) what is the most efficient representation for transfinite arrays.

1177 *Recursive imap*s. Strict data parallel languages like SAC rarely support recursive *imap* constructs,
 1178 even if the shape of the result is finite. There are two difficulties: (i) the evaluation of recursive
 1179 *imaps* results in the necessity to support *imap* closures; (ii) parallel implementation of a recurs-
 1180 ive *imap* becomes trickier because of potential dependencies between the elements of an array.
 1181 In [Anonymous-2 2018] we propose an elegant solution to this problem. We introduce a mechanism
 1182 that switches from strict to lazy evaluation of a potentially recursive *imap*. It is demonstrated that
 1183 the lifetime of *imap* closures is kept to a minimum and that a parallel implementation is possible.
 1184 Furthermore, the proposed solution enables the detection of cyclic array definitions that diverge
 1185 under strict semantics.

1186 *Data structures*. The current semantics prescribes that, when evaluating selections into a lazy
 1187 *imap*, the partition that contains the index that is to be selected from has to be split into a single-
 1188 element partition and the remainder. This means that, as the number of selections into the *imap*
 1189 increases, the structure that stores partitions of the *imap* will have to deal with a large number
 1190 of single-element arrays. Partitions can be stored in a tree, providing $O(\log n)$ look-up; however
 1191 triggering a memory allocation for every scalar is likely to be very inefficient. An alternate approach
 1192 would be to allocate larger chunks, each of which would store a subregion of the index space of
 1193 an *imap*. When doing so, we would need to establish a policy on the size of chunks and chose
 1194 a mechanism on how to indicate evaluated elements in a chunk. Another possibility would be
 1195 to combine the chunking with some strictness speculation, using a technique similar to the one
 1196 presented in [Anonymous-2 2018]. That way, a single element selection could trigger the evaluation
 1197 of an entire chunk.

1199 *Memory management*. An efficient memory management model is not obvious. In case of strict
 1200 arrays, reference counting is known to be an efficient solution [Cann 1989; Grellck and Scholz 2006].
 1201 For lazy data structures, garbage collection is usually preferable. Most likely, the answer lies in a
 1202 combination of those two techniques.

1203 The *imap* construct offers an opportunity for garbage collection at the level of partitions. Consider
 1204 a lazy *imap* of boolean values with a partition that has a constant expression:

```
1205 imap [ $\omega$ ] { ... , 1 <= iv < u: false , ...
```

1206 Assume further that neighbouring partitions evaluate to *false*. In this case, we can merge the
 1207 boundaries of partitions and instead of keeping values in memory, the partition can be treated as a
 1208 generator. However, an efficient implementation of such a technique is non-trivial.

1210 *Ordinals*. An efficient implementation of ordinals and their operations is also essential. Here, we
 1211 could make use of the fact that λ_ω is limited to ordinals up to ω^ω . For further details see [Anonymous-
 1212 1 2018, Sec. 4]

1214 7 RELATED WORK

1215 Several works propose to extend the index domain of arrays to increase expressibility of a language.
 1216 A straightforward way to do this is to stay within cardinal numbers but add a notion of ∞ , similarly
 1217 to what we have proposed in λ_α^∞ . Similar approach is described in [McDonnell and Shallit 1980]; in
 1218 J [Jsoftware 2016] infinity is supported as a value, but infinite arrays are not allowed. As we have
 1219 seen, by doing so we lose a number of array equalities.

1220 In [More 1973, page 137] we read: ‘A restriction of indices to the finite ordinal numbers is a
 1221 needless limitation that obscures the essential process of counting and indexing.’ We cannot agree
 1222 more. [More 1973] describes an axiomatic array theory that combines set theory and APL. The
 1223 theory is self contained and gives rise to a number of array equalities. However, the question on
 1224 how this theory can be implemented (if at all) is not discussed.

1225

1226 In [Taylor 1982] the authors propose to extend the domain of array indices with real numbers.
 1227 More specifically, a real-valued function gives rise to an array in which valid indices are those that
 1228 belong to the domain of that function. The authors investigate expressibility of such arrays and
 1229 they identify classes of problems where this could be useful, but neither provide a full theory nor
 1230 discuss any implementation-related details.

1231 Besides the related work that stems from APL and the plethora of array languages that evolved
 1232 from it, there is an even larger body of work that has its origins in lists and streams. One of the
 1233 best-known fundamental works on the theory of lists using ordered pairs can be found in [McCarthy
 1234 1960, sec. 3], where a class of S-expressions is defined. The concepts of *nil* and *cons* are introduced,
 1235 as well as *car* and *cdr*, for accessing the constituents of *cons*.

1236 The Theory of Lists [Bird 1987] defines lists abstractly as linearly ordered collections of data. The
 1237 empty list and operations like length of the list, concatenation, filter, map and reduce are introduced
 1238 axiomatically. Lists are assumed to be finite. The questions of representation of this data structure
 1239 in memory, or strictness of evaluation, are not discussed.

1240 Concrete Stream Calculus [Hinze 2010] introduces streams as codata. Streams are similar to
 1241 McCarthy’s definition of lists, in that they have functions *head* and *tail*, but they lack *nil*. This
 1242 requires streams to be infinite structures only. The calculus is presented within Haskell, rendering
 1243 all evaluation lazy.

1244 Coinduction and codata are the usual way to introduce infinite data structures in programming
 1245 languages [Jeannin et al. 2012; Kozen and Silva 2016]. Key to the introduction of codata typically is
 1246 the use of coinductive semantics [Leroy and Grall 2009]. In our paper, the use of ordinals keeps
 1247 the semantics inductive and deals with infinite objects by means of ordinals. In [Turner 1995], the
 1248 author investigates a model of a total functional language, in which codata is used to define infinite
 1249 data objects.

1250 Streams are also related to dataflow models, such as [Estrin and Turn 1963; Kahn 1974; Petri 1962].
 1251 The computation graphs in the latter can be seen as recursive expressions on potentially infinite
 1252 streams. As demonstrated in [Beck et al. 2015], there is a demand to consider multi-dimensional
 1253 infinite streams that cache their parts for better efficiency.

1254 Two array representations, called *push arrays* and *pull arrays*, are presented in [Svensson and
 1255 Svenningsson 2014]. Pull arrays are treated as objects that have a length and an index-mapping
 1256 function; push arrays are structures that keep sequences of element-wise updates. The *imap* defined
 1257 here can be considered an advanced version of a pull array, with partitions and transfinite shape.
 1258 The availability of partitions circumvents a number of inefficiencies, (e.g. embedded conditionals)
 1259 of classical pull arrays; the ordinals, in the context of the *imap*-construct, enable the expression of
 1260 streaming algorithms.

1261 The #Id language, presented in [Heller 1989], is similar to λ_ω ; It combines the idea of lazy data
 1262 structures with an eager execution context.

1263 In [Atkey and McBride 2013; Møgelberg 2014], the authors propose a system that makes it
 1264 possible to reason whether a computation defined on an infinite stream is productive⁷ — a question
 1265 that can be transferred directly to λ_ω . Their technique lies in the introduction of a clock abstraction
 1266 which limits the number of operations that can be made before a value must be returned. This
 1267 approach has some analogies with defining explicit “windows” on arrays, as for example proposed
 1268 in [Hammes et al. 1999], or guarantees that programs run in constant space in [Lippmeier et al.
 1269 2016].

1270 One of the key features of the array language described in this paper is the availability of strict
 1271 shape for any expression of the language. Combining this with updates in place, which can be
 1272

1273 ⁷The computation will eventually produce the next item, i.e. it is not stuck.
 1274

1275 achieved by means of monads [Wadler 1995], uniqueness typing [Barendsen and Smetsers 1996] or
1276 reference counting [Grelck and Scholz 2006], very efficient code generation becomes possible.

1277 Strict shapes can be encoded in types as well. Specifically in the dependently-typed system, such
1278 an approach can be very powerful. The work on container theory [Abbott et al. 2005] allows a very
1279 generic description of indexed objects capturing ideas of shapes and indices in types. A very similar
1280 idea in the context of arrays is described in [Gibbons 2017]. The work on dependent type systems
1281 for array languages include [Slepek et al. 2014; Trojahner and Grelck 2009; Xi and Pfenning 1998].
1282 Finally, a way to extend a type theory to include the notion of ordinals can be found in [Hancock
1283 2000].

1284 8 CONCLUSIONS AND FUTURE WORK

1286 This paper proposes *transfinite arrays* as a basis for an applied λ -calculus λ_ω . The distinctive
1287 feature of transfinite arrays is their ability to capture arrays with infinitely many elements, while
1288 maintaining structure within that infiniteness. The number of axes is preserved, and individual
1289 axes can contain infinitely many infinite subsequences of elements. This capability extends many
1290 structural properties that hold for finite arrays into the transfinite space.

1291 The embedding of transfinite arrays into λ_ω allows for recursive array definitions, offering an
1292 opportunity to transliterate typical list-based algorithms, including algorithms on infinite lists
1293 for stream processing, into a generic array-based form. The paper presents several examples to
1294 this effect, and provides some efficiency considerations for them. It remains to be seen if these
1295 considerations, in practice, enable a truly unified view of arrays, lists, and streams.

1296 The array-based setting of λ_ω allows this recursive style of defining infinite structures to be taken
1297 into a multi-dimensional context, enabling elegant specification of inherently multi-dimensional
1298 problems on infinite arrays. As an example, we present an implementation of Conway's Game of
1299 Life which, despite looking very similar to a formulation for finite arrays, is defined for positive
1300 infinities on both axes. Within λ_ω , accessing neighbouring elements along both axes can be specified
1301 without requiring traversals of nested cons lists.

1302 We also present an implementation for the Ackerman function, using a 2-dimensional transfinite
1303 array, one axis per parameter. The resulting code adheres closely to the abstract declarative formu-
1304 lation of the function, while also implicitly generating a basis for a memoising implementation of
1305 the algorithm.

1306 An interesting aspect of transfinite arrays is that ordinal-based indexing opens up an avenue
1307 to express transfinite induction in data in very much the same way as *nil* and *cons* are duals to
1308 the principle of mathematical induction. This can not be done easily using *cons* lists as there is no
1309 concept of a limit ordinal in that context. It may be possible to encode this principle by means of
1310 nesting, but then one would need a type system or some sort of annotations to distinguish lists of
1311 transfinite length from nested lists. The *imap* construct from the proposed formalism can be seen
1312 as an elegant solution to this.

1313 The fact that *imap* supports random access and is powerful enough to capture list and stream
1314 expressions alike opens up an exciting perspective for the implementation of λ_ω . When arrays
1315 are finite, it is possible to reuse one of the existing efficient array-based implementations. When
1316 arrays are infinite, we can use list or stream implementations to encode λ_ω , but at the same time
1317 the properties of the original λ_ω programs open the door to rich program analysis and alternative
1318 representations. We believe that many functional languages striving for performance could benefit
1319 from the proposed design, at least when the distinction between finite and infinite arrays can be
1320 statically determined be it through annotation or inference.

1321 The concept of transfinite arrays as proposed in this paper offers several new and interesting
1322 possibilities for further investigation. As discussed in the implementation section, it is not yet
1323

1324 clear what the most efficient implementations for our proposed infinite structures are. Choices of
 1325 representation affect both memory management design and the guarantees that our semantics can
 1326 provide.

1327 Further research into type systems for λ_ω would also be interesting. Not only could type systems
 1328 guarantee absence of indexing errors but they could deliver the distinction between finite and
 1329 infinite cases as well. The type system we describe in the paper can serve as a starting point.
 1330 Decidability aspects around ordinals have raised interest independently. The first-order theory of
 1331 ordinal addition is known to be decidable [Büchi 1990], but more complex ordinal theories can
 1332 quickly get undecidable [Choffrut 2002].
 1333

1334 ACKNOWLEDGMENTS

1335 This material is based upon work supported by the National Science Foundation under Grant
 1336 No. nnnnnnn and Grant No. mmmmmmm. Any opinions, findings, and conclusions or recommend-
 1337 ations expressed in this material are those of the author and do not necessarily reflect the views of
 1338 the National Science Foundation.
 1339

1340 REFERENCES

- 1341 Michael Abbott, Thorsten Altenkirch, and Neil Ghani. 2005. Containers: Constructing strictly positive types. *Theoretical*
 1342 *Computer Science* 342, 1 (2005), 3 – 27. <https://doi.org/10.1016/j.tcs.2005.06.002> Applied Semantics: Selected Topics.
- 1343 Anonymous-1. 2018. Operational Semantics of Lambda-omega. Can be found in the anonymous supplementary materials.
 1344 (2018). [heh/semantics/lambda-omega-semantics.pdf](https://heh.semantics/lambda-omega-semantics.pdf)
- 1345 Anonymous-2. 2018. Recursive Array Comprehensions in a Call-by-Value Language. Can be found in the anonymous
 1346 supplementary materials. (2018). heh/array-comprehensions/rec-array-comprehensions.pdf
- 1347 Robert Atkey and Conor McBride. 2013. Productive Coprogramming with Guarded Recursion. In *Proceedings of the 18th*
 1348 *ACM SIGPLAN International Conference on Functional Programming (ICFP '13)*. ACM, New York, NY, USA, 197–208.
<https://doi.org/10.1145/2500365.2500597>
- 1349 Erik Barendsen and Sjaak Smetsers. 1996. Uniqueness Typing for Functional Languages with Graph Rewriting Semantics.
 1350 *Mathematical Structures in Computer Science* 6, 6 (1996), 579–612.
- 1351 Jarryd P. Beck, John Plaice, and William W. Wadge. 2015. Multidimensional infinite data in the language Lucid. *Mathematical*
 1352 *Structures in Computer Science* 25, 7 (2015), 1546–1568. <https://doi.org/10.1017/S0960129513000388>
- 1353 Robert Bernecky. 1987. An Introduction to Function Rank. *ACM SIGAPL Quote Quad* 18, 2 (Dec. 1987), 39–43.
- 1354 Robert Bernecky. 1993. The Role of APL and J in High-performance Computation. *ACM SIGAPL Quote Quad* 24, 1 (Aug.
 1355 1993), 17–32.
- 1356 Robert Bernecky and Paul Berry. 1993. *SHARP APL Reference Manual* (2nd ed.). Iverson Software Inc., Iverson Software Inc.,
 1357 33 Major St., Toronto, Canada.
- 1358 Robert Bernecky and Kenneth E. Iverson. 1980. Operators and Enclosed Arrays. In *APL Users Meeting 1980*. I.P. Sharp
 1359 Associates Limited, I.P. Sharp Associates Limited, Toronto, Canada, 319–331.
- 1360 R. S. Bird. 1987. An Introduction to the Theory of Lists. In *Proceedings of the NATO Advanced Study Institute on Logic of*
 1361 *Programming and Calculi of Discrete Design*. Springer-Verlag New York, Inc., New York, NY, USA, 5–42. <http://dl.acm.org/citation.cfm?id=42675.42676>
- 1362 Larry M. Breed, Roger D. Moore, Luther Woodrum, Richard Lathwell, and Adin Falkoff. 1972. (1972). <http://www.computerhistory.org/atcm/the-apl-programming-language-source-code> The APL360 source code is available at the
 1363 Computer History Museum.
- 1364 J. Richard Büchi. 1990. *Transfinite Automata Recursions and Weak Second Order Theory of Ordinals*. Springer New York, New
 1365 York, NY, 437–457. https://doi.org/10.1007/978-1-4613-8928-6_24
- 1366 D.C. Cann. 1989. *Compilation Techniques for High Performance Applicative Computation*. Technical Report CS-89-108.
 1367 Lawrence Livermore National Laboratory, LLNL, Livermore California.
- 1368 Christian Choffrut. 2002. *Elementary Theory of Ordinals with Addition and Left Translation by ω* . Springer Berlin Heidelberg,
 1369 Berlin, Heidelberg, 15–20. https://doi.org/10.1007/3-540-46011-X_2
- 1370 K. Ciesielski. 1997. *Set Theory for the Working Mathematician*. Cambridge University Press.
- 1371 Martin Elsmann and Martin Dybdal. 2014. Compiling a Subset of APL Into a Typed Intermediate Language. In *Proceedings of*
 1372 *ACM SIGPLAN International Workshop on Libraries, Languages, and Compilers for Array Programming (ARRAY'14)*. ACM,
 1373 New York, NY, USA, Article 101, 6 pages. <https://doi.org/10.1145/2627373.2627390>

- 1373 G. Estrin and R. Turn. 1963. Automatic Assignment of Computations in a Variable Structure Computer System. *IEEE*
 1374 *Transactions on Electronic Computers* EC-12, 6 (Dec 1963), 755–773. <https://doi.org/10.1109/PGEC.1963.263559>
- 1375 Peter Falster and Michael Jenkins. 1999. *Array Theory and Nial*.
- 1376 Jeremy Gibbons. 2017. APlicative Programming with Naperian Functors. In *European Symposium on Programming (LNCS)*,
 1377 Hongseok Yang (Ed.), Vol. 10201. 568–583. https://doi.org/10.1007/978-3-662-54434-1_21
- 1378 J. I. Glasgow and M. A. Jenkins. 1988. Array theory, logic and the Nial language. In *Proceedings. 1988 International Conference*
 1379 *on Computer Languages*. 296–303. <https://doi.org/10.1109/ICCL.1988.13077>
- 1380 Clemens Grelck and Sven-Bodo Scholz. 2006. SAC - A Functional Array Language for Efficient Multi-threaded Execution.
 1381 *International Journal of Parallel Programming* 34, 4 (2006), 383–427. <https://doi.org/10.1007/s10766-006-0018-x>
- 1382 Jeffrey P. Hammes, Bruce A. Draper, and A. P. Willem Böhm. 1999. *Sassy: A Language and Optimizing Compiler for*
 1383 *Image Processing on Reconfigurable Computing Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 83–97. https://doi.org/10.1007/3-540-49256-9_6
- 1384 Peter Hancock. 2000. *Ordinals and interactive programs*. Ph.D. Dissertation.
- 1385 S.K. Heller. 1989. *Efficient lazy data structures on a data-flow machine. Technical report*. Number TR-438.
- 1386 Ralf Hinze. 2010. Concrete Stream Calculus: An Extended Study. *J. Funct. Program.* 20, 5-6 (Nov. 2010), 463–535. <https://doi.org/10.1017/S0956796810000213>
- 1387 Roger K.W. Hui and Kenneth E. Iverson. 1998. *J Dictionary*.
- 1388 IBM. 1994. *APL2 Programming: Language Reference* (second ed.). IBM Corporation. SH20-9227.
- 1389 Kenneth E. Iverson. 1962. *A Programming Language*. John Wiley & Sons, Inc., New York, NY, USA.
- 1390 Jean-Baptiste Jeannin, Dexter Kozen, and Alexandra Silva. 2012. *CoCaml: Programming with Coinductive Types*. Technical
 1391 Report <http://hdl.handle.net/1813/30798>. Computing and Information Science, Cornell University.
- 1392 Michael A. Jenkins and Lenore R. Mullin. 1991. *A Comparison of Array Theory and a Mathematics of Arrays*. Springer US,
 1393 Boston, MA, 237–267. https://doi.org/10.1007/978-1-4615-4002-1_14
- 1394 Inc. Jsoftware. 2016. Jsoftware: High performance development platform. <http://www.jsoftware.com/>. (2016).
- 1395 Gilles Kahn. 1974. The Semantics of Simple Language for Parallel Programming.. In *IFIP Congress*. 471–475.
- 1396 G. Kahn. 1987. Natural semantics. In *STACS 87*, FranzJ. Brandenburg, Guy Vidal-Naquet, and Martin Wirsing (Eds.). Lecture
 1397 Notes in Computer Science, Vol. 247. Springer Berlin Heidelberg, 22–39. <https://doi.org/10.1007/BFb0039592>
- 1398 Dexter Kozen and Alexandra Silva. 2016. Practical coinduction. *Mathematical Structures in Computer Science* FirstView
 1399 (February 2016), 1–21. <https://doi.org/10.1017/S0960129515000493>
- 1400 Xavier Leroy and Hervé Grall. 2009. Coinductive big-step operational semantics. *Information and Computation* 207, 2 (Feb.
 1401 2009), 284–304. <https://doi.org/10.1016/j.ic.2007.12.004>
- 1402 Ben Lippmeier, Fil Mackay, and Amos Robinson. 2016. Polarized Data Parallel Data Flow. In *Proceedings of the 5th*
 1403 *International Workshop on Functional High-Performance Computing (FHPC 2016)*. ACM, New York, NY, USA, 52–57.
 1404 <https://doi.org/10.1145/2975991.2975999>
- 1405 Panagiotis Manolios and Daron Vroon. 2005. Ordinal Arithmetic: Algorithms and Mechanization. *Journal of Automated*
 1406 *Reasoning* 34, 4 (2005), 387–423. <https://doi.org/10.1007/s10817-005-9023-9>
- 1407 John McCarthy. 1960. Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part I. *Commun.*
 1408 *ACM* 3, 4 (April 1960), 184–195. <https://doi.org/10.1145/367177.367199>
- 1409 Eugene E. McDonnell and Jeffrey O. Shallit. 1980. Extending APL to Infinity. In *APL 80 : International Conference on APL*,
 1410 Gijsbert van der Linden (Ed.). Amsterdam ; New York : North-Holland Pub. Co. : sole distributors for the USA and Canada,
 1411 Elsevier North-Holland, 123–132.
- 1412 Rasmus Ejlers Møgelberg. 2014. A Type Theory for Productive Coprogramming via Guarded Recursion. In *Proceedings of*
 1413 *the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth*
 1414 *Annual ACM/IEEE Symposium on Logic in Computer Science (LICS) (CSL-LICS '14)*. ACM, New York, NY, USA, Article 71,
 1415 10 pages. <https://doi.org/10.1145/2603088.2603132>
- 1416 Trenchard More. 1973. Axioms and Theorems for a Theory of Arrays. *IBM J. Res. Dev.* 17, 2 (March 1973), 135–175.
 1417 <https://doi.org/10.1147/rd.172.0135>
- 1418 Lenore Mullin and Scott Thibault. 1994. *A Reduction Semantics for Array Expressions: The PSI Compiler*. Technical Report
 1419 CSC-94-05. University of Missouri-Rolla.
- 1420 Lenore M. Restifo Mullin. 1988. *A Mathematics of Arrays*. Ph.D. Dissertation. Syracuse University.
- 1421 Carl Adam Petri. 1962. *Kommunikation mit Automaten*. Ph.D. Dissertation. Universität Hamburg.
- Justin Slepak, Olin Shivers, and Panagiotis Manolios. 2014. An Array-Oriented Language with Static Rank Polymorphism.
 In *Proceedings of the 23rd European Symposium on Programming Languages and Systems - Volume 8410*. Springer-Verlag
 New York, Inc., New York, NY, USA, 27–46. https://doi.org/10.1007/978-3-642-54833-8_3
- Robert Stephens. 1997. A survey of stream processing. *Acta Informatica* 34, 7 (01 Jul 1997), 491–541. <https://doi.org/10.1007/s002360050095>

- 1422 Bo Joel Svensson and Josef Svenningsson. 2014. Defunctionalizing Push Arrays. In *Proceedings of the 3rd ACM SIGPLAN*
1423 *Workshop on Functional High-performance Computing (FHPC '14)*. ACM, New York, NY, USA, 43–52. <https://doi.org/10.1145/2636228.2636231>
- 1424 R. W.W. Taylor. 1982. Indexing Infinite Arrays: Non-finite Mathematics in APL. *SIGAPL APL Quote Quad* 13, 1 (July 1982),
1425 351–355. <https://doi.org/10.1145/390006.802264>
- 1426 William Thies, Michal Karczmarek, and Saman P. Amarasinghe. 2002. StreamIt: A Language for Streaming Applications. In
1427 *Proceedings of the 11th International Conference on Compiler Construction (CC '02)*. Springer-Verlag, London, UK, UK,
1428 179–196. <http://dl.acm.org/citation.cfm?id=647478.727935>
- 1429 Kai Trojahner and Clemens Grelck. 2009. Dependently typed array programs don't go wrong. *The Journal of Logic and*
1430 *Algebraic Programming* 78, 7 (2009), 643 – 664. <https://doi.org/10.1016/j.jlap.2009.03.002> The 19th Nordic Workshop on
1431 Programming Theory (NWPT 2007).
- 1432 D. A. Turner. 1995. *Elementary strong functional programming*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1–13.
1433 https://doi.org/10.1007/3-540-60675-0_35
- 1434 A. Šinkarovs, S.B. Scholz, R. Bernecky, R. Douma, and C. Grelck. 2013. SAC/C Formulations of the All-Pairs N-Body
1435 Problem and their Performance on SMPs and GPGPUs. *Concurrency and Computation: Practice and Experience* (2013).
1436 <https://doi.org/10.1002/cpe.3078>
- 1437 Philip Wadler. 1995. *Monads for functional programming*. Springer Berlin Heidelberg, Berlin, Heidelberg, 24–52. https://doi.org/10.1007/3-540-59451-5_2
- 1438 V. Wieser, C. Grelck, P. Haslinger, J. Guo, F. Korzeniowski, R. Bernecky, B. Moser, and S.B. Scholz. 2012. Combining High
1439 Productivity and High Performance in Image Processing Using Single Assignment C on Multi-core CPUs and Many-core
1440 GPUs. *Journal of Electronic Imaging* 21, 2 (2012). <https://doi.org/10.1117/1.JEI.21.2.021116>
- 1441 Hongwei Xi and Frank Pfenning. 1998. Eliminating Array Bound Checking Through Dependent Types. In *Proceedings of the*
1442 *ACM SIGPLAN 1998 Conference on Programming Language Design and Implementation (PLDI '98)*. ACM, New York, NY,
1443 USA, 249–257. <https://doi.org/10.1145/277650.277732>
- 1444
- 1445
- 1446
- 1447
- 1448
- 1449
- 1450
- 1451
- 1452
- 1453
- 1454
- 1455
- 1456
- 1457
- 1458
- 1459
- 1460
- 1461
- 1462
- 1463
- 1464
- 1465
- 1466
- 1467
- 1468
- 1469
- 1470