

# Extending a C-like Language for Portable SIMD Programming

Roland Leißa    Sebastian Hack

Compiler Design Lab, Saarland University  
{leissa, hack}@cs.uni-saarland.de

Ingo Wald

Visual Applications Research, Intel Corporation  
ingo.wald@intel.com

## Abstract

SIMD instructions are common in CPUs for years now. Using these instructions effectively requires not only vectorization of *code*, but also modifications to the *data layout*. However, automatic vectorization techniques are often not powerful enough and suffer from restricted scope of applicability; hence, programmers often vectorize their programs manually by using intrinsics: compiler-known functions that directly expand to machine instructions. They significantly decrease programmer productivity by enforcing a very error-prone and hard-to-read assembly-like programming style. Furthermore, intrinsics are not portable because they are tied to a specific instruction set.

In this paper, we show how a C-like language can be extended to allow for portable and efficient SIMD programming. Our extension puts the programmer in total control over where and how control-flow vectorization is triggered. We present a type system and a formal semantics of our extension and prove the soundness of the type system. Using our prototype implementation IVL that targets Intel’s MIC architecture and SSE instruction set, we show that the generated code is roughly on par with handwritten intrinsic code.

**Categories and Subject Descriptors** D.1.3 [Concurrent Programming]: Parallel programming; D.3.1 [Formal Definitions and Theory]: Semantics, Syntax

**General Terms** Languages, Performance, Theory

**Keywords** language theory, parallel programming, polymorphism, semantics, SIMD, SIMT, type system, vectorization

## 1. Introduction

SIMD instructions are available on commodity processors for more than a decade now. Many developers from various domains (for example, high-performance graphics [7], databases [36], or bioinformatics [6]) use SIMD instructions to speed up their applications. Many of those algorithms are not massively data-parallel but contain data-parallel and scalar parts intermixed at a fine granularity. To illustrate this, let us run through the fundamental techniques which are necessary in order to port traditional code to SIMD architectures by considering the small example in Figure 1. It is taken (and slightly modified) from a paper by Zhou and Ross about the use of SIMD instructions to speed up data base operations [36]. The program scans data records for a certain criterion and returns the value of the first record for which the criterion holds. In this par-

```
1 struct data_t {
2     int key;
3     int other;
4 };
5
6 int search(data_t* data, int N) {
7     for (int i = 0; i < N; i++) {
8         int x = data[i].key;
9         if (4 < x & x <= 8) return x;
10    }
11    return -1;
12 }
```

Figure 1. Example program (by Zhou and Ross [36])

```
1 struct simd_data_t {
2     simd_int key;
3     simd_int other;
4 };
5
6 int search(simd_data_t* data, int N) {
7     for (int i = 0; i < N/L; ++i) {
8         simd_int x = load(data[i].key);
9         simd_int cmp = simd_and(simd_lt(4, x),
10                                simd_le(x, 8));
11         int mask = simd_to_mask(cmp);
12         if (mask != 0) {
13             simd_int result = simd_and(mask, x);
14             for (int j = 0; j < log2(L); j++)
15                 result = simd_or(result,
16                                 whole_reg_shr(result, 1 << j));
17             return simd_extract(result, 0);
18         }
19     }
20     return -1;
21 }
```

Figure 2. Manually vectorized example program

ticular example, we assume that at most one element in the array fulfills the criterion.<sup>1</sup>

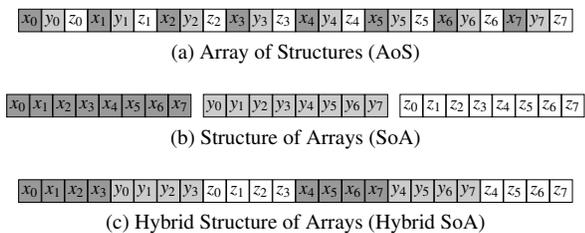
The hand-vectorized version of the program is shown in Figure 2. We assume that  $L$  is the native vector length of the target machine, i.e., there are  $L$  slots in one SIMD register and each slot can hold exactly one `int`. When performing vector computations we refer to the computations applied to a certain slot index as *SIMD* or *vector lane*. The program takes  $L$  data items from the array at once and processes them in parallel. It features the three most common techniques used to adapt a program for SIMD instructions:

**Adapting the data layout.** Fetching  $L$  values in parallel from memory is tremendously faster, if they lie consecutively in memory. Thus, the traditional *array of structures* layout (AoS, see Figure 3a) is not suitable for SIMD programming, as it requires non-contiguous memory access—called *scatter* and *gather* (S/G)—that severely degrades performance. A *structure of arrays* (SoA, see Figure 3b) provides a better solution. However, a loop over an SoA must maintain one pointer for each

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PPoPP’12, February 25–29, 2012, New Orleans, Louisiana, USA.  
Copyright © 2012 ACM 978-1-4503-1160-1/12/02...\$10.00

<sup>1</sup> Zhou and Ross mention searching in  $B^+$  trees with unique keys as an application.



**Figure 3.** Different array layouts with vector length 4 for a struct with three members:  $x$ ,  $y$  and  $z$

member. For this reason and due to data locality problems, the *hybrid SoA* (see Figure 3c) is superior, as this layout does not suffer from these problems. It emerges by first inflating the original struct to the vector length (we call this unit a *block* and the process of doing so *type vectorization*) and then grouping this block in an array. This data layout is regarded best practice in SIMD programming and is recommended by various application notes of all major CPU manufacturers, e.g. [13]. These blocks can be used to build other vectorized data structures like linked lists or trees just as well.

**Control-flow to data-flow conversion.** In line 9, the comparison is evaluated on all  $L$  loaded items in parallel. The result is a vector of booleans being *true* for those vector lanes for which the comparison was *true*. That mask has then to be converted into an ordinary `int` in order to use it in the conditional branch in line 12. In line 13, the mask is used to discard all values for which the comparison was *false*.

**Seamless integration of scalar and vectorized code.** While several statements of our example program perform data-parallel computations, the inner loop (lines 14–16) is *not* vectorized. It performs a scan-like operation: by successively shifting and OR-ing the *whole* SIMD variable, the first non-zero element of the vector is propagated to element 0 of the vector. It is typical for manually vectorized code to have a strong interaction between scalar and vector code.

This example illustrates the current state of the art in hand-tuning programs for SIMD instruction sets and clearly shows its limitations. First, vector code has to be written using assembly-like intrinsics: compiler-known functions that directly correspond to an assembly instruction. While operator overloading can improve the readability of expressions, control structures can usually not be overloaded. Hence, the control to data-flow conversion has to be performed manually by explicit handling of masks. This severely degrades the readability of the code because control statements are replaced by clumsy mask handling code. Furthermore, intrinsics directly expose a concrete instruction set architecture (ISA) to the programmer. Accordingly, intrinsics code is not only inherently unportable between different processor architectures but also within different revisions of the ISA *within* the same processor family.

Second, the definition of data structures is no longer compositional: To vectorize a struct, the programmer has to create a new type where the members of the struct are vectorized (see Figure 2). Often she is interested in the scalar *and* the vectorized type because she often refers to a single component of a vector. Hence, all types would have to be written twice. However, it is reasonable for the programmer to assume that she can obtain the vectorized version of a data type  $T$  by applying a type constructor, say `T block[N]`, which yields the desired block type. Furthermore, the compiler *knows* that `T block[N]` is the  $N$  times vectorized version of  $T$ . This is important for parametric polymorphism (see below) and correct handling of S/G in the type system.

Third, the hand-vectorized function of Figure 2 works *only* in a vector context. Many functions however should work in a scalar

*and* in a vector context. Hence, functions should be polymorphic with respect to vectorization where possible.

We believe it to be clear that the code transformations presented above cannot be carried out automatically. Current optimizing compilers vectorize loop nests which have to satisfy strong prerequisites. Automatic interprocedural data-structure rewriting or automatic scalar/vectorial code switching are out of the scope of what an optimizing compiler for a traditional C-like language can do. Furthermore, porting an algorithm to SIMD instructions does not only change the mapping of a *given* implementation to another hardware but often requires adapting the implementation of the algorithm itself. For example, note lines 13–17 in Figure 2 differ from the corresponding code in Figure 1.

On GPUs, which internally also make extensive use of SIMD, this issue is addressed by not using a traditional scalar programming paradigm like in C, but rather using the SIMT paradigm (single instruction, multiple threads) as employed in CUDA [29] or OpenCL [17]: In the SIMT paradigm each program—called *kernel*—is always instantiated  $n$  times to  $n$  logical *threads*. These threads all execute this same program at the same time but on different data; SIMD is then employed by always processing  $k$  such threads together where  $k$  is the SIMD width. A group of  $k$  threads is also called *warp*. Each of a warp’s  $k$  threads run one SIMD lane. In this programming paradigm, the user writes scalar-looking, C-like code, and the control-flow to data-flow conversion along with the necessary predication is mostly handled by the hardware.

While SIMT is a very powerful paradigm that allows to express parallelism in a scalar-looking way, it also has its disadvantages. In particular, in a SIMT language *everything* is always vectorized, which closely matches the execution model of such GPUs, but does not easily allow to express scalar parts of a program. Executing truly scalar code then requires explicitly masking off all but one thread of a block and costly explicit synchronization between the warps in a block, since the hardware usually needs to run multiple such warps in parallel. Furthermore, by explicitly *hiding* the effects of vectorization, these languages have no means of allowing the programmer to define vectorized data layouts, thus leaving her with the choice between inefficient AoS data layouts (and costly AoS to (hybrid) SoA conversions on access), or manual vectorization of data structures. Above all, writing all the glue code in order to merge the host code with the kernels is very inconvenient.

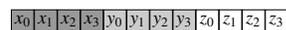
## 1.1 Contributions

In this paper, we propose a different programming model for data-parallel programming that combines the advantages of both scalar and SIMT programming, and is especially suited for machines with explicit SIMD instructions. We do this by defining a small core language called `VECIMP` (see Sections 2 and 4), which extends the type system and formal semantics of an existing formalization of a substantial subset of C, called `IMP` (see Section 3).

- In contrast to the SIMT model, we do not require the strict separation in host code and kernels but integrate vectorization seamlessly into the language. Moreover, we allow to mix scalar and SIMT programming (and scalar and vectorized data structures) in the same basic blocks.
- We provide a type qualifier `block[N]` to obtain a block of vector length  $N$ . The example code

```
struct vec3 { float x, y, z; };
vec3 block[4] v;
```

produces the following memory layout for  $v$ :



- In contrast to the vectorize-everything paradigm of SIMT, a `VECIMP` program starts in a *scalar* context. We use the type

system of VecIMP to trigger vectorization: An if-statement with a vector-type condition is vectorized.

```
int block[4] v, w;
if (v < w) /* vectorized */ else /* vectorized */
```

- We implemented several features of VecIMP in our research language IVL and evaluated the performance of the vectorized code experimentally on several benchmarks. Our target machine was a 16-wide Intel® Many Integrated Core (MIC) processor [14, 31]. We are within 4% to 15% of the performance of hand-written intrinsic code. Writing the IVL code was almost as easy as writing traditional scalar code. Besides that, IVL code is fully portable because it does not contain any architecture-dependent details (see Section 5). The programs can just be recompiled for the SSE back-end.

## 2. VecIMP at a Glance

As outlined in the introduction, a programmer often wants to perform SIMD operations only on certain lanes and not on the whole vector. We call these lanes *active lanes*. In hardware, this is either accomplished by *predicated execution* [31] or *blending* (SSE, AVX). In the first case, each vector instruction is equipped with a mask, which indicates active lanes, and a vector instruction performs the operation only for active lanes. In the latter case, vector instructions always perform operations for all lanes. Unwanted effects must be later on masked out by so-called blend operations. The programming model of VecIMP exposes predicated execution to the programmer, although blending can be used by the compiler in order to emulate this execution model. For this reason, every statement in VecIMP is executed in a *scalar* or *vector context*. For typing, the context corresponds the vector length of the current predication mask—also called *current mask*. The programmer has read access to this variable by the `current_mask` keyword. A context of length one is called *scalar context*. In VecIMP, evaluation begins in a scalar context and control-flow statements like `if` or `while` spawn different vector contexts:

```
// vector length of context = 1; current_mask = T
int block[4] v = <0,3,4,1>;
int block[4] w = 3; // <3,3,3,3> via broadcast
bool block[4] m = v < w; // <T,F,F,T>
++v; // <1,4,5,2>
if (m) {
    // vector length of context = 4; current_mask = m
    v += 2; // <3,4,5,4>
} else {
    // vector length of context = 4; current_mask = ~m
    v += 3; // <3,7,8,4>
}
// vector length of context = 1; current_mask = T
```

As can be seen in this example, a type modifier `block[N]` is available. This qualifier recursively inflates the data type to be of vector length `N`. If a qualifier is elided, the variable is *unbound*. This means that the vector length is inferred by the vector length of the right-hand side. If a right-hand side does not exist, the vector length of the current context is used:

```
struct vec3 { float x, y, z; };
// scalar context
vec3 block[4] v;
vec3 w = v; // vec3 block[4];
vec3 x; // vec3 block[1];
```

Let us now implement a dot product:

```
float dot(vec3 a, vec3 b) {
    return a.x*b.x + a.y*b.y + a.z*b.z;
}
```

All *unbound* parameters of functions in VecIMP are polymorphic in their vector length and so are `a` and `b`. The vector length of an *unbound* return type is automatically inferred. Thus, a call-site in a scalar context

```
1 struct data_t {
2     int key;
3     int other;
4 };
5
6 int search(data_t *scalar data, int scalar N) {
7     int L = lengthof(*data);
8     for (int i = 0; i < N/L; ++i) {
9         int x = data[i].key;
10        if (4 < x & x <= 8)
11            int block[L] result = [x, 0];
12            scalar {
13                for (int j = 0; j < log2(L); ++j)
14                    result |= whole_reg_shr(result, 1 << j);
15                return get(x, 0);
16            }
17        }
18    return -1;
19 }
```

Figure 4. Example program in VecIMP

```
vec3 block[4] a = /*...*/;
vec3 scalar b = /*...*/;
float block[4] r = dot(a, b);
```

instantiates a version of `dot` with this signature:

```
float block[4] dot(vec3 block[4] a, vec3 scalar b);
```

If a scalar and a vectorial value are mixed in an operation, the scalar value is broadcast. Thus, the uses of `b.x`, `b.y` and `b.z` get implicitly broadcast in this version. But what happens, if we call the same function from a vectorial context?

```
if (a.x < b.x) { // enter context of length 4
    float block[4] r = dot(a, b); /*...*/
}
```

This time `dot` must know of the current mask. For this reason, we pass the current mask as hidden parameter to `dot`. In VecIMP, functions are also polymorphic in the context they are called in. Thus, the signature of this `dot` instance is:

```
float block[4] dot(bool block[4] current_mask,
    vec3 block[4] a, vec3 scalar b);
```

As a more sophisticated example, let us reconsider the introductory example in VecIMP (Figure 4). The function `search` can be invoked with a pointer to an ordinary array of `data_t` or a hybrid SoA of `data_t`. However, the pointer itself is scalar, denoted by the `scalar` modifier which is just shorthand for `block[1]`. Hence, one cannot pass a vector of pointers to `search`. Similarly, `N` stays scalar in all possible instances of `search`. Consider the call-site in a scalar context:

```
data_t block[4] *data; /*...*/
int x = search(data, N);
```

The type of `data` is an unbound pointer to a hybrid SoA of length 4 of `data_t`. Since `data` is declared in a scalar context, `data` itself is instantiated as scalar while the referenced type remains vectorial. Since `search` is invoked from a scalar context, the context of this particular `search` instance created for this call-site is scalar, too. As `L` and `i` in line 7 and 8, respectively, inherit the vector length from the right-hand side, these variables are scalar. Now consider line 9: Since the type of the right-hand side is `int block[4]` and `x` is unbound, `x` is also of vector type `int block[4]`. As `x` is now of vector type, the `if` statement in line 10 is vectorized. Thus, reading from `x` just obtains the (by precondition unique) found key which is—via the assignment in line 11—written into the vector `result`. However, the inactive lanes of `result` would stay undefined and we therefore set them explicitly to 0. This is important because subsequently, the vector `result` is processed as a whole and not in a vectorized mode: The algorithm propagates the found key to position 0. For this reason, we want to enter a scalar context again. The `scalar`

keyword in front of a statement stores the current mask on a stack and resets it to the scalar mask. Hence, the for loop beginning in line 13 is executed in a scalar context. After the for loop finishes execution, the scalar element at position 0 of the vector is returned.

In the case that `search` is invoked on scalar data, the whole procedure degenerates to a scalar program whose semantics is equivalent to the program in Figure 1. Although we cannot gain an optimized vectorial program from a scalar one, we can still extract the scalar version from a program written in a polymorphic way.

### 3. Formalization

In this section, we first introduce the notation used throughout this paper and then the imperative language IMP.

The set which holds all powers of two is referred to as  $\mathbb{P} := \{2^i \mid i \in \mathbb{N}\}$ . We use  $\bar{a}$  as shorthand for a sequence  $a_0, \dots, a_{|\bar{a}|-1}$  with zero or more elements, while  $|\bar{a}|$  denotes the number of items in a sequence. The term  $\bar{a}^+$  represents a sequence with one or more elements. Sequences are propagated into more complex expressions while global values and constants remain untouched. Thus,  $(\bar{a}, 3)$  expands to  $(a_0, 3, \dots, a_{|\bar{a}|-1}, 3)$ .

#### 3.1 Types

We use calligraphic letters to indicate sets of types and use small Greek letters for types. The size of a type  $\tau$  given in bytes is referred to as  $|\tau|$ . Furthermore, we make use of the following conventions:

We use standard C-like primitive integer and floating point types whose sizes (in bytes) must be a power of two. Booleans are denoted by `bool` := {T, F}. The size of a `bool` is implementation-dependent and not transparent to the programmer. We pool these types into a set of *atomic types*  $\mathcal{A}$ .

A pointer  $\tau^*$  points to a location which is interpreted to be of type  $\tau$ ; the size of a pointer is also implementation-dependent.

A record  $\rho := \{\bar{\tau}^+\}$  is compounded of other element types  $\bar{\tau}^+$ . A record may not (in)directly contain an element of its own type, but cycles induced by pointers are allowed (as in C). A projection  $\rho.i$ , where  $i \in \mathbb{N} \wedge 0 \leq i < |\bar{\tau}^+|$ , grabs  $\tau_i$ . An implementation is free to add additional padding space between the elements of a record. This implies that the size of a record is machine dependent, too. The set of all types which can be built with atomic types, records and pointers is denoted by  $\mathcal{T}$ .

A *vector type*  $v$  with element type  $\tau \in \mathcal{A} \cup \{\pi^* \mid \pi \in \mathcal{V}\}$  and  $l \in \mathbb{P}$  elements is written as  $(\tau \mid l)$ . An instance of such a type  $v$  is written as  $\langle \bar{v} \rangle$ . The set of all types which can be built with scalar types and vector types is denoted by  $\mathcal{V}$ .

#### 3.2 IMP

Before describing the vector extensions, we briefly review the imperative language IMP we extend. IMP is a substantial subset of C. Its formalization is closely related to Norrish [26].

The syntax is given in Figure 5. To simplify matters, we assume that each function has exactly one return statement and function calls are used instead of typical in- and prefix operators (although we silently make use of operators in later `VecIMP` examples, for the sake of readability; moreover, we use named structs instead of tuples and other syntactic sugar for the same reason). A representative set of type inference and evaluation rules is given in Figure 7. See Figure 13 for the complete version.

**Typing rules.** Static semantics maintains a type map  $\Gamma$ , which maps identifiers to types, and has read-access to a function map  $\Phi$ , which maps identifiers to a function’s signature and body. The type map is altered in a declaration (see T-DECL). On the other hand the occurrence of a variable is resolved by a lookup of the given identifier in this map (see T-VAR). Type inference rules must differentiate between statements and expressions. Statements yield the

$e ::=$	$c_\alpha$	Expressions:
	$\text{id}$	constant of type $\alpha \in \mathcal{A}$
	$e.i$	variable
	$\&e$	projection
	$*e$	address
	$e_1[e_2]$	dereferencing
	$\text{id}(\bar{e})$	indexing
		invoke
$s ::=$		Statements:
	$;$	skip
	$\{ \hat{s} \}$	scope
	$e;$	expression-statement
	$t \text{ id};$	declaration
	$e_2 = e_1;$	assignment
	$t \text{ id} = e;$	initialization
	$\text{if } (e) \ s_1 \ \text{else } \ s_2$	if-else
	$\text{while } (e) \ s$	while
$\hat{s} ::=$		Scoped statements:
	$s$	statement
	$\hat{s} \ s$	sequence
$t ::=$		Type:
	$\alpha \in \mathcal{A}$	base
	$\tau^*$	pointer
	$\{ \bar{\tau}^+ \}$	record
$f ::=$		Function:
	$\epsilon$	empty
	$f_1 \ f_2$	sequence
	$t \ \text{id}(t \ \text{id}) \{ s \ \text{return } e; \}$	function

Figure 5. Syntax of IMP

dummy type  $\epsilon$ . Since expressions can return lvalues, type inference rules acknowledge this fact by an internal lvalue type  ${}^l\tau$  (read *lvalue of type*  $\tau$ ). If we want to state that a type must not be an lvalue, we write  ${}^v\tau$ . In many typing rules both lvalue and value types are allowed in a given premise. We annotate such a type as  ${}^l\tau$  as shorthand (see T-ASSIGN).

**Evaluation rules.** Dynamic semantics uses small step semantics evaluating a program  $a$  in scope  $\sigma$  and state  $M$  in one step to a program  $a'$  in scope  $\sigma'$  in a new state  $M'$  with the help of the function map  $\Phi$ . The state  $M$  is a memory map, mapping an address to an `is8`. We assume  $M$  to be infinite and initialized with random values. A scope  $\sigma$  consists of an address map  $A$  which maps an identifier of a variable to the starting address in memory and a type map  $T$  which maps an identifier to a type. In order to extract a subsequence of  $n$  bytes starting at address/offset  $a$  of a memory map  $M$  or a byte sequence  $m$ , we use  $M\langle a, n \rangle$  or  $m\langle a, n \rangle$  as shorthand, respectively.

As already outlined, expressions can evaluate to values and lvalues. A value consists of a sequence of bytes  $m$  and a type  $\tau$  and is written as  $\mathcal{V}(m, \tau)$ . An lvalue  $\mathcal{L}(a, \tau)$  consists of an address  $a$  and a type  $\tau$  and represents a reference to a memory location. Both constructs are added to the syntax (but are not directly usable by the programmer and hence do not appear in Figure 5). For instance, `E-VAR` looks up the identifier of a variable in the address and type map to build an lvalue. Statements can evaluate to the final configuration “;” (read *skip*).

A declaration (see E-DECL) allocates memory and updates the type and address maps accordingly. The function `alloc`( $\sigma, \tau$ ) is used for this task. It returns an address pointing to  $|\tau|$  bytes of usable memory. An assignment (see E-ASSIGN) updates the memory map at the position pointed to by the given left-hand side lvalue with the given right-hand side value.

In order to memorize an old scope  $\sigma$  when entering a new one we use the special construct “{  $s \}$ ” which is also added to the “internal” syntax (see E-INSCOPE and E-OUTSCOPE).

$e ::= \dots$	<b>current_mask</b>	<i>Expressions:</i> current mask
$s ::= \dots$	<b>scalar</b> $s$	<i>Statements:</i> scalar
$t ::=$		<i>Type:</i>
	$\alpha \in \mathcal{A} \ q$	base
	$t^* \ q$	pointer
	$\{\bar{t}^+\} \ q$	record
$q ::=$		<i>Qualifiers:</i>
	$\epsilon$	unbound
	<b>block</b>	auto block
	<b>block</b> [ $i$ ]	block

**Figure 6.** Syntax changes of IMP in order to obtain VecIMP: The current-mask expression and the scalar statement are added. Types follow a different grammar along with the new qualifier rule.

**Theorem 1** (Soundness of IMP). *Each well-typed IMP program is either in its final configuration or*

1. *there exists an evaluation rule which progresses evaluation and*
2. *the resulting program after application of an evaluation rule preserves well-typedness.*

See Appendix A for details and the complete proof of this theorem.

## 4. Extending IMP

In the following sections, we show how to extend IMP to its vector counterpart VecIMP. The needed modifications of IMP’s syntax in order to obtain VecIMP are given in Figure 6. During discussion it is important to understand that VecIMP is just a substantial intentionally small calculus. Practical languages may add as much syntactic sugar as needed. Furthermore, languages with powerful abstraction mechanisms can build higher-level vector programming constructs (see Section 7).

An excerpt of the typing and evaluation rules, which are referred to in this section, is shown in Figure 8. We have already seen in the introductory example that we must keep track of the current mask while doing SIMD computations. This has several implications. First, all typing rules need a new variable  $\Lambda \in \mathbb{P}$  called *current vector length* or *context*. Likewise, all evaluation rules need a new scope variable  $P_\sigma \in \langle \text{bool} \mid l \rangle$  ( $l \in \mathbb{P}$ ), which reflects predication. We extend our notation of extracting a byte sequence to a predicated extraction  $M\langle P, a, n \rangle$  or  $m\langle P, a, n \rangle$ , respectively, where  $P$  is a mask. This *current mask* or *predication variable* can be queried via the **current\_mask** keyword, which is added as expression to the syntax. Generally, a statement is only executed if the current mask is not entirely false. In the special case that  $\Lambda = 1$  and accordingly  $P_\sigma = \langle T \rangle$ , we are in a scalar context.

### 4.1 Types

The context influences the generation of types while the programmer retains tight control over this process with the help of few qualifiers (see T-, E-DeclL). Type vectorization is a recursive process:

- A **block** type adopts the vector length of the current context, i.e., in a scalar context the type stays scalar.
- A **block**[ $i$ ] type always receives vector length  $i$  where  $i$  is a compile time constant. The special case  $i = 1$  means a scalar type and we write **scalar** as syntactic sugar for **block**[1]. In the case that  $i > 1$  does not match the current vector length, a type error occurs.
- Types without explicit qualifier are qualified as *unbound*. These types adopt the vector length of the current recursion. Top level unbound types either infer the vector length from the right-hand

side (in an initialization, see T- and E-Init) or adopt the current vector length (in a declaration statement, see T- and E-Decl).

- Records are never vectorized. Instead the qualifier of a record controls recursively vectorization for *unbound* element types.

The type constructor of a qualified type  $t$  in context  $\Lambda$  is defined as follows:<sup>2</sup>

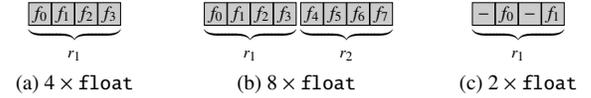
$$\text{vec}_\Lambda : t \times \mathbb{P} \rightarrow \mathcal{V} \cup \{\perp\}$$

$$(v, n) \mapsto \begin{cases} \perp & \text{if } v = \tau \in \mathcal{T} \text{ block}[i] \wedge \Lambda \neq i > 1, \\ \langle \alpha \mid n \rangle & \text{if } v = \alpha \in \mathcal{A}, \\ \langle \alpha \mid \Lambda \rangle & \text{if } v = \alpha \in \mathcal{A} \text{ block}, \\ \langle \alpha \mid i \rangle & \text{if } v = \alpha \in \mathcal{A} \text{ block}[i], \\ \langle \text{vec}_\Lambda(\tau, n)^* \mid n \rangle & \text{if } v = \tau^*, \\ \langle \text{vec}_\Lambda(\tau, n)^* \mid \Lambda \rangle & \text{if } v = \tau^* \text{ block}, \\ \langle \text{vec}_\Lambda(\tau, n)^* \mid i \rangle & \text{if } v = \tau^* \text{ block}[i], \\ \overline{\langle \text{vec}_\Lambda(\tau, n) \rangle} & \text{if } v = \{\bar{\tau}^+\}, \\ \overline{\langle \text{vec}_\Lambda(\tau, \Lambda) \rangle} & \text{if } v = \{\bar{\tau}^+\} \text{ block}, \\ \overline{\langle \text{vec}_\Lambda(\tau, i) \rangle} & \text{if } v = \{\bar{\tau}^+\} \text{ block}[i]. \end{cases}$$

Consider the following example:

```
struct Bar { int c; int block d; };
struct Foo { int a; int block b; Bar scalar bar; };
// current vector length = 4
Foo foo;
int i = 0;
```

Upon the declaration of `foo`  $\text{vec}_4(\text{Foo}, 4)$  is applied in order to vectorize the type. Member `a` gets vectorized as it sees 4 as parameter; `b` gets vectorized, because it sees 4 as current vector length; however, `c` stays scalar, because it sees length 1 passed to  $\text{vec}_4$  via `Bar scalar bar`; finally, `d` gets vectorized for the same reason as `b`. The variable `i` stays scalar as the right-hand side is also scalar.



**Figure 9.** Different packagings of floats into hardware SIMD registers of 16 byte width

In order to exploit SIMD hardware efficiently, it is very important to organize data in a way that maximum coherence is achieved, i.e., optimally all SIMD lanes are occupied and none lie idle. If the desired vector length is equal to the one which is available in hardware, all elements can be put seamlessly into one register (Figure 9a). If the desired vector length exceeds the native one, several registers must be allocated (Figure 9b). In the case that the requested length is smaller than the native one, padding space must be inserted. One possible solution is depicted in Figure 9c, albeit other possibilities may be more advantageous for the underlying hardware. As case b) usually increases register pressure and case c) decreases throughput, it is important to choose the vector length wisely. A compiler can provide a built-in mechanism for querying the preferred vector length of a type for the given target machine:

```
struct T { double d; float f; };
T block[preferred_lengthof(T)] t;
```

See Section 7 for additional suggestions in order to abstract away other low-level details.

The current vector length also restricts types that may occur. Access of type  $\tau$  in context  $\Lambda$  is only granted in a scalar context, or if the type is scalar, or if the current vector length and the one of

<sup>2</sup> Note that cyclic dependencies in types induced by pointers can be resolved by first defining an opaque type and not to follow that pointer and later on refining that field.

	$\Phi, \Gamma \vdash e : \tau, \Gamma$		$\Phi, a, \sigma, M \rightarrow a', \sigma', M'$
$\frac{\text{id} \in \text{dom}(\Gamma)}{\Phi, \Gamma \vdash \text{id} : \tau(\text{id}), \Gamma}$	T-VAR	$\frac{\text{id} \in \text{dom}(T_\sigma)}{\Phi, \text{id}, \sigma, M \rightarrow \mathcal{L}(A_\sigma(\text{id}), T_\sigma(\text{id})), \sigma, M}$	E-VAR
$\frac{\Phi, \Gamma \vdash s : \epsilon, \Gamma'}{\Phi, \Gamma \vdash \{s\} : \epsilon, \Gamma}$	T-SCOPE	$\frac{\Phi, \{s\}, \sigma, M \rightarrow \{s\}^{\sigma'}, \sigma', M}{\Phi, \{s\}^{\sigma'}, \sigma', M \rightarrow ;, \sigma', M}$	E-INSCOPE
$\frac{\tau = t \quad \text{id} \notin \text{dom}(\Gamma)}{\Phi, \Gamma \vdash t \text{id}; : \epsilon, \Gamma[\text{id} := \tau]}$	T-DECL	$\frac{\tau = t \quad \text{alloc}(\sigma, \tau) = a}{\Phi, t \text{id};, \sigma, M \rightarrow ;, \sigma \left[ \begin{array}{l} T := T[\text{id} := \tau], \\ A := A[\text{id} := a] \end{array} \right], M}$	E-DECL
$\frac{\Phi, \Gamma \vdash e_1 : \tau, \Gamma \quad \Phi, \Gamma \vdash e_2 : \tau, \Gamma}{\Phi, \Gamma \vdash e_2 = e_1; : \epsilon, \Gamma}$	T-ASSIGN	$\frac{\Phi, \mathcal{L}(a, \tau) = \mathcal{V}(m, \tau);, \sigma, M \rightarrow ;, \sigma, M := M[\langle a, \tau \rangle := m]}$	E-ASSIGN
$\frac{\Phi, \Gamma \vdash e : \text{bool}, \Gamma \quad \Phi, \Gamma \vdash s_1 : \epsilon, \Gamma_1 \quad \Phi, \Gamma \vdash s_2 : \epsilon, \Gamma_2}{\Phi, \Gamma \vdash \text{if}(e) s_1 \text{ else } s_2 : \epsilon, \Gamma}$	T-IF	$\frac{\Phi, \text{if}(\mathcal{V}(T, \text{bool})) s_1 \text{ else } s_2, \sigma, M \rightarrow \{s_1\}^{\sigma'}, \sigma', M}{\Phi, \text{if}(\mathcal{V}(F, \text{bool})) s_1 \text{ else } s_2, \sigma, M \rightarrow \{s_2\}^{\sigma'}, \sigma', M}$	E-IF E-IF
$\frac{\Phi, \Gamma \vdash e : \text{bool}, \Gamma \quad \Phi, \Gamma \vdash s : \epsilon, \Gamma'}{\Phi, \Gamma \vdash \text{while}(e) s : \epsilon, \Gamma}$	T-WHILE	$\frac{\Phi, \text{while}(e) s, \sigma, M \rightarrow \text{if}(e) \{s \text{ while}(e) s\}, \sigma, M}$	E-WHILE

**Figure 7.** Static (left) and dynamic (right) semantics of IMP (excerpt)

	$\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma$		$\Phi, a, \sigma, M \rightarrow a', \sigma', M'$
$\frac{\text{id} \in \text{dom}(\Gamma) \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash \text{id} : \tau(\text{id}), \Gamma}$	T-VAR	$\frac{\text{id} \in \text{dom}(T_\sigma) \quad \vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, \text{id}, \sigma, M \rightarrow \mathcal{L}(A_\sigma(\text{id}), T_\sigma(\text{id})), \sigma, M}$	E-VAR
$\frac{\Phi, \Gamma, \Lambda \vdash \text{id}; : \epsilon, \Gamma'}{\Phi, \Gamma, \Lambda \vdash t \text{id}; : \epsilon, \Gamma'}$	T-DECL	$\frac{\Phi, t \text{id};, \sigma, M \rightarrow t^{\ P_\sigma\ } \text{id};, \sigma, M}{\Phi, t^{\ P_\sigma\ } \text{id};, \sigma, M \rightarrow ;, \sigma \left[ \begin{array}{l} T := T[\text{id} := \tau], \\ A := A[\text{id} := a] \end{array} \right], M}$	E-DECL
$\frac{\tau = \text{vec}_\Lambda(t, l) \quad \tau \neq \perp \quad \text{id} \notin \text{dom}(\Gamma) \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash t \text{id}; : \epsilon, \Gamma[T := T[\text{id} := \tau]]}$	T-DECLL	$\frac{\tau = \text{vec}_\Lambda(t, l) \quad \text{alloc}(\sigma, \tau) = a \quad \vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, t^{\ P_\sigma\ } \text{id};, \sigma, M \rightarrow ;, \sigma, M := M[\langle P_\sigma, a, \tau \rangle := m]}$	E-DECLL
$\frac{\Phi, \Gamma, \Lambda \vdash e_1 : \tau, \Gamma \quad \Phi, \Gamma, \Lambda \vdash e_2 : \tau, \Gamma \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash e_2 = e_1; : \epsilon, \Gamma}$	T-ASSIGN	$\frac{\vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, \mathcal{L}(a, \tau) = \mathcal{V}(m, \tau);, \sigma, M \rightarrow ;, \sigma, M := M[\langle P_\sigma, a, \tau \rangle := m]}$	E-ASSIGN
$\frac{\Phi, \Gamma, \Lambda \vdash e_1 : \tau, \Gamma \quad \Phi, \Gamma, \Lambda \vdash e_2 : \nu, \Gamma \quad \Lambda = 1 \vee \Lambda = \ \nu\  \quad \text{vec}_{\ P_\sigma\ }(\tau, \ \nu\ ) = \nu}{\Phi, \Gamma, \Lambda \vdash e_2 = e_1; : \epsilon, \Gamma}$	T-BROADCAST	$\frac{\ P_\sigma\  = 1 \vee \ P_\sigma\  = \ \nu\  \quad \text{vec}_{\ P_\sigma\ }(\tau, \ \nu\ ) = \nu}{\Phi, \mathcal{L}(a, \nu) = \mathcal{V}(m, \tau);, \sigma, M \rightarrow ;, \sigma, M := M[\langle P_\sigma, a, \nu \rangle := \langle \vec{m}^+ \rangle]}$	E-BROADCAST
$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \Phi, \Gamma, \Lambda \vdash t^{\ \text{id}\ } \text{id}; \text{id} = e; : \epsilon, \Gamma'}{\Phi, \Gamma, \Lambda \vdash t \text{id} = e; : \epsilon, \Gamma'}$	T-INIT	$\frac{\Phi, t \text{id} = \mathcal{V}(m, \tau);, \sigma, M \rightarrow t^{\ \text{id}\ } \text{id}; \text{id} = \mathcal{V}(m, \tau);, \sigma, M}{\Phi, t \text{id} = \mathcal{V}(m, \tau);, \sigma, M \rightarrow t^{\ \text{id}\ } \text{id}; \text{id} = \mathcal{V}(m, \tau);, \sigma, M}$	E-INIT
$\frac{\Phi, \Gamma, \Lambda \vdash e : \text{bool} \mid l, \Gamma \quad \Lambda = 1 \vee \Lambda = l \quad \Phi, \Gamma, \Lambda := l \vdash s_1 : \epsilon, \Gamma_1 \quad \Phi, \Gamma, \Lambda := l \vdash s_2 : \epsilon, \Gamma_2}{\Phi, \Gamma, \Lambda \vdash \text{if}(e) s_1 \text{ else } s_2 : \epsilon, \Gamma}$	T-IF	$\frac{\Lambda = 1 \vee \Lambda = l}{\Phi, \text{if}(\mathcal{V}(m, \langle \text{bool} \mid l \rangle)) s_1 \text{ else } s_2, \sigma, M \rightarrow \{s_1\}^{\sigma'} \text{ if } (\mathcal{V}(\text{not } m, \langle \text{bool} \mid l \rangle)) s_2 \text{ else } ;, \sigma[P := P \text{ and } m], M}$	E-IF
$\frac{\Phi, \Gamma, \Lambda \vdash e : \text{bool} \mid l, \Gamma \quad \Lambda = 1 \vee \Lambda = l \quad \Phi, \Gamma, \Lambda := l \vdash s : \epsilon, \Gamma'}{\Phi, \Gamma, \Lambda \vdash \text{while}(e) s : \epsilon, \Gamma}$	T-WHILE	$\frac{\Phi, \text{while}(e) s, \sigma, M \rightarrow \text{if}(e) \{s \text{ while}(e) s\}, \sigma, M}{\Phi, \text{while}(e) s, \sigma, M \rightarrow \text{if}(e) \{s \text{ while}(e) s\}, \sigma, M}$	E-WHILE

**Figure 8.** Static (left) and dynamic (right) semantics of VecIMP (excerpt)

the type match:

$$\vec{\tau}: \mathbb{P} \times \mathcal{V} \rightarrow \text{bool}$$

$$(\Lambda, \tau) \mapsto \Lambda = 1 \vee \|\tau\| = 1 \vee \Lambda = \|\tau\|$$

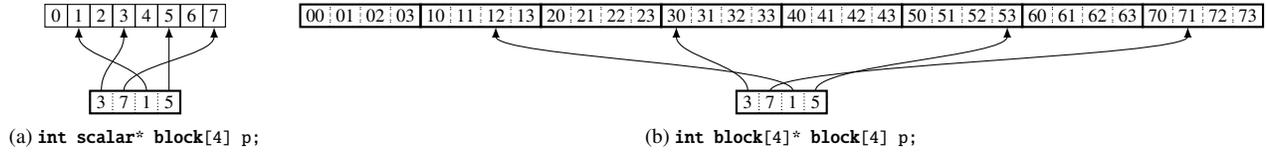
where  $\|\tau\|$  denotes the vector length of type  $\tau$ . Moreover, a scalar variable on the right-hand side of an assignment can be *broadcast* to a vector variable on the left-hand side. This is accomplished by replicating the given value to all lanes (see T- and E-BROADCAST).

## 4.2 Triggering Control-Flow Vectorization

Conditionally executing code is usually achieved by testing a boolean and either running one or another code path. IMP supports typical if- and while-statements. In VecIMP we enhance these constructs by allowing boolean vectors as conditional expressions as well. Generally, the vector length of the header must either match

the one of the current mask, or a scalar conditional occurs in a vectorial context, or the conditional is in a scalar context. The latter case triggers vectorization. This means that control-flow is converted to data-flow which gives the programmer the illusion that different SIMD lanes are executed on different control-flow branches. Care must be taken when different branches have side-effects like I/O-accesses. These side-effects may happen in an unexpected order. For this reason, it is important that the programmer at least roughly understands, how the control-flow vectorization is accomplished. For the sake of simplicity, we concentrate on simple structural control-flow in this paper. For vectorization of arbitrary control-flow we refer to Karrenberg and Hack [16].

**If/Else.** In VecIMP, both the then-block and the else-block are evaluated (see T-IF and E-IF). In the then-block the current mask and the test mask are combined with an AND operation whereas



**Figure 10.** Involved address calculation when indexing a vectorial pointer  $p$  with an index vector  $\text{int } \text{block}[4] \ x = \langle 3, 7, 1, 5 \rangle$ . A read  $p[x]$  obtains in case a)  $\langle 3, 7, 1, 5 \rangle$  and in case b)  $\langle 30, 71, 12, 53 \rangle$ . Note that both variants yield an  $\text{int } \text{block}[4]$ .

the else-block uses the complementary test mask. Note that an if statement degenerates to standard IMP semantics if an ordinary boolean is used in the header, since statements are not executed if the predication variable is completely false.

**While.** We define the semantics of the while-construct as a recursive program transformation to an if-statement containing a while-statement (see T-WHILE and E-WHILE). If every element of this mask evaluates to false, the body is not executed at all. Otherwise, the body must be entered while adjusting the current mask. Analogously to the if-statement, a while-statement degenerates to an ordinary while-statement when testing a scalar boolean.

#### 4.2.1 The Scalar Statement

Sometimes it is desirable to deactivate vectorization within a kernel and proceed with scalar computations (see the introductory example in Section 1). This can be realized by adding a *scalar statement* to the syntax. This statement saves the current mask, sets the current mask to the scalar *true* value, executes the inner statement in a scalar context, and restores the saved mask afterwards. Using the scalar statement, several sequential patterns of algorithms using SIMD instructions can be implemented.

**Reactivating the old mask.** Although the old mask is reset after termination of the scalar statement, it may be nevertheless useful to temporarily reactivate the mask within this statement. This can be achieved by manually saving the current mask before executing the scalar statement and using the saved mask in an if-statement. Since a store just sets the active lanes, it is also important to set the inactive lanes manually to *false*. The example in Figure 4 makes use of some syntactic sugar in order to initialize elements in both active and inactive lanes. This syntactic sugar

```
T var = [true_values, false_values];
```

translates to:

```
bool mask; // a mask; we are in a vector context
scalar { mask = false; } // set all lanes to false
mask = current_mask; // set all active lanes to true
T var;
scalar {
  if (mask) var = true_values; // set active lanes
  else var = false_values; // set inactive lanes
}
```

**Scalarize each active lane.** Using a consecutive bit scan over the old mask can yield the index of each prior active lane. In the following example each active component of  $v$  of type  $T \ \text{block}[N]$  is extracted:

```
bool mask = [current_mask, current_mask];
scalar
  for(int i=bitscan(mask,0); i>=0; i=bitscan(mask,i+1)) {
    T s = get(v, i);
    //...
  }
```

A `bitscan` scans a given mask, beginning at the given index, from left to right for the next entry set to *true*. In order to indicate the case that no further *true* follows, a negative value is returned.

The loop body can now do scalar computations with  $v$ , broadcast this value in order to vectorize another computation, or collect such values in another data structure which can be later on used for further vectorizations, for instance.

**Uniquely scalarize each active lane.** A more sophisticated technique only extracts unique values from active lanes. In the example above the following adjustment of the mask in each iteration before the next `bitscan` blends out duplicates:

```
mask &= s != v;
```

#### 4.3 Address Calculation

Since both a pointer itself and its referenced type may be scalar or vectorial, four possibilities emerge: Scalar pointers to scalar data are ordinary pointers as known from C. Scalar pointers to vectorial data integrate seamlessly into IMP as the layout of the referenced type is irrelevant for typing and semantics. However, when dealing with vectorial pointers, the programmer must use an index vector. Vectorial indexing to scalar data effectively performs a S/G (see Figure 10a) and obtains a vector. Vectorial indexing to vectorial data is more involved as depicted in Figure 10b: Each lane  $i$  “sees” the  $i^{\text{th}}$  component of a vector located at address  $\text{get}(p, i) + \text{get}(x, i)$  and hence read/write access with this vectorial indexing gathers/scatters an  $\text{int } \text{block}[4]$ .

#### 4.4 Parametric Polymorphism

An important feature of VecIMP is that all functions are parametric towards vectorization. This means that all unbound types which appear in the signature and the vector context itself are parametric. Upon invocation a new version of the function is instantiated (like templates in C++). This has several consequences:

- Type checking and evaluation is done with the current context/mask. This changes the vector length of all `block` variables and unbound variables in declaration statements.
- All unbound function arguments try to adopt the vector length of the parameter (like an initialization).
- If the return type is unbound the type checker infers its vector length from the return expression.
- As the semantics of all syntactical constructs depends on the instantiated types, the behaviour of the function is also dependent on these parameters.

#### 4.5 Soundness

All these additions and changes to IMP do not harm the type system and using the same definition of soundness as in Theorem 1 (see Appendix B for details) we can state:

**Theorem 2** (Soundness of VecIMP). *VecIMP’s type system is sound.*

## 5. The IVL Vectorizing Language

IVL is a prototypical R&D compiler that applies many of the proposed concepts to generate vectorized code for the Intel<sup>®</sup> MIC architecture and SSE instruction set.

### 5.1 IVL’s Back-Ends

MIC is a many-core x86 architecture in which each x86 core is augmented with a 16-wide vector unit accessed through a rich vector instruction set. MIC supports both S/G as well as efficient masking/predication via a separate set of 16-bit mask registers.

In a vectorized context IVL emits MIC code that runs on  $L$  lanes in parallel, while all *scalar* constructs are executed in the core’s scalar pipe, and scalar data is held in scalar registers. Rather than allowing arbitrary vector lengths, IVL currently only supports a small set of vectorization lengths per compilation unit. In particular, on MIC we currently support 16-wide and (“double-pumped”) 32-wide vectorization for MIC (with 16 being the default). The SSE back-end supports 4 wide, and a future Intel® AVX back-end will also support 8 wide vectorization.

Note that MIC and SSE instruction sets fundamentally differ in many aspects: SSE and MIC have different native vector lengths, SSE does not support hardware accelerated S/G and uses blending in order to emulate predicated execution. Given our experience in implementing the SSE and MIC back-ends, we think implementing additional back-ends supporting other SIMD CPU types and ISAs is straightforward.

IVL acts as source-to-source compiler in the sense that it translates IVL code to a special “C++ with intrinsics” code that is then passed to the official Intel® C/C++ compiler [11] (icc) for MIC. This allows the programmer to visually inspect (and possibly modify) the emitted code, and to use this code with tools like debuggers, performance analyzers, etc.

## 5.2 Supported Types and Language Constructs

IVL currently supports a significant subset of C and a small set of additional keywords to guide vectorization. In terms of types, IVL fully supports bools as well as (32-bit) ints, uints, and floats, but currently only partially supports 8-bit, 16-bit, and 64-bit data types. IVL also fully supports structs, arrays, and references (including vector references to vectorized types), but only partially supports pointers.

As IVL currently only supports one global vector length  $N$  per compilation unit, `VecImp`’s `block[N]` is called **varying**, and `block[1]` is called **uniform**. Like `VecImp`, IVL supports *unbound* variables, and full polymorphism in parameter type length. In terms of control-flow IVL supports all of `if/else`, `do`, `while`, `for`, `break`, `continue`, and `return`, with the sole exception of `goto`. In addition, IVL also supports some simple reduction operations like “**varying** `bool` `any`(**varying** `bool`)” or “**varying** `bool` `none`(**varying** `bool`)” in order to determine whether at least one or no element in a boolean vector is *true*.

## 5.3 On-demand Vectorization

Vectorization in IVL is done *on demand*: Similar to the behavior of templates in C++, IVL parses struct and function definitions but does not emit anything until *instances* of those types and functions are required. For example, when IVL encounters a global variable of a vectorized type, it emits a vectorized form of this type (in C++ code). Vectorization of code is triggered when IVL encounters a function with the `kernel` keyword; it will then emit a C function for this kernel (plus some additional helper functions to allow calling this function from the host machine if required), and will vectorize this kernel’s body, which in turn will on-demand emit all functions called by this body, etc.

Statements and expressions operating on varying types will emit vector intrinsics, while purely uniform expressions/statements will emit only scalar code even when inside a vectorized function (like proposed by `VecImp`). This is highly desirable in that it uses precious vector registers and costly vector instructions only where required, and enables a mix of scalar and vector expressions that MICs superscalar architecture (with parallel scalar  $U$  and vector  $V$  pipes) is particularly good at. For example, if a for-loop in a (vectorized) function uses a loop condition that only depends on a uniform function parameter that respective loop control code will use only scalar x86 instructions.

**Polymorphism in vector length.** Polymorphism also works like proposed in `VecImp`: Unbound function arguments are bound when IVL encounters its call site, and are then bound to the vector length

of the respective argument. This means that the same *logical* function can end up being emitted in multiple incarnations (depending on which of its parameters are uniform or varying). This, too, is highly desirable in that information about which data are uniform is *preserved* across function calls—usually leading to significantly more efficient code and storage than if all scalar data had been broadcast to vector form upon the first function call (as would happen in a pure SIMT paradigm). In particular, a given control-flow statement in an input function can get emitted in scalar form in one function instance, and vectorized in another.

Like polymorphism in C++, this mechanism requires the function to be known at compile-time; to use a function across different compilation units one either has to explicitly specify all parameters’ vector lengths, or explicitly instantiate this function in one compilation unit.

## 5.4 IVL Examples

Though a full performance analysis is beyond the scope of this paper, we give a brief overview over some examples realized with our current IVL compiler, all running on a 32-core 900MHz Knights Ferry prototype board. Since IVL also has an SSE back-end, the exactly same examples will also run on any SSE-enabled architecture, but for the sake of brevity we restrict ourselves to only the Knights-Ferry results.

**Proof-of-concept examples.** We ported Mandelbrot, `nbody`, and `VolumeRender` from the CUDA SDK. These examples ran more or less “out of the box”. Since IVL does not currently support any native hardware-texturing, the `VolumeRender` example has to resort to “manual” tri-linear interpolation to sample the volume, but nevertheless already reaches roughly 30 frames per second. For all other examples, both ease of porting and resulting performance matched or exceeded expectations. As just one example, the publicly available `aobench` benchmark [3] required only trivial modification to port to IVL,<sup>3</sup> while rendering a 1024<sup>2</sup> frame (with 16 samples per pixel and 16 rays per sample) in an impressive 402ms.

**Ray tracer examples.** As a somewhat more challenging example, we also implemented an IVL ray tracer with various shaders *into* an existing Knight’s Ferry ray tracing system [35]. The IVL-based traversal, intersection, and shading code was linked together with manual intrinsics code for data structure construction and other renderers. In this setting, IVL and manual C code were actually sharing the same data structures! First, we integrated an intentionally simple *eyelight* shader into the framework. As a next step, we included an *ambient occlusion* (AO) renderer requiring random number generation, quasi-Monte Carlo sampling, CDF inversions and recursion, involving both incoherent data access patterns and significant SIMD divergence.

**Comparison to hand-written code.** To better quantify the performance of IVL’s code we also ran some experiments where we compared IVL-generated code to manually-written reference intrinsics code (in all cases, the reference code was written *before* the IVL code). For the *eyelight* ray tracer, IVL renders a 1600 × 1024 frame in 13.8 million cycles, vs 13.28m cycles in reference intrinsics code, a difference of only 4% (an exact reference version for the AO renderer is not available, due to that code’s complexity).

Finally, we also ran an artificial *k-nearest neighbor* benchmark for which we had reference code for a variety of architectures. This benchmark is highly non-trivial in both control-flow and data access patterns. In this workload, IVL requires 755 million cycles (for 1 million 50-neighbor queries in a 1 million point dataset), as compared to 656m cycles for the hand-coded version—a difference of 15%.

<sup>3</sup> For example, IVL does not support doubles, yet, so both the IVL and the reference implementation have been modified to use floats.

The performance gap is mostly encumbered by that fact that IVL is a source-to-source compiler emitting vectorized code and relies on `icc` for optimizations. This compiler does not yet recognize some patterns which are suitable for optimizations but would have been applied by a human intrinsics programmer. `ispc` (see Section 6) on the other hand generates LLVM assembly code and ships its own set of specific optimizations for such patterns.

## 6. Related Work

While some efforts exist to integrate short vector types into languages, over the past years several data parallel languages have been developed in order to tackle programming for SIMD hardware. These languages can be basically subdivided into *data-parallel languages* and *vector (or array) programming languages*. Independent from that, several automatic vectorization techniques exist.

**Short vector data types.** With the help of C++ operator overloading or built-in support for short vector types in C [9, sec. 6.49], standard operator syntax can be used instead of writing intrinsics. This does not only increase readability but also portability of the source code as each back-end can resort to its specific SIMD ISA. However, type and control-flow vectorization must still be handled manually by the programmer. Especially, manual conversion of control-flow to data dependence is problematic as this—besides being very error-prone—again introduces architectural dependent low-level details to the source code: This approach has no means of abstracting blending/predicated execution.

**Data-parallel languages.** OpenCL and CUDA have already been discussed in Section 1. These languages are strongly influenced from shading languages whose programming model is similar but more limited than the former ones. One of the first shading languages implemented is RenderMan [10] which originally targeted a virtual SIMD array machine. Programs for this language are also written for a single scalar thread, while the compiler and the runtime system assured to instantiate this thread for all available SIMD lanes. RenderMan pioneered the concept of *uniform* and *varying* variables. *Uniform* variables are only set once for each set of properties. *Varying* variables on the other hand are allowed to change as a function of position. For instance, vertex normals in a Phong shader must be interpolated and must therefore be declared as *varying*.

Our work is close to the Intel® SPMD Program Compiler (`ispc`) [12]. This LLVM based compiler, where a program is instantiated  $n$  times to run on  $n$  SIMD lanes, currently supports several SSE revisions and AVX. Although IVL and `ispc` are two independent projects, we fostered an ongoing discussion with the `ispc` authors, which influenced each other’s work. Thus, `ispc` and IVL/VecImp share many concepts: implicit handling of blending, automatic control-flow vectorization and *uniform* and *varying* variables among other features. There are also some differences: `ispc` does not support parametric polymorphism while `ispc`’s pointer support is more mature than IVL’s one, for instance.

Another data-parallel language is  $C^n$  [21] which is in contrast to many other languages a real extension to C.  $C^n$  targets a special kind of SIMD machine, similar to IBM’s Cell processor. There, a main processor controls  $n$  processing elements (PEs) that all have their own local memory.  $C^n$  extends the type system by two modifiers `mono` and `poly`. `poly` values sit in the memory of the PEs. The type system is used to trigger DMA transfers from the main CPU to the PEs. However, due to this different hardware model,  $C^n$  does not support `poly` values in compound data structures. More precisely, `mono` and `poly` are *storage-class specifiers* and not *type qualifiers*.

**Vector programming languages.** These languages are very powerful when dealing with mathematical vectors and matrices, i.e., arrays of atomic types. But this is also the main problem: Since

arrays are a too high-level abstraction for short vector SIMD processors, it is very difficult to write efficient code for these machines when dealing with higher data structures. There does not exist any built-in way in order to deal efficiently with compound data, as this would contradict the array programming philosophy.

In some languages of this category like APL [15] or Matlab [22], the programmer has to manually vectorize control-flow. Others like Vector C [20], Vector Pascal [24], Intel’s Array Building Blocks (ArBB [23], formerly known as Ct [8]) or Nesl [4] provide special constructs and/or enhance scalar constructs to work with vectors for this task.

Nesl integrates nested vectors as first class citizens in a pure functional language similar to ML. However, Nesl does not integrate imperative control structures and low-level constructs like pointers as we do.

ArBB/Ct is inspired by Nesl and can be seen as a Nesl library for C/C++. ArBB is a functional array language with an interface to C++. At first glance, it looks like a language extension to C++ for vector processing. However, ArBB uses operator overloading to construct an intermediate representation of the vector program which is compiled and executed *lazily* by a just-in-time compiler in the ArBB library. While this has several advantages, like adaption of vectorization to the system the program is currently running on, it has the disadvantage of having a second language within another language: ArBB requires own types (`i32` instead of `int`, etc.), own control structures (`_for ... _end_for` instead of the common `for` loop), and sometimes own operators (`(call)(f)(a, b)` instead of `f(a, b)`). Furthermore, the user has to understand the execution model precisely: ArBB expressions are used to *construct* a program. Putting `a+=b` in an *ordinary* loop will create as many expressions in the IR as the loop iterates. VecImp however combines both, scalar and vector computing in the same language.

**Automatic vectorization.** Vectorizing compilers go back to the early 1980s [1, 2]. Basic predication techniques have been introduced in this line of work [1]. Many different automatic vectorization techniques have been developed since then. The majority of approaches target vectorization of straight-line code loop bodies [27, 34]. There, instruction-level parallelism (ILP) is created by loop unrolling, or more sophisticated transformations (e.g. polyhedral techniques [5]) are applied. Other approaches exploit ILP in code that can contain control structures (like conditionals and loops) [18, 19, 32] for vectorization or vectorize a whole function [16]. Often, those approaches rely on preceding transformations that create this ILP. Outer-loop vectorization [25, 28, 30] goes beyond a loop body by integrating inner loops into vectorization. Another idea is to use the SIMD register file as fast cache [33].

As stated in the introduction, automatic vectorization is not applicable in our setting. The modifications to the data structure layout and the resulting changes in the code are not local to some code part. They often affect the whole program. Furthermore, algorithms substantially change, when they are programmed with SIMD instructions in mind. Current static analyses and code transformations are not able to perform such transformations automatically in a satisfactory way. Our work focuses on the programmer who explicitly wants to program for a CPU equipped with a SIMD instruction set without exposing her to assembly-level programming. This clientele is not satisfied with automatic techniques that only work on small code parts and depend on a preceding analysis that is not transparent to the programmer.

## 7. Conclusions and Future Work

The first practical experiences show that code emitted by IVL is comparable to hand-written intrinsic code. Smarter compiler transformations and back-ends would even produce higher quality code.

**A hybrid SoA container class.** As a next step, we would like to integrate this model into a C/C++ compiler. VecImp only changes

IMP where necessary and the added features are only minimum requirements in order to describe low-level SIMD programming in an abstract way. Often used programming patterns can be further abstracted with standard higher-level language features in order to reduce boilerplate code and increase portability. In particular, the C++ template mechanism would allow to define a hybrid SoA container class: Similar to `std::vector` which abstracts a traditional C array, one could implement a wrapper around a `T block[N]*`:

```
// scalar context throughout this example
struct vec3 { float x, y, z; };
// vec3 block[N]* pointing to ceil(n/N) elements
hsoa<vec3> vecs(n);
// preferred vector length of vec3 automatically derived
static const int N = hsoa<vec3>::vector_length;
int i = /*...*/
hsoa<vec3>::block_index ii = /*...*/
vec3 v = vecs[i]; // gather
vecs[i] = v; // scatter
vec3 block[N] w = vecs[ii]; // fetch whole block
hsoa<vec3>::ref r = vecs[i]; // get proxy to a scalar
r = v; // pipe through proxy
// for each element
vecs.foreach([](vec3& scalar v) { /*...*/ });
```

The inner type `hsoa<vec3>::ref` consists of a pointer pointing to the block in question and a lane index. Furthermore, this class is equipped with some conversion operators in order to automatically convert this proxy to the scalar base type (`vec3` in this case).

The `foreach` method is particularly interesting: Most of the time, i.e., while not processing one of the two border areas, each iteration can fetch a whole `vec3 block[N]& scalar`. The border areas must be handled specially—either by setting up an appropriate mask, or by scalarizing the border sections. This special setup is factored out in the `foreach` method which invokes different versions of the polymorphic lambda function and the user of `foreach` does not have to care about border areas anymore.

**Block hierarchy.** One could also imagine having a hierarchy of blocks. This information would also be annotated at the type and could be used in order to abstract blocks/work groups in CUDA/OpenCL or the layers of the cache hierarchy like cache lines and page sizes. A long term goal would be a language which emits SIMD CPU code and PTX/CUDA code dependent on the types which trigger the invocation of a kernel. Additional setup code would also be automatically generated.

## Acknowledgments

The authors would like to thank Sigurd Schneider and Jan Schwinghammer who helped with the formalization of the type system. Furthermore, Ralf Karrenberg, Christoph Mallon and Phillip Slusallek were involved in the discussion of our work. William M. Mark, Matt Pharr and Sven Woop from Intel also gave a lot of valuable feedback and in particular the aforementioned collaboration regarding Matt's `ispc` project was crucial to the realization of IVL. Much of this work has been funded by the Intel Visual Computing Institute and the German ministry for education and research via the ECOUSS project.

## References

- [1] J. R. Allen, K. Kennedy, C. Porterfield, and J. Warren. Conversion of Control Dependence to Data Dependence. In *POPL*, 1983.
- [2] R. Allen and K. Kennedy. Automatic Translation of FORTRAN Programs to Vector Form. *ACM Trans. Program. Lang. Syst.*, 1987.
- [3] aobench. URL <http://code.google.com/p/aobench/>.
- [4] G. E. Blelloch et al. Implementation of a Portable Nested Data-Parallel Language. In *PPOPP*, 1993.
- [5] A. Darte, Y. Robert, and F. Vivien. *Scheduling and Automatic Parallelization*. Birkhauser Boston, 2000.
- [6] M. Farrar. Striped Smith–Waterman speeds database searches six times over other SIMD implementations. *Bioinformatics*, 23:156–161, January 2007.

- [7] I. Georgiev and P. Slusallek. RTfact: Generic Concepts for Flexible and High Performance Ray Tracing. In *IEEE/Eurographics Symposium on Interactive Ray Tracing*, 2008.
- [8] A. Ghuloum et al. Future-Proof Data Parallel Algorithms and Software on Intel Multi-Core Architecture. *Intel Technology Journal*, 11(04), November 2007.
- [9] GNU Press. Using the GNU Compiler Collection. For GCC version 4.6.2.
- [10] P. Hanrahan and J. Lawson. A Language for Shading and Lighting Calculations. In *SIGGRAPH*, 1990.
- [11] Intel Corp. Intel® Compilers and Libraries. URL <http://software.intel.com/en-us/articles/intel-compilers>.
- [12] Intel Corp. Intel SPMD Program Compiler. URL <http://ispc.github.com>.
- [13] Intel Corp. Intel® 64 and IA-32 Architectures Optimization Reference Manual, 2009.
- [14] Intel Corp. The Intel Many Integrated Core (MIC) Architecture, 2010.
- [15] K. E. Iverson. *A Programming Language*. John Wiley & Sons, Inc., 1962.
- [16] R. Karrenberg and S. Hack. Whole Function Vectorization. In *CGO*, 2011.
- [17] Khronos Group. *OpenCL 1.0 Specification*, 2009.
- [18] A. Krall and S. Lelait. Compilation Techniques for Multimedia Processors. *Int. J. Parallel Program.*, 28(4):347–361, 2000.
- [19] S. Larsen and S. Amarasinghe. Exploiting Superword Level Parallelism with Multimedia Instruction Sets. *PLDI*, 35(5):145–156, 2000.
- [20] K.-C. Li and H. Schwetman. Vector C—A Vector Processing Language. *Journal of Parallel and Distributed Computing*, 2(2):132 – 169, 1985.
- [21] A. Lokhmotov, B. R. Gaster, A. Mycroft, N. Hickey, and D. Stuttard. Revisiting SIMD Programming. In *LCPC*, pages 32–46, 2007.
- [22] MatLab. URL <http://www.mathworks.com/products/matlab>.
- [23] M. McCool. A Retargetable, Dynamic Compiler and Embedded language. In *CGO*, 2011.
- [24] G. Michaelson and P. Cockshott. Vector Pascal, an array language, 2002.
- [25] V. Ngo. *Parallel Loop Transformation Techniques For Vector-Based Multiprocessor Systems*. PhD thesis, University of Minnesota, 1994.
- [26] M. Norrish. *C formalised in HOL*. PhD thesis, University of Cambridge, 1998.
- [27] D. Nuzman and R. Henderson. Multi-platform Auto-vectorization. In *CGO*, 2006.
- [28] D. Nuzman and A. Zaks. Outer-Loop Vectorization: Revisited for Short SIMD Architectures. In *PACT*, 2008.
- [29] NVIDIA. *CUDA Programming Guide*, 2009.
- [30] R. G. Scarborough and H. G. Kolsky. A vectorizing Fortran compiler. *IBM J. Res. Dev.*, 30(2):163–171, 1986.
- [31] L. Seiler et al. Larrabee: A Many-Core x86 Architecture for Visual Computing. In *SIGGRAPH*, 2008.
- [32] J. Shin. Introducing Control Flow into Vectorized Code. In *PACT '07*, 2007.
- [33] J. Shin, C. Jacqueline, and M. W. Hall. Compiler-Controlled Caching in Superword Register Files for Multimedia Extension Architectures. In *PACT*, 2002.
- [34] N. Sreeramam and R. Govindarajan. A Vectorizing Compiler for Multimedia Extensions. *Int. J. Parallel Program.*, 28(4):363–400, 2000.
- [35] I. Wald. Fast Construction of SAH BVHs on the Intel® Many Integrated Core (MIC) Architecture. *IEEE Transactions on Visualization and Computer Graphics*, 99, 2010.
- [36] J. Zhou and K. A. Ross. Implementing Database Operations Using SIMD Instructions. In *SIGMOD*, 2002.

## A. The Semantics of IMP

In this section we present we introduce the remaining semantic rules of IMP along with the used notation. Then, we formally prove the soundness of IMP's type system.

IMP's type checker works by first registering each function in a function map  $\Phi$  (see Figure 11). The function map  $\Phi$  maps the function's identifier  $\text{id}$  to the tuple  $(\overline{t \text{id}}, s, e_r, t)$  where  $\overline{t \text{id}}$  is the function's parameter list,  $s$  its body,  $e_r$  its return expression and  $t_r$  its return type. Then, all functions are type checked in order to examine whether the program is well-formed:

**Definition 1** (Well-formed Program). *A function map  $\Phi$  is called well-formed iff*

$$\forall (\overline{t \text{id}}, s, e_r, t_r) \in \Phi : \\ \Phi, \Gamma := \emptyset \vdash \overline{t \text{id}}; s \text{ return } e_r; ; {}_r \tau_r, \Gamma' \quad \text{with } t_r = \tau_r.$$

Since the evaluation of an expression may be needed as lvalue we cannot simply write a rule which transforms all lvalues to values. For this reason, rule E-TOVAL only converts an lvalue which is not the root of an expression tree to a value. General expression evaluation order is determined by rule E-EVAL. These rules act as templates for various statements/expressions. The used evaluation contexts are defined as follows:

$$\begin{aligned} \mathcal{E}[\star] &::= \mathcal{L}[\star] \mid \& \star \mid \star. i \mid \star = \mathcal{V}(m, \tau) \mid \star s \mid \{ \star \}^\sigma \\ \mathcal{L}[\star] &::= \star \mid \star[e] \mid \mathcal{V}(m, \tau)[\star] \mid \text{id}(\overline{\mathcal{V}(m, \tau)}, \star, \bar{e}) \\ &\mid \star; \mid e = \star \mid t \text{id} = \star; \\ &\mid \text{if } (\star) s_1; \text{else } s_2; \mid \text{return}^\sigma \star; \end{aligned}$$

Projections of records are not caught by rule E-TOVAL. Instead two versions are needed: One evaluating to values (see E-RPROJ) and another one to lvalues (see E-LPROJ). We use a function

$$\text{offset}_{\lceil \tau \rceil} : \{0, \dots, \lceil \tau \rceil - 1\} \rightarrow \mathbb{N}$$

in order to obtain the (implementation-dependent) offset in bytes of the given component.

The function call rule E-VOKE reduces the call to the body of the called function: First, all parameters are initialized (via E-INIT) with their associated arguments. Next, the body is executed and finally, the return expression is evaluated. The return expression memorized the old scope  $\sigma$  before the call and evaluating "return $^\sigma e$ ;" restores the old state again (E-RETURN). This ensures that scope updates are simply "forgotten" when returning to the caller while the memory state stays preserved. During evaluation the type and address map of a scope must stay in sync:

**Definition 2** (Well-formed scope list). *A scope  $\sigma$  is called well-formed iff*

$$\forall \text{id} \in T_\sigma : \text{id} \in A_\sigma \quad \vee \quad \forall \text{id} \in A_\sigma : \text{id} \in T_\sigma.$$

Some rules make use of a function  $V_\tau : \mathbb{I}8^{|\tau|} \rightarrow \tau$  which maps a sequence of  $|\tau|$  bytes to a value (E-CONST and E-INDEX). Indexing is allowed with all atomic integer types—denoted by  $\mathcal{A}_{\text{int}}$ .

### A.1 The Soundness Proof for IMP's Type System

**Lemma 1** (Progress of IMP). *Let  $\Phi$  be a well-formed function map,  $\sigma$  be a well-formed scope,  $a$  be an IMP program and  $M$  be a memory state. Then holds*

$$\begin{aligned} \Phi, \Gamma := T_\sigma \vdash a : v, \Gamma' \\ \Rightarrow \begin{cases} a \text{ is a value} & \text{if } v = {}_r \tau \\ a \text{ is an lvalue} & \text{if } v = {}_l \tau \\ a \text{ is a value or an lvalue} & \text{if } v = {}_r \tau \\ a \text{ is " ; " } & \text{if } v = \epsilon \end{cases} \\ \vee \quad \exists a', \sigma', M' : \Phi, a, \sigma, M \rightarrow a', \sigma', M' . \end{aligned}$$

Functions	$\Phi \vdash f \triangleright \Phi'$
$\overline{\Phi \vdash \epsilon \triangleright \Phi}$	F-EMPTY
$\frac{\Phi_0 \vdash f_1 \triangleright \Phi_1 \quad \Phi_1 \vdash f_2 \triangleright \Phi_2}{\Phi_0 \vdash f_1 f_2 \triangleright \Phi_2}$	F-SEQ
$\frac{\text{id} \notin \text{dom}(\Phi)}{\Phi \vdash t_r \text{id}(\overline{t \text{id}}) \{ s \text{ return } e_r; \} \triangleright \Phi[\text{id} := (\overline{t \text{id}}, s, e_r, t_r)]}$	F-FCT

**Figure 11.** Functions are registered in the function map  $\Phi$  in IMP and VECIMP.

*Proof.* By induction on a derivation of  $\Phi, \Gamma \vdash a : v, \Gamma$ .

The proof is conducted in two parts. First, we show that  $\sigma'$  is also well-formed:

- E-VOKE sets the output scope to  $[T := \emptyset, A := \emptyset]$  which is by definition well-formed.
- E-DECL updates the scope to  $[T := T[\text{id} := \tau], A := A[\text{id} := a]]$ . As  $\sigma$  was well-formed beforehand (by the induction hypothesis) and a synced addition happens the output scope is well-formed.
- All other rules do not modify the scope and so the output scope is well-formed by the induction hypothesis.

Now to the statement/lvalue/value type property:

*Case (T-VAL):* This is a final configuration.

*Case (T-LVAL):* This is a final configuration.

*Case (T-CONST):* A constant is axiomatically a value by rule E-CONST.

*Case (T-VAR id in scope  $\sigma$ ):* Since  $\text{id} \in T_\sigma$  by prerequisite, a well-typed variable is an lvalue.

*Case (T-LPROJ/T-RPROJ  $\bar{e} := e.i$ ):* By the induction hypothesis  $e$  is either a value (meaning E-RPROJ applies), an lvalue (meaning E-LPROJ applies), or  $\bar{e}$  progresses via E-EVAL.

*Case (T-ADDR  $\bar{e} := \&e$ ):* By the induction hypothesis  $e$  is either an lvalue (meaning E-ADDR applies), or  $\bar{e}$  progresses via E-EVAL.

*Case (T-DEREF  $\bar{e} := *e$ ):* By the induction hypothesis  $e$  is either a value (meaning E-DEREF applies), an lvalue (meaning E-TOVAL applies) or  $\bar{e}$  progresses via E-EVAL.

*Case (T-INDEX  $\bar{e} := e_1[e_2]$ ):* By the induction hypothesis  $e_1$  can be a value. There are three sub cases for  $e_2$ . By the induction hypothesis  $e_2$  can be a value. If  $e_2$  is not of integer type then  $\bar{e}$  is ill-typed. Otherwise rule E-INDEX applies. In the case that  $e_2$  is an lvalue, rule E-TOVAL applies. If  $e_2$  is neither a value nor an lvalue,  $\bar{e}$  progresses via E-EVAL. Back to the cases for  $e_1$ . If  $e_1$  is an lvalue, rule E-TOVAL applies. If  $e_1$  is neither a value nor an lvalue,  $\bar{e}$  progresses via E-EVAL.

*Case (T-VOKE  $\bar{e} := \text{id}(\bar{e})$ ):* Arguments are evaluated from left to right (see E-EVAL).

*Subcase (Let  $e_i$  be the left most argument which is not a value):* By the induction hypothesis  $e_i$  can also be an lvalue (meaning E-TOVAL applies) or—if  $e_i$  is neither a value nor an lvalue— $\bar{e}$  progresses via E-EVAL.

*Subcase (There are no arguments at all or all arguments are values):* Via the function map, a lookup for the function to call can be done:  $\Phi(\text{id}) = (\overline{t \text{id}}, s, e_r, \tau_r)$ . Otherwise  $\bar{e}$  would be ill-typed. Now rule E-VOKE applies.

*Case (T-RETURN  $\bar{e} := \text{return } e$ ):* By the induction hypothesis  $e$  is either a value (meaning E-RETURN applies), an lvalue (meaning E-TOVAL applies) or  $\bar{e}$  progresses via E-EVAL.

*Case (T-SKIP ;):* This is a final configuration.

*Case (T-SEQ  $\bar{s} := s_1 s_2$ ):* By the induction hypothesis  $s_1$  is either " ; " (meaning rule E-SEQ applies) or  $\bar{s}$  progresses via E-EVAL.

*Case (T-SCOPE):* Axiomatically rule E-INSCOPE applies rendering  $\{s\}^\sigma$ . By the induction hypothesis  $s$  is either “;” (meaning rule E-OUTSCOPE applies) or the whole statement progresses via E-EVAL.

*Case (T-EXPRSTMNT  $\tilde{s} := e$ ):* By the induction hypothesis  $e$  is either a value (meaning E-EXPRSTMNT applies), an lvalue (meaning E-EVAL<sub>e</sub> applies), or  $\tilde{s}$  progresses via E-EVAL.

*Case (T-DECL  $t \text{ id}$ ):* Axiomatically rule E-DECL applies rendering “;”.

*Case (T-ASSIGN  $\tilde{s} := e_2 = e_1$ ):* By the induction hypothesis  $e_1$  can be a value. There are two sub cases for  $e_2$ . By the induction hypothesis  $e_2$  can be an lvalue, where the types of the right- and left-hand side match (meaning E-ASSIGN applies), or  $\tilde{s}$  progresses via E-EVAL. Back to the cases for  $e_1$ . If  $e_1$  is an lvalue, rule E-TOVAL applies. If  $e_1$  is neither a value nor an lvalue,  $\tilde{s}$  progresses via E-EVAL.

*Case (T-INIT  $\tilde{s} := t \text{ id} = e$ ):* By the induction hypothesis  $e$  is either a value (meaning E-INIT applies), an lvalue (meaning E-TOVAL applies) or  $\tilde{s}$  progresses via E-EVAL.

*Case (T-IF  $s := \text{if } (e) s_1 \text{ else } s_2$ ):* By the induction hypothesis  $e$  is either a value. This value is either a boolean T (meaning E-IFT applies), a boolean F (meaning E-IFF applies) or not a boolean at all (meaning  $\tilde{s}$  is ill-typed). Back to the cases for  $e$ . If  $e$  is an lvalue rule E-TOVAL triggers, otherwise  $e$  progresses via E-EVAL.

*Case (T-WHILE  $s := \text{while } (e) s$ ):* By the induction hypothesis  $e$  axiomatically progresses via E-WHILE.  $\square$

**Lemma 2** (Preservation of IMP). *Let  $\Phi$  be a well-formed function map,  $\sigma$  and  $\sigma'$  be well-formed scopes,  $a$  and  $a'$  be IMP programs. Then holds*

$$\Phi, \Gamma := T_\sigma \vdash a : v, \Gamma' \quad \wedge \quad \Phi, a, \sigma, M \rightarrow a', \sigma', M' \\ \Rightarrow \quad \Phi, \Gamma' \vdash a' : v', \Gamma''$$

with

$$v = v' = \epsilon \quad \vee \quad v = v' = {}^l\tau \\ \vee \quad (v = v' = {}_r\tau \quad \vee \quad v = {}_r\tau \quad \wedge \quad v' = {}^l\tau).$$

*Proof.* By induction on a derivation of  $\Phi, \Gamma \vdash a : v, \Gamma$ .

The proof is conducted in two parts. First, we show the property that  $\Gamma$  and  $T_\sigma$  stay in sync:

- The scoping rules are correct: Since the old scope  $\sigma$  (which is correct by the induction hypothesis) is stored in the special form “ $\{s\}^\sigma$ ” (by E-INSCOPE) and is later on restored (by E-OUTSCOPE) as demanded by T-SCOPE.
- Similarly, the old scope (which is correct by the induction hypothesis) during a function call is stored in the special form “ $\text{return}^\sigma e$ ” (by E-INVOKE) and later on restored (by E-RETURN) as demanded by T-INVOKE.
- Scoping in T-IF and T-WHILE is correct, as the corresponding evaluation rules (E-IFT, E-IFF and E-WHILE, respectively) reduce this issue by a program transformation to conventional scoping.
- Finally, rules T-DECL and E-DECL put the same identifier/type pair into their maps. By the induction hypothesis the scope was correct beforehand and is thus correct afterwards.
- All other rules do not change the type maps at all and do not need further consideration.

Now to the statement/lvalue/value type property:

*Case (T-VAL):* There does not exist a next evaluation step.

*Case (T-LVAL):* There does not exist a next evaluation step.

*Case (T-CONST):* There does not exist a next evaluation step.

*Case (T-VAR):* There does not exist a next evaluation step.

*Case (T-LPROJ/T-RPROJ  $e.i$ ):* Since  $e$  type checks to a value or an lvalue (by the premise of T-LPROJ or T-RPROJ, respectively)  $e$  can be a value. Applying E-RPROJ renders a value which type checks via T-VAL. The expression  $e$  can also be an lvalue. Applying E-LPROJ renders an lvalue which type checks via T-LVAL. Otherwise  $e$  still type checks to a value or an lvalue and a newly rendered expression type checks via T-LPROJ or T-RPROJ, respectively, by the induction hypothesis.

*Case (T-ADDR  $\&e$ ):* Since  $e$  type checks to an lvalue (by the premise of T-ADDR),  $e$  can be an lvalue. Applying E-ADDR renders a value which type checks via T-VAL. Otherwise  $e$  still type checks to an lvalue and a newly rendered expression type checks via T-ADDR by the induction hypothesis.

*Case (T-DEREF  $*e$ ):* Since  $e$  type checks to a pointer value or lvalue (by the premise of T-DEREF),  $e$  can be a pointer value. Applying E-DEREF renders an lvalue which type checks via T-LVAL. Otherwise  $e$  still type checks to a pointer value or lvalue and a newly rendered expression type checks via T-DEREF by the induction hypothesis.

*Case (T-INDEX  $e_1[e_2]$ ):* Since  $e_1$  type checks to a pointer value or lvalue and  $e_2$  type checks to an integer lvalue or value (by the premise of T-INDEX),  $e_1$  can be an index value and  $e_2$  can be an integer value. Applying E-INDEX renders an lvalue which type checks via T-LVAL. Otherwise  $e_1$  still type checks to a pointer value or lvalue and  $e_2$  still type checks to an integer lvalue or value and a newly rendered expression type checks via T-INDEX by the induction hypothesis.

*Case (T-INVOKE  $\tilde{e} := \text{id}(\tilde{e})$ ):* Since all arguments type check to value or lvalues, they all can be values (or there are no arguments at all). Applying E-INVOKE inlines the function (a  $\beta$ -reduction). As  $\Phi$  is well-defined by precondition, the new program also type checks. If one or more arguments are not values, these expressions still type check to values or lvalues and a newly rendered expression type checks via T-INVOKE by the induction hypothesis.

*Case (T-RETURN  $\tilde{s} := \text{return } e$ ):* Since  $e$  type checks to a value or an lvalue (by the premise of T-RETURN),  $e$  can be a value. Applying E-RETURN renders a value which type checks via T-VAL. Otherwise  $e$  still type checks to a value or an lvalue and a newly rendered statement type checks via T-RETURN by the induction hypothesis.

*Case (T-SKIP ;):* There does not exist a next evaluation step.

*Case (T-SEQ  $s_1 s_2$ ):* Since  $s_1$  type checks to  $\epsilon$  (by the premise of T-SEQ),  $s_1$  can be “;”. Applying E-SEQ renders  $s_2$  which type checks by the premise of T-SEQ. Otherwise  $s_1$  still type checks to  $\epsilon$  and a newly rendered statement type checks via T-SEQ by the induction hypothesis.

*Case (T-SCOPE):* Axiomatically rule E-INSCOPE applies rendering  $\{s\}^\sigma$ . Since  $s$  type checks to  $\epsilon$  (by the premise of T-SCOPE),  $s$  can be “;”. Applying E-OUTSCOPE renders “;” which type checks via T-SKIP. Otherwise  $s$  still type checks to  $\epsilon$  and a newly rendered statement type checks via T-SCOPE by the induction hypothesis.

*Case (T-EXPRSTMNT  $e$ ):* Since  $e$  type checks to a value or an lvalue (by the premise of T-EXPRSTMNT),  $e$  can be a value. Applying E-EXPRSTMNT renders “;” which type checks via T-SKIP. Otherwise  $e$  still type checks to a value or an lvalue and a newly rendered statement type checks via T-EXPRSTMNT by the induction hypothesis.

*Case (T-DECL  $t \text{ id}$ ):* Axiomatically rule E-DECL applies rendering “;” which type checks via T-SKIP.

*Case (T-ASSIGN  $e_2 = e_1$ ):* Since  $e_1$  type checks to a value or an lvalue and  $e_2$  type checks to an lvalue of the same type  $\tau$ ,  $e_1$  can be  $\mathcal{V}(m, \tau)$  and  $e_2$  can be  $\mathcal{L}(a, \tau)$  (for some byte sequence  $m$  and address  $a$ ). Applying rule E-ASSIGN renders “;” which type checks via T-SKIP. Otherwise  $e_1$  still type checks to a value or an lvalue and  $e_2$  type checks to an lvalue of the same type and a newly rendered statement type checks via T-ASSIGN by the induction hypothesis.

Case (T-INIT  $t \text{ id} = e$ ): Since  $e$  type checks to a value or an lvalue (by the premise of T-ASSIGN which is implied by the premise of T-INIT),  $e$  can be a value. Applying rule E-INIT renders a program consisting of a declaration and an assignment. This program type checks by the premise of T-INIT. Otherwise  $e$  still type checks to a value or an lvalue and a newly rendered statement type checks via T-INIT by the induction hypothesis.

Case (T-IF  $s := \text{if } (e) s_1 \text{ else } s_2$ ): Since  $e$  type checks to a boolean value or lvalue (by the premise of T-IF),  $e$  can be the boolean value T. Applying rule E-IFT renders  $\{ s_1 \}^\sigma$  which type checks via T-SCOPE. The expression  $e$  can also be the boolean value F. Applying rule E-IF renders  $\{ s_2 \}^\sigma$  which type checks via T-SCOPE. Otherwise  $e$  still type checks to a boolean value or lvalue and a newly rendered statement type checks via T-IF by the induction hypothesis.

Case (T-WHILE  $s := \text{while } (e) s$ ): Axiomatically rule E-WHILE applies rendering  $\text{if } (e) \{ s \text{ while } (e) s \}$ . This new program again type checks by the induction hypothesis and the premise of T-WHILE.  $\square$

*Proof.* Theorem 1 directly follows from Lemmas 1 and 2.  $\square$

## B. The Semantics of VecIMP

The semantics of VecIMP is very similar to the one of IMP, but the type system must keep track of the current vector length whereas evaluation rules have to maintain a predication mask as already outlined in Section 4.

The most intrusive addition to VecIMP is its polymorphism. In contrast to IMP, function calls are only type checked on demand. For this reason, no definition of a well-formed function map is given for VecIMP. Instead, the call-site is responsible to explicitly initiate type checking for the particular instance of the function (see T-INVOKe). In the case of recursion, it is important to introduce a termination condition for the type checker in order to prevent the type checker running infinitely. This is achieved by splitting  $\Gamma$  into the former type map—now  $T_\Gamma$ —and a function instance map  $F_\Gamma$ . The latter map serves two purposes: First, T-INVOKe memorizes with the help of  $F_\Gamma$  that a specific function instance is currently being type checked. Second, a specific instance maps to the inferred return type as the return type may be *unbound*. Initially, T-INVOKe sets this type to *unknown* (denoted by  $\epsilon$ ). If the return type is still unknown after examining the whole function, the vector length of the return expression is used to deduce the return type. If, however, a recursion occurs, T-INVOKeREC is triggered and this rule heuristically decides to use the current vector length  $\Lambda$  as vector length for the return type without type checking the specific instance again.

Furthermore, the evaluation contexts of rules E-EVAL and E-TOVAL, must be adjusted in order to reflect the changed syntax:

$$\begin{aligned} \mathcal{E}[\star] &::= \mathcal{L}[\star] \mid \& \star \mid \star. \dot{\text{i}} \mid \star = \mathcal{V}(m, \tau) \\ \mathcal{L}[\star] &::= \star \star \mid \star[e] \mid \mathcal{V}(m, \tau)[\star] \mid \text{id}(\overline{\mathcal{V}(m, \tau)}, \star, \bar{e}) \\ &\mid \star; \mid e = \star \mid t \text{ id} = \star; \\ &\mid \text{if } (\star) s_1; \text{ else } s_2; \mid \text{return}^\sigma \star; \end{aligned}$$

The sequence and scope statements are not contained in these evaluation contexts. Code regions where the mask is entirely set to *false* must be skipped (see Figure 12). Rules E-SEQT, E-SEQF, E-SCOPE and E-SCOPEF take care of this behavior.

### B.1 The Soundness Proof for VecIMP's Type System

**Lemma 3** (Progress of VecIMP). *Let  $\Phi$  be a function map,  $\sigma$  be a well-formed scope,  $I$  be a function instance map,  $a$  be a VecIMP*

*program and  $M$  be a memory state. Then holds*

$$\begin{aligned} \Phi, \Gamma := [T := T_\sigma, I], \Lambda := \|P_\sigma\| \vdash a : v, \Gamma' \\ \Rightarrow \begin{cases} a \text{ is a value} & \text{if } v = {}_r\tau \\ a \text{ is an lvalue} & \text{if } v = {}_l\tau \\ a \text{ is a value or an lvalue} & \text{if } v = {}_r\tau \\ a \text{ is ``;''} & \text{if } v = \epsilon \end{cases} \\ \vee \exists a', \sigma', M' : \Phi, a, \sigma, M \rightarrow a', \sigma', M'. \end{aligned}$$

*Proof.* Analogue to the proof of Lemma 1 with the following changes/additions:

Principally, vector lengths which are allowed by the static semantics must be also allowed by the dynamic semantics and vice versa. Since these checks (mostly done by  $\bar{v}$ ) are the same for typing and evaluation rules, this property is trivial.

Case (T-INVOKe/T-INVOKeREC): Analogue to IMP's cases, but either T-INVOKe applies—if the current function instance is not in the function instance map—or T-INVOKeREC applies otherwise. In both cases E-INVOKe triggers.

Case (T-SEQ  $s_1 s_2$ ): By the induction hypothesis  $s_1$  is either “;” (meaning rule E-SEQ applies) or  $s_1$  progresses via E-SEQT or E-SEQF.

Case (T-SCOPE): Axiomatically rule E-INSCOPE applies rendering  $\{ s \}^\sigma$ . By the induction hypothesis  $s$  is either “;” (meaning rule E-OUT-SCOPE applies) or  $s$  progresses via E-SCOPEF or E-SCOPE.

Case (T-DECL  $t \text{ id};$ ): Axiomatically rule E-DECLL applies rendering  $t^\wedge \text{ id};$ .

Case (T-DECLL  $t^\wedge \text{ id};$ ): Since  $\Lambda = \|P_\sigma\|$  rule E-DECLL applies.

Case (T-ASSIGN/T-BROADCAST  $\bar{s} := e_2 = e_1$ ): By the induction hypothesis  $e_1$  can be a value. There are three sub cases for  $e_2$ . By the induction hypothesis  $e_2$  can be an lvalue, where the types of the right- and left-hand side match (meaning E-ASSIGN applies), or, the left-hand side is vectorial whereas the right-hand side is scalar, while the base type  $v$  matches and  $\|P_\sigma\| \equiv \Lambda = 1 \vee \Lambda = \|v\|$  (meaning E-BROADCAST applies), or,  $\bar{s}$  progresses via E-EVAL. Back to the cases for  $e_1$ . If  $e_1$  is an lvalue, rule E-TOVAL applies. If  $e_1$  is neither a value nor an lvalue,  $\bar{s}$  progresses via E-EVAL.

Case (T-IF  $s := \text{if } (e) s_1 \text{ else } s_2$ ): By the induction hypothesis  $e$  is either a value. This value is a boolean vector of length  $l$  while  $\|P_\sigma\| \equiv \Lambda = 1 \vee \Lambda = l$  meaning E-IF applies or not a boolean vector at all (meaning  $\bar{s}$  is ill-typed). Back to the cases for  $e$ . If  $e$  is an lvalue rule E-TOVAL triggers, otherwise  $e$  progresses via E-EVAL.  $\square$

**Lemma 4** (Preservation of VecIMP). *Let  $\Phi$  be a function map,  $\sigma$  and  $\sigma'$  be well-formed scopes,  $a$  and  $a'$  be IMP programs. Then holds*

$$\begin{aligned} \Phi, \Gamma := [T := T_\sigma, I], \Lambda := \|P_\sigma\| \vdash a : v, \Gamma' := [T' := T_{\sigma'}, I'] \\ \wedge \Phi, a, \sigma, M \rightarrow a', \sigma', M' \\ \Rightarrow \Phi, \Gamma', \Lambda' := \|P_{\sigma'}\| \vdash a' : v, \Gamma'' \end{aligned}$$

with

$$\begin{aligned} v = v' = \epsilon \quad \vee \quad v = v' = {}_l\tau \\ \vee \quad (v = v' = {}_r\tau \quad \vee \quad v = {}_r\tau \wedge v' = {}_l\tau). \end{aligned}$$

*Proof.* Analogue to the proof of Lemma 2 with the following changes/additions:

$\Lambda$  and  $\|P_\sigma\|$  must stay in sync: This property is trivial as all evaluation rules corresponding to typing rules changing  $\Lambda$ , change  $\|P_\sigma\|$  in the same way and vice versa.

Case (T-INVOKe/T-INVOKeREC  $\bar{e} := \text{id}(\bar{e})$ ): Since all arguments type check to value or lvalues, they all can be values (or there are

$\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma$	$\Phi, a, \sigma, M \rightarrow a', \sigma', M'$
$\frac{\Phi, \Gamma_0, \Lambda \vdash s_1 : \epsilon, \Gamma_1 \quad \Phi, \Gamma_1, \Lambda \vdash s_2 : \epsilon, \Gamma_2}{\Phi, \Gamma_0, \Lambda \vdash s_1 s_2 : \epsilon, \Gamma_2} \quad \text{T-SEQ}$	$\frac{P_\sigma = \langle \bar{F}^+ \rangle}{\Phi, s_1 s_2, \sigma, M \rightarrow s_2, \sigma, M} \quad \text{E-SEQF}$
$\frac{\Phi, \Gamma, \Lambda \vdash s : \epsilon, \Gamma'}{\Phi, \Gamma, \Lambda \vdash \{s\} : \epsilon, \Gamma} \quad \text{T-SCOPE}$	$\frac{P_\sigma \neq \langle \bar{F}^+ \rangle \quad \Phi, s_1, \sigma, M \rightarrow s'_1, \sigma', M'}{\Phi, s_1 s_2, \sigma, M \rightarrow s'_1 s_2, \sigma', M'} \quad \text{E-SEQT}$
	$\frac{}{\Phi, ; s, \sigma, M \rightarrow s, \sigma, M} \quad \text{E-SEQ}$
	$\frac{}{\Phi, \{s\}, \sigma, M \rightarrow \{s\}^\sigma, \sigma, M} \quad \text{E-INSCOPE}$
	$\frac{P_\sigma = \langle \bar{F}^+ \rangle}{\Phi, \{s\}^{\sigma'}, \sigma, M \rightarrow ;, \sigma', M} \quad \text{E-SCOPEF}$
	$\frac{P_\sigma \neq \langle \bar{F}^+ \rangle \quad \Phi, s, \sigma, M \rightarrow s', \sigma', M'}{\Phi, \{s\}^{\hat{\sigma}}, \sigma, M \rightarrow \{s'\}^{\hat{\sigma}}, \sigma', M'} \quad \text{E-SCOPE T}$
	$\frac{}{\Phi, \{ ; \}^{\sigma'}, \sigma, M \rightarrow ;, \sigma', M} \quad \text{E-OUTSCOPE}$

**Figure 12.** Skipping evaluation if the current mask is entirely false

no arguments at all). Applying E-INVOKE inlines the function (a  $\beta$ -reduction). This new program type checks by the premise of T-INVOKE or by the induction hypothesis in the case of T-INVOKEREC. If one or more arguments are not values, these expressions still type check to values or lvalues and a newly rendered expression type checks via T-INVOKE by the induction hypothesis.

*Case (T-SEQ  $s_1 s_2$ ):* Since  $s_1$  type checks to  $\epsilon$  (by the premise of T-SEQ),  $s_1$  can be “;”. Applying E-SEQ renders  $s_2$  which type checks by the premise of T-SEQ. Otherwise  $s_1$  still type checks to  $\epsilon$  and a newly rendered statement by E-SEQT or E-SEQF, respectively, type checks via T-SEQ by the induction hypothesis.

*Case (T-SCOPE):* Axiomatically rule E-INSCOPE applies rendering  $\{s\}^\sigma$ . Since  $s$  type checks to  $\epsilon$  (by the premise of T-SCOPE),  $s$  can be “;”. Applying E-OUTSCOPE renders “;” which type checks via T-SKIP. Otherwise  $s$  still type checks to  $\epsilon$  and a newly rendered statement by E-SCOPE T or E-SCOPE F, respectively, type checks via T-SCOPE by the induction hypothesis.

*Case (T-DECL  $t \text{ id};$ ):* Axiomatically rule E-DECLL applies rendering  $t^\wedge \text{ id};$ . This new program again type checks via T-DECLL by premise of T-DECL.

*Case (T-DECLL  $t^\wedge \text{ id};$ ):* Since  $\Lambda = \|\|P_\sigma\|\|$  rule E-DECLL applies rendering “;” which type checks via T-SKIP.

*Case (T-ASSIGN/T-BROADCAST  $e_2 = e_1$ ):* Since  $e_1$  type checks to a value or an lvalue and  $e_2$  type checks to an lvalue of the same type  $\tau$ ,  $e_1$  can be  $\mathcal{V}(m, \tau)$  and  $e_2$  can be  $\mathcal{L}(a, \tau)$  (for some byte sequence  $m$  and address  $a$ ). Applying rule E-ASSIGN renders “;” which type checks via T-SKIP. The left-hand side  $e_1$  can also be  $\mathcal{V}(m, \nu)$  and  $e_2$  can be  $\mathcal{L}(a, \tau)$  (for some byte sequence  $m$  and address  $a$ ) while  $\text{vec}_{\Lambda = \|\|P_\sigma\|\|}(\tau, \|\nu\|) = \nu$ . Applying rule E-BROADCAST renders “;” which type checks via T-SKIP. Otherwise  $e_1$  still type checks to a value or an lvalue and  $e_2$  type checks to an lvalue of the same type or  $e_2$  is broadcastable to  $e_1$ 's type and a newly rendered statement type checks via T-ASSIGN or T-BROADCAST, respectively, by the induction hypothesis.  $\square$

*Proof.* Theorem 2 directly follows from Lemmas 3 and 4.  $\square$

## B.2 Vectorial Pointers

Indexing vectorial pointers (as described in Section 4.3) is not explicitly modeled in VECLMP. However, the involved address calculations can be implemented on top of the VECLMP calculus. For example, here is the pseudo code for implementing the gather from Figure 10a:

```
T gather(T scalar* block[N] p) {
  T block[N] result;
  for (int i = 0; i < N; ++i) {
    T scalar* s = get(p, i);
    set(result, *s, i);
  }
  return result;
}
```

This is a possible implementation for the read-access from Figure 10b:

```
T gather(T block[N]* block[N] p) {
  T block[N] result;
  for (int i = 0; i < N; ++i) {
    T block[n]* v = get(p, i);
    set(result, get(*v, i), i);
  }
  return result;
}
```

The write/scatter operations almost look identical. Also, versions which acknowledge the current mask are just slightly more complex (the for-each-active construct from Section 4.2.1 can be used for this task).

A practical language may want to overload the standard array access operator with this mechanism. This is useful in combination with polymorphism: In a scalar setting, scalar indexing has the usual semantics; however, in a vectorial setting the index and possibly the referenced data becomes vectorial and now the array access still has a sane semantics. Some of our ray tracer examples make use of this technique.

## B.3 Shallow vs. Deep Copies

Some of the presented examples make use of get/set operations in order to extract/insert vector elements. When pointers are involved some subtle memory problems arise:

```
struct Foo { int i; int* scalar p; }
Foo block[4] v = /*...*/;
Foo scalar s = get(v, 2); // where does s.p point to?
```

A *shallow copy* would simply copy  $v.p$  to  $s.p$  while a *deep copy* propagates into the pointed memory:

```
Foo scalar get(Foo f, int scalar i) {
  Foo scalar result;
  set(result.i, f.i, i);
  result.p = new int(get(f->p, i));
  return result;
}
```

This is akin to the problematic of copy constructors in C++ and a programming language can deal with this issue in a similar way:

By default *get/set* operations can be made available by the compiler which are implemented as shallow copies. Programmers are free to override this behaviour by implementing custom *get/set* operations. A broadcast suffers from the same problematic. `VECIMP` uses a shallow copy for the sake of simplicity (see `E-BROADCAST`).

	$\Phi, \Gamma \vdash e : \tau, \Gamma$		$\Phi, a, \sigma, M \rightarrow a', \sigma', M'$	
$\frac{}{\Phi, \Gamma \vdash \mathcal{V}(m, \tau) : \tau, \Gamma}$	T-VAL		$\frac{\Phi, a, \sigma, M \rightarrow a', \sigma', M'}{\mathcal{E}[a], \sigma, M \rightarrow \mathcal{E}[a'], \sigma', M'}$	E-EVAL
$\frac{}{\Phi, \Gamma \vdash \mathcal{L}(a, \tau) : \tau, \Gamma}$	T-LVAL		$\frac{}{\mathcal{L}[\mathcal{L}(a, \tau)], \sigma, M \rightarrow \mathcal{L}[\mathcal{V}(M_\sigma(a,  \tau ), \tau)], \sigma, M}$	E-TOVAL
$\frac{}{\Phi, \Gamma \vdash c_\tau : \tau, \Gamma}$	T-CONST		$\frac{V_\tau(c_\tau) = m}{\Phi, c_\tau, \sigma, M \rightarrow \mathcal{V}(m, \tau), \sigma, M}$	E-CONST
$\frac{\text{id} \in \text{dom}(\Gamma)}{\Phi, \Gamma \vdash \text{id} : \tau(\text{id}), \Gamma}$	T-VAR		$\frac{\text{id} \in \text{dom}(T_\sigma)}{\Phi, \text{id}, \sigma, M \rightarrow \mathcal{L}(A_\sigma(\text{id}), T_\sigma(\text{id})), \sigma, M}$	E-VAR
$\frac{\Phi, \Gamma \vdash e : \tau_i, \Gamma \quad \rho \in \{\bar{\tau}^+\} \quad 0 \leq i <  \bar{\tau}^+ }{\Phi, \Gamma \vdash e \cdot i : \tau_i, \Gamma}$	T-LPROJ		$\frac{\rho \in \{\bar{\tau}^+\} \quad 0 \leq i <  \bar{\tau}^+ }{\Phi, \mathcal{L}(a, \rho) \cdot i, \sigma, M \rightarrow \mathcal{L}(a + \text{offset}_\rho(i), \tau_i), \sigma, M}$	E-LPROJ
$\frac{\Phi, \Gamma \vdash e : \tau_i, \Gamma \quad \rho \in \{\bar{\tau}^+\} \quad 0 \leq i <  \bar{\tau}^+ }{\Phi, \Gamma \vdash e \cdot i : \tau_i, \Gamma}$	T-RPROJ		$\frac{\rho \in \{\bar{\tau}^+\} \quad 0 \leq i <  \bar{\tau}^+ }{\Phi, \mathcal{V}(m, \rho) \cdot i, \sigma, M \rightarrow \mathcal{V}(m(\text{offset}_\rho(i),  \tau_i ), \tau_i), \sigma, M}$	E-RPROJ
$\frac{\Phi, \Gamma \vdash e : \tau, \Gamma}{\Phi, \Gamma \vdash \&e : \tau^*, \Gamma}$	T-ADDR		$\frac{V_{\tau^*}(m) = a}{\Phi, \&\mathcal{L}(a, \tau), \sigma, M \rightarrow \mathcal{V}(m, \tau^*), \sigma, M}$	E-ADDR
$\frac{\Phi, \Gamma \vdash e : \tau^*, \Gamma}{\Phi, \Gamma \vdash *e : \tau, \Gamma}$	T-DEREF		$\frac{}{\Phi, *\mathcal{V}(m, \tau^*), \sigma, M \rightarrow \mathcal{L}(V_{\tau^*}(m), \tau), \sigma, M}$	E-DEREF
$\frac{\Phi, \Gamma \vdash e_1 : \tau_i^*, \Gamma \quad \Phi, \Gamma \vdash e_2 : \tau_i^*, \Gamma \quad \iota \in \mathcal{A}_{\text{int}}}{\Phi, \Gamma \vdash e_1 [e_2] : \tau_i, \Gamma}$	T-INDEX		$\frac{\iota \in \mathcal{A}_{\text{int}}}{\Phi, \mathcal{V}(m_1, \tau^*)[\mathcal{V}(m_2, \iota)], \sigma, M \rightarrow \mathcal{L}(V_{\tau^*}(m_1) +  \tau V_\tau(m_2), \tau), \sigma, M}$	E-INDEX
$\frac{\forall 0 \leq i <  \bar{\tau}  : \Phi, \Gamma \vdash e_i : \tau_i, \Gamma}{\Phi(\text{id}) = (t \text{id}, s, e_r, \tau_r) \quad \tau_r = \tau_r}$	T-INVOKE		$\frac{\Phi(\text{id}) = (t \text{id}, s, e_r, \tau_r)}{\Phi, \text{id}(\mathcal{V}(m, \tau)), \sigma, M \rightarrow \left[ \begin{array}{l} T := \emptyset, \\ A := \emptyset \end{array} \right], M}$	E-INVOKE
$\frac{\Phi, \Gamma \vdash e : \tau, \Gamma}{\Phi, \Gamma \vdash \text{return } e; : \tau, \Gamma}$	T-RETURN		$\frac{}{\Phi, \text{return}^{\sigma'} \mathcal{V}(m, \tau);, \sigma, M \rightarrow \mathcal{V}(m, \tau), \sigma, M}$	E-RETURN
$\frac{}{\Phi, \Gamma \vdash ; : \epsilon, \Gamma}$	T-SKIP			
$\frac{\Phi, \Gamma_0 \vdash s_1 : \epsilon, \Gamma_1 \quad \Phi, \Gamma_1 \vdash s_2 : \epsilon, \Gamma_2}{\Phi, \Gamma_0 \vdash s_1 s_2 : \epsilon, \Gamma_2}$	T-SEQ		$\frac{}{\Phi; s, \sigma, M \rightarrow s, \sigma, M}$	E-SEQ
$\frac{\Phi, \Gamma \vdash s : \epsilon, \Gamma'}{\Phi, \Gamma \vdash \{s\} : \epsilon, \Gamma}$	T-SCOPE		$\frac{}{\Phi, \{s\}, \sigma, M \rightarrow \{s\}^{\sigma}, \sigma, M}$	E-INSCOPE
			$\frac{}{\Phi, \{;\}^{\sigma'}, \sigma, M \rightarrow ;, \sigma', M}$	E-OUTSCOPE
$\frac{\Phi, \Gamma \vdash e : \tau, \Gamma}{\Phi, \Gamma \vdash e; : \epsilon, \Gamma}$	T-EXPRSTMT		$\frac{}{\Phi, \mathcal{V}(m, \tau);, \sigma, M \rightarrow ;, \sigma, M}$	E-EXPRSTMT
$\frac{\tau = t \quad \text{id} \notin \text{dom}(\Gamma)}{\Phi, \Gamma \vdash t \text{id}; : \epsilon, \Gamma[\text{id} := \tau]}$	T-DECL		$\frac{\tau = t \quad \text{alloc}(\sigma, \tau) = a}{\Phi, t \text{id};, \sigma, M \rightarrow ;, \sigma \left[ \begin{array}{l} T := T[\text{id} := \tau], \\ A := A[\text{id} := a] \end{array} \right], M}$	E-DECL
$\frac{\Phi, \Gamma \vdash e_1 : \tau, \Gamma \quad \Phi, \Gamma \vdash e_2 : \tau, \Gamma}{\Phi, \Gamma \vdash e_2 = e_1; : \epsilon, \Gamma}$	T-ASSIGN		$\frac{}{\Phi, \mathcal{L}(a, \tau) = \mathcal{V}(m, \tau);, \sigma, M \rightarrow ;, \sigma, M := M[\langle a,  \tau  \rangle := m]}$	E-ASSIGN
$\frac{\Phi, \Gamma \vdash t \text{id}; \text{id} = e; : \epsilon, \Gamma'}{\Phi, \Gamma \vdash t \text{id} = e; : \epsilon, \Gamma'}$	T-INIT		$\frac{}{\Phi, t \text{id} = \mathcal{V}(m, \tau);, \sigma, M \rightarrow t \text{id}; \text{id} = \mathcal{V}(m, \tau);, \sigma, M}$	E-INIT
$\frac{\Phi, \Gamma \vdash e : \text{bool}, \Gamma \quad \Phi, \Gamma \vdash s_1 : \epsilon, \Gamma_1 \quad \Phi, \Gamma \vdash s_2 : \epsilon, \Gamma_2}{\Phi, \Gamma \vdash \text{if}(e) s_1 \text{ else } s_2 : \epsilon, \Gamma}$	T-IF		$\frac{}{\Phi, \text{if}(\mathcal{V}(T, \text{bool})) s_1 \text{ else } s_2, \sigma, M \rightarrow \{s_1\}^{\sigma}, \sigma, M}$	E-IFT
			$\frac{}{\Phi, \text{if}(\mathcal{V}(F, \text{bool})) s_1 \text{ else } s_2, \sigma, M \rightarrow \{s_2\}^{\sigma}, \sigma, M}$	E-IFF
$\frac{\Phi, \Gamma \vdash e : \text{bool}, \Gamma \quad \Phi, \Gamma \vdash s : \epsilon, \Gamma'}{\Phi, \Gamma \vdash \text{while}(e) s : \epsilon, \Gamma}$	T-WHILE		$\frac{}{\Phi, \text{while}(e) s, \sigma, M \rightarrow \text{if}(e) \{s \text{ while}(e) s\}, \sigma, M}$	E-WHILE

Figure 13. Static (left) and dynamic (right) semantics of IMP

	$\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma$		$\Phi, a, \sigma, M \rightarrow a', \sigma', M'$
	$\frac{\vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash \mathcal{V}(m, \tau) : \tau, \Gamma}$	T-VAL	$\frac{\Phi, a, \sigma, M \rightarrow a', \sigma', M'}{\mathcal{E}[a], \sigma, M \rightarrow \mathcal{E}[a'], \sigma', M'}$ E-EVAL
	$\frac{\vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash \mathcal{L}(a, \tau) : \tau, \Gamma}$	T-LVAL	$\frac{\vec{\tau}(\ P_\sigma\ , \tau)}{\mathcal{L}[\mathcal{L}(a, \tau)], \sigma, M \rightarrow \mathcal{L}[\mathcal{V}(M_\sigma(a,  \tau ), \tau)], \sigma, M}$ E-TOVAL
	$\frac{\vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash c_\tau : \tau, \Gamma}$	T-CONST	$\frac{V_\tau(c_\tau) = m \quad \vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, c_\tau, \sigma, M \rightarrow \mathcal{V}(m, \tau), \sigma, M}$ E-CONST
	$\frac{\text{id} \in \text{dom}(\Gamma) \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash \text{id} : \tau, \Gamma}$	T-VAR	$\frac{\text{id} \in \text{dom}(T_\sigma) \quad \vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, \text{id}, \sigma, M \rightarrow \mathcal{L}(A_\sigma(\text{id}), T_\sigma(\text{id})), \sigma, M}$ E-VAR
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \rho \in \{\vec{\tau}^+\} \quad 0 \leq i <  \vec{\tau}^+  \quad \vec{\tau}(\Lambda, \tau_i)}{\Phi, \Gamma, \Lambda \vdash e.i : \tau_i, \Gamma}$	T-LPROJ	$\frac{\rho \in \{\vec{\tau}^+\} \quad 0 \leq i <  \vec{\tau}^+  \quad \vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, \mathcal{L}(a, \rho).i, \sigma, M \rightarrow \mathcal{L}(a + \text{offset}_\rho(i), \tau_i), \sigma, M}$ E-LPROJ
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \rho \in \{\vec{\tau}^+\} \quad 0 \leq i <  \vec{\tau}^+  \quad \vec{\tau}(\Lambda, \tau_i)}{\Phi, \Gamma, \Lambda \vdash e.i : \tau_i, \Gamma}$	T-RPROJ	$\frac{\rho \in \{\vec{\tau}^+\} \quad 0 \leq i <  \vec{\tau}^+  \quad \vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, \mathcal{V}(m, \rho).i, \sigma, M \rightarrow \mathcal{V}(m(\text{offset}_\rho(i),  \tau_i ), \tau_i), \sigma, M}$ E-RPROJ
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash \&e : \tau^*, \Gamma}$	T-ADDR	$\frac{V_{\tau^*}(m) = a}{\Phi, \&\mathcal{L}(a, \tau), \sigma, M \rightarrow \mathcal{V}(m, \tau^*), \sigma, M}$ E-ADDR
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash *e : \tau, \Gamma}$	T-DEREF	$\frac{}{\Phi, * \mathcal{V}(m, \tau^*), \sigma, M \rightarrow \mathcal{L}(V_{\tau^*}(m), \tau), \sigma, M}$ E-DEREF
	$\frac{\Phi, \Gamma, \Lambda \vdash e_1 : \tau, \Gamma \quad \vec{\tau}(\Lambda, \tau) \quad \Phi, \Gamma, \Lambda \vdash e_2 : \tau, \Gamma \quad \iota \in \mathcal{A}_{\text{int}}}{\Phi, \Gamma, \Lambda \vdash e_1[e_2] : \tau, \Gamma}$	T-INDEX	$\frac{\iota \in \mathcal{A}_{\text{int}} \quad \vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, \mathcal{V}(m_1, \tau^*)[\mathcal{V}(m_2, \iota)], \sigma, M \rightarrow \mathcal{L}(V_{\tau^*}(m_1) +  \tau V_\tau(m_2), \tau), \sigma, M}$ E-INDEX
	$\frac{\forall 0 \leq i <  \vec{\tau}  : \Phi, \Gamma, \Lambda \vdash e_i : \tau_i, \Gamma \quad \vec{\tau}(\Lambda, \tau_i) \quad \Phi(\text{id}) = (\vec{\tau} \text{ id}, s, e_r, \tau_r) \quad f := (\text{id}, \Lambda, \text{vec}_\Lambda(t, \ \tau\ )) \notin \text{dom}(F_\Gamma) \quad \Phi, \Gamma[F := F[f := \epsilon]], \Lambda \vdash \text{id} = \mathcal{V}(m, \tau); s; \text{return } e_r; : \tau, \Gamma, \Gamma'$ $\tau_r := \begin{cases} \text{vec}_\Lambda(t_r, \ \tau_r\ ) & \text{if } F_\Gamma(f) = \epsilon \\ F_{\Gamma'}(f) & \text{otherwise} \end{cases}$	T-INVOKE	$\frac{\Phi(\text{id}) = (\vec{\tau} \text{ id}, s, e_r, \tau_r)}{\Phi, \text{id}(\mathcal{V}(m, \tau)), \sigma, M \rightarrow \left[ \begin{array}{l} T := \emptyset, \\ A := \emptyset \end{array} \right], M}$ E-INVOKE
	$\frac{\forall 0 \leq i <  \vec{\tau}  : \Phi, \Gamma, \Lambda \vdash e_i : \tau_i, \Gamma \quad \vec{\tau}(\Lambda, \tau_i) \quad \Phi(\text{id}) = (\vec{\tau} \text{ id}, s, e_r, \tau_r) \quad f := (\text{id}, \Lambda, \text{vec}_\Lambda(t, \ \tau\ )) \in \text{dom}(F_\Gamma) \quad \tau_r := \begin{cases} \text{vec}_\Lambda(t_r, \Lambda) & \text{if } F_\Gamma(f) = \epsilon \\ F_\Gamma(f) & \text{otherwise} \end{cases}$	T-INVOKEREC	$\frac{}{\Phi, \Gamma, \Lambda \vdash \text{id}(\vec{\tau}) : \tau, \Gamma[F := F[f := \tau_r]]}$
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash \text{return } e; : \tau, \Gamma}$	T-RETURN	$\frac{}{\Phi, \text{return}^{\sigma'} \mathcal{V}(m, \tau); \sigma, M \rightarrow \mathcal{V}(m, \tau), \sigma', M}$ E-RETURN
	$\frac{}{\Phi, \Gamma, \Lambda \vdash ; : \epsilon, \Gamma}$	T-SKIP	
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \Phi, \Gamma, \Lambda \vdash e; : \epsilon, \Gamma}{\Phi, \Gamma, \Lambda \vdash e; : \epsilon, \Gamma}$	T-EXPRSTMT	$\frac{}{\Phi, \mathcal{V}(m, \tau); \sigma, M \rightarrow ; \sigma, M}$ E-EXPRSTMT
	$\frac{\Phi, \Gamma, \Lambda \vdash t^\Lambda \text{id}; : \epsilon, \Gamma' \quad \Phi, \Gamma, \Lambda \vdash t \text{id}; : \epsilon, \Gamma'}{\Phi, \Gamma, \Lambda \vdash t^\Lambda \text{id}; : \epsilon, \Gamma'}$	T-DECL	$\frac{}{\Phi, t \text{id}; \sigma, M \rightarrow t^{\ P_\sigma\ } \text{id}; \sigma, M}$ E-DECL
	$\frac{\tau = \text{vec}_\Lambda(t, l) \quad \tau \neq \perp \quad \text{id} \notin \text{dom}(T_\Gamma) \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash t^\perp \text{id}; : \epsilon, \Gamma[T := T[\text{id} := \tau]]}$	T-DECLL	$\frac{\tau = \text{vec}_\Lambda(t, l) \quad \text{alloc}(\sigma, \tau) = a \quad \vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, t^\perp \text{id}; \sigma, M \rightarrow ; \sigma \left[ \begin{array}{l} T := T[\text{id} := \tau], \\ A := A[\text{id} := a] \end{array} \right], M}$ E-DECLL
	$\frac{\Phi, \Gamma, \Lambda \vdash e_1 : \tau, \Gamma \quad \Phi, \Gamma, \Lambda \vdash e_2 : \tau, \Gamma \quad \vec{\tau}(\Lambda, \tau)}{\Phi, \Gamma, \Lambda \vdash e_2 = e_1; : \epsilon, \Gamma}$	T-ASSIGN	$\frac{\vec{\tau}(\ P_\sigma\ , \tau)}{\Phi, \mathcal{L}(a, \tau) = \mathcal{V}(m, \tau); \sigma, M \rightarrow ; \sigma, M := M[\langle P_\sigma, a,  \tau  \rangle := m]}$ E-ASSIGN
	$\frac{\Phi, \Gamma, \Lambda \vdash e_1 : \tau, \Gamma \quad \Phi, \Gamma, \Lambda \vdash e_2 : \tau, \Gamma \quad \Lambda = 1 \vee \Lambda = \ \nu\  \quad \text{vec}_{\ P_\sigma\ }(\tau, \ \nu\ ) = \nu}{\Phi, \Gamma, \Lambda \vdash e_2 = e_1; : \epsilon, \Gamma}$	T-BROADCAST	$\frac{\ P_\sigma\  = 1 \vee \ P_\sigma\  = \ \nu\  \quad \text{vec}_{\ P_\sigma\ }(\tau, \ \nu\ ) = \nu}{\Phi, \mathcal{L}(a, \nu) = \mathcal{V}(m, \tau); \sigma, M \rightarrow ; \sigma, M := M[\langle P_\sigma, a,  \nu  \rangle := \langle \vec{m}^+ \rangle]}$ E-BROADCAST
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \Phi, \Gamma, \Lambda \vdash t^{\ \tau\ } \text{id}; \text{id} = e; : \epsilon, \Gamma'}{\Phi, \Gamma, \Lambda \vdash t^{\ \tau\ } \text{id} = e; : \epsilon, \Gamma'}$	T-INIT	$\frac{}{\Phi, t \text{id} = \mathcal{V}(m, \tau); \sigma, M \rightarrow t^{\ \tau\ } \text{id}; \text{id} = \mathcal{V}(m, \tau); \sigma, M}$ E-INIT
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \Lambda = 1 \vee \Lambda = l \quad \Phi, \Gamma, \Lambda := l \vdash s_1 : \epsilon, \Gamma_1 \quad \Phi, \Gamma, \Lambda := l \vdash s_2 : \epsilon, \Gamma_2}{\Phi, \Gamma, \Lambda \vdash \text{if } (e) s_1 \text{ else } s_2 : \epsilon, \Gamma}$	T-IF	$\frac{\Lambda = 1 \vee \Lambda = l}{\Phi, \text{if } (\mathcal{V}(m, \langle \text{bool} \mid l \rangle)) s_1 \text{ else } s_2, \sigma, M \rightarrow \{s_1\}^{\sigma'} \text{ if } (\mathcal{V}(\text{not } m, \langle \text{bool} \mid l \rangle)) s_2 \text{ else } ; \sigma[P := P \text{ and } m], M}$ E-IF
	$\frac{\Phi, \Gamma, \Lambda \vdash e : \tau, \Gamma \quad \Lambda = 1 \vee \Lambda = l \quad \Phi, \Gamma, \Lambda := l \vdash s : \epsilon, \Gamma'}{\Phi, \Gamma, \Lambda \vdash \text{while } (e) s : \epsilon, \Gamma}$	T-WHILE	$\frac{}{\Phi, \text{while } (e) s, \sigma, M \rightarrow \text{if } (e) \{s \text{ while } (e) s\}, \sigma, M}$ E-WHILE

Figure 14. Static (left) and dynamic (right) semantics of VECIMP