

HALIVER: Deductive Verification and Scheduling Languages Join Forces

Lars B. van den Haak¹[0000-0002-0330-5016], Anton Wijs¹[0000-0002-2071-9624],
Marieke Huisman²[0000-0003-4467-072X], and Mark van den
Brand¹[0000-0003-3529-6182]

¹ Eindhoven University of Technology, The Netherlands

{l.b.v.d.haak, a.j.wijs, m.g.j.v.d.brand}@tue.nl

² University of Twente, The Netherlands

m.huisman@utwente.nl

Abstract. The HALIVER tool integrates deductive verification into the popular scheduling language HALIDE, used for image processing pipelines and array computations. HALIVER uses VERCORS, a separation logic-based verifier, to verify the correctness of (1) the HALIDE algorithms and (2) the optimised parallel code produced by HALIDE when an optimisation schedule is applied to an algorithm. This allows proving complex, optimised code correct while reducing the effort to provide the required verification annotations. For both approaches, the same specification is used. We evaluated the tool on several optimised programs generated from characteristic HALIDE algorithms, using all but one of the essential scheduling directives available in HALIDE. Without annotation effort, HALIVER proves memory safety in almost all programs. With annotations HALIVER, additionally, proves functional correctness properties. We show that the approach is viable and reduces the manual annotation effort by an order of magnitude.

Keywords: Program correctness · Deductive verification · Scheduling language.

1 Introduction

To meet the continuously growing demands on software performance, parallelism is increasingly often needed [15]. However, introducing parallelism tends to increase the risk of introducing errors, as the interactions between parallel computations can be hard to predict. Moreover, a plethora of optimisation techniques exists [12], so identifying when an optimisation can be applied safely, without breaking correctness, can be very challenging. Also, applying optimisations tends to make a program more complex, making it harder to reason about.

To address this, on the one hand, various domain-specific languages (DSLs) have been proposed that separate the *algorithm* (*what* it does) from the parallelisation *schedule* (*how* it does it). These are called *scheduling languages* [3, 7, 8, 10, 24, 25, 30]. Given an algorithm and a schedule, a compiler generates an




Fig. 1. High level overview of our approach.

optimised parallel program. This approach crucially depends on the schedule not introducing any errors in the functionality, which is not always obvious.

On the other hand, *deductive program verification* [11] has been successfully applied to verify the functionality of parallel programs [5]. This requires that the intended functionality is formalised as a contract, for instance using *permission-based separation logic* [1,6]. A major hurdle, preventing this technique from being adopted at a large scale, is that if a program becomes more complicated, the required annotations rapidly grow in size and complexity [27, 28].

In this paper, we combine the best of both worlds. We propose the HALIVER tool, which focusses on HALIDE [24,25], a scheduling language for portable image computations and array processing. It has been widely adopted in industry, for instance to produce parts of Adobe Photoshop and to implement the YouTube video-ingestion pipeline. For verification, we use the VERCORS program verifier [5]. In this paper we define two verification approaches (1) *front-end* and (2) *back-end*, as seen in Figure 1. Our approaches verify that the program adheres to the *same* functional specification. This specification is detailed by annotating the algorithmic part of a HALIDE program, thereby keeping the annotations focussed on the functionality, and therefore relatively straightforward. With the *front-end* verification approach we verify the correctness of the algorithmic part of a HALIDE program. HALIVER transforms the algorithm and the annotations to an annotated VERCORS program. With *back-end* verification approach we verify the C code that the HALIDE compiler generates, given a HALIDE algorithm and a schedule. HALIVER transforms the given annotations to match the generated code. Furthermore, where possible, HALIVER generates annotations, such as permission specifications, to relieve the user from having to manually write these. This contributes to making the annotation process straightforward.

In this way, HALIVER allows the user to succinctly specify the intended functionality of optimised, parallel code, and it checks that the resulting program indeed has the desired functionality. A major advantage of our approach is that it is flexible to use in a setting where multiple compiler passes are made. Also, it can be easily extended if a new compiler pass or schedule optimisation is added. An alternative would be to prove correctness of the compiler, but this would require a large amount of initial work and additionally for each change to the compiler.

Concretely, this paper provides the following contributions:

- An annotation language to describe the functionality of HALIDE algorithms, which is integrated into the HALIDE algorithm language;

Listing 1. HALIDE blur example with annotations added to verify the code.

```

1  requires inp.x.min == blur_y.x.min  $\wedge$  inp.x.max == blur_y.x.max+2  $\wedge$  inp.y.min ==
   blur_y.y.min  $\wedge$  inp.y.max == blur_y.y.max+2;
2  ensures  $\forall$  x, y . blur_y.x.min  $\leq$  x  $<$  blur_y.x.max  $\wedge$  blur_y.y.min  $\leq$  y  $<$  blur_y.y.max  $\Rightarrow$ 
   blur_y(x,y) == ((inp(x,y)+inp(x+1,y)+inp(x+2,y))/3 + (inp(x,y+1) +inp(x+1,y+1)+
   inp(x+2,y+1))/3 + (inp(x,y+2)+inp(x+1,y+2) + inp(x+2,y+2))/3)/3);
3  void blur(Buffer<2,int> inp, Func &blur_y){
4  Func blur_x; Var x, y;
5  blur_x(x,y) = (inp(x,y) + inp(x+1,y) + inp(x+2,y))/3
6  blur_x.ensures(blur_x(x,y) == (inp(x,y) + inp(x+1,y) + inp(x+2,y))/3);
7  blur_y(x,y) = (blur_x(x,y) + blur_x(x,y+1) + blur_x(x,y+2))/3;
8  blur_y.ensures(blur_y(x,y) == ((inp(x,y)+inp(x+1,y)+inp(x+2,y))/3 + (inp(x,y+1)+inp
   (x+1,y+1) + inp(x+2,y+1))/3 + (inp(x,y+2)+inp(x+1,y+2)+inp(x+2,y+2))/3)/3);}

```

- Tool support for the front-end verification approach of HALIDE algorithms;
- Tool support for the back-end verification approach, which can verify programs generated by the HALIDE compiler from an algorithm and a schedule;
- Evaluation of the HALIVER tool on HALIDE examples using all but one of the essential scheduling directives available in various combinations. We evaluated the tool on 23 different optimised programs, generated from eight characteristic HALIDE algorithms, to prove memory safety with no annotation effort. For 21 cases, HALIVER proves safety, for the remaining two cases we discuss the limitations. For 20 programs, based on five algorithms, we also add annotations for functional correctness properties. For 19 of these programs HALIVER proves correctness, for the remaining one we run into a similar limitation.

The remainder of this paper is organised as follows. Section 2 gives brief background information on HALIDE and VERCORS. Section 3 introduces HALIDE annotations, and describes how HALIVER supports the verification of an algorithm and an optimised program. The approach is illustrated on characteristic examples. Section 4 evaluates the HALIVER tool, and Sections 5 and 6 address related work, conclusions and future work.

2 Background

HALIDE. HALIDE is a DSL embedded in C++, targeting image processing pipelines and array computations [24, 25].³ HALIDE separates the *algorithm*, defining what you want to calculate, from the *schedule*, defining how the calculation should be performed. Typically, when optimising code for a specific architecture, the code becomes much more complex and loses portability. By separating the schedule, the code expressing the functionality is not altered.

Listing 1 presents the HALIDE algorithm for a box filter, or blur function. The reader can ignore the **requires** and **ensures** annotations for now. Images are represented as pure (side-effect free) functions that point-wise map coordinates to values. A blur function defines how every pixel, referred to by its two-

³ A HALIDE tutorial can be found here: <https://halide-lang.org/tutorials/>.

Listing 2. A reduction to count the positive numbers of each row in matrix `inp`.

```

1 void cnt(Buffer<2,int> inp, Func count) {
2   Var x; RDom r(0,10);
3   count(x) = 0;
4   count.ensures(count(x) == 0);
5   count(x) = select(inp(x, r) > 0, count(x)+1, count(x))
6   count.invariant(0 ≤ count(x) ≤ r);
7   count.ensures(0 ≤ count(x) ≤ 10);}

```

dimensional coordinates, should be updated. In the example, the coordinates are represented by the variables `x` and `y`. HALIDE uses a functional style, allowing algorithms to be compact and loop-free. HALIDE functions are denoted by the keyword `Func`. In the example, the input image is stored in a two-dimensional integer buffer `inp`, and the output is given by defining the function `blur_y`, a reference to which is a parameter of `blur`. A *pipeline* of function calls is defined: the function `blur_x` is applied on the input image (line 5). The output of that function is used to compute the final image with the function `blur_y` (line 7). With `inp.x.min` and `inp.x.max` we refer to the minimum and maximum value of the dimension `inp.x`, respectively.

A function may involve *update definitions*, which (partially) update the value of a function. A *reduction domain* is a way to apply an update a finite number of times and is typically used to express sums or histograms in HALIDE. A function is called a *reduction* when such a domain is used, and an initialisation and an update definition are given. Listing 2 presents a reduction example. For now, ignore the `ensures` and `invariant` lines. The reduction domain (`RDom`) `r` ranges from 0 to 9, i.e. it consists of 10 values. The initial value of the `count` function is defined at line 3, and line 5 is executed once for every value in `r`. The statement `select(a, b, c)` returns `b` if `a` evaluates to `true`, `c` otherwise. For a given matrix of integers `inp`, `cnt` counts the number of non-zeros at the first ten positions of each row in `inp`.

A HALIDE *schedule* is given in Listing 5 and further explained in Section 3.3.

VERCORS. VERCORS⁴ [5] is a deductive verifier to verify the functional correctness of, possibly concurrent, software. Its specification language uses permission-based separation logic [6], a combination of first-order logic and read/write permissions. The latter are used for concurrency-related verification, to express which data can be accessed by a thread at which moment. Programs written in a number of languages, such as JAVA and C, can be verified. VERCORS also has its own language, PVL. VERCORS's verification engine relies on VIPER [18], which applies symbolic execution to analyse programs with persistent mutable state.

Intended functional behaviour can be specified by means of pre- and post-conditions, indicated by the keywords `requires` and `ensures`, respectively. The statement `context P` is an abbreviation for `requires P`; `ensures P`. Loop in-

⁴ An online tutorial can be found at <https://vercors.ewi.utwente.nl/wiki/>.

variants and assertions can be added to the code to help VERCORS in proving the pre- and postconditions. We refer to the pre- and postconditions, loop invariants and assertions together as the *annotations* of a code fragment. A permission `Perm(x, f)` gives permission to memory location `x`, where `f` is a fractional, with `1\1` indicating a write and anything between `0\1` and `1\1` a read. For a statement `s`, we have the Hoare triple $\{P\}s\{Q\}$. This indicates that if P holds in the *pre-state* then after `s`, Q holds in the *post-state*. A *pure* function is without side-effects, thus can be used in annotations. It has the keyword `pure` in the header, and its body is a single expression. Annotations and pure function definitions in C files are given in special comments, like `//@` or `/*@. . .@*/` for multi-line comments. (See Listing 6 for examples.)

VERCORS can prove termination of recursive functions. Whenever the clause `decreases r` is added to a function contract, VERCORS will try to prove that the function terminates, by showing that all recursive calls will strictly decrease the value of `r` while `r` has a lower bound.

3 Verification of Scheduling Languages with HALIVER

HALIVER works directly on a HALIDE program and its intermediate representations, adding and transforming annotations where necessary. The tool is embedded in the HALIDE compiler. From a user’s point of view, the general approach is as follows, using the front-end and back-end approach as in Figure 1.

1. **Write a HALIDE algorithm and add annotations.** Annotations are the functional specification of the HALIDE algorithm. Since a user can write an incorrect HALIDE algorithm, its correctness is ideally checked against a user-supplied specification.
2. **The front-end approach produces a PVL encoding.** This encoding contains the algorithm and the specified annotations.
3. **VERCORS verifies the encoding.** If verification succeeds, we know that the front-end algorithm conforms to the functional specification. Otherwise, the verification fails; VERCORS produces a counterexample and we return to step 1.
4. **Write a HALIDE schedule.**
5. **The back-end approach produces an annotated C file.** The tool automatically generates permission annotations. These allow us to prove data-race freedom and the absence of out-of-bound errors. The tool transforms the annotations and generates additional annotations to match the scheduled back-end code. This is highly non-trivial, as each for-loop requires precise annotations to guide VERCORS in the verification. However, it is ensured that the same property is verified.
6. **VERCORS verifies the back-end C file.** If the verification fails, the lines of C code that caused the failure are given, which can be traced back to the HALIDE algorithm. The cause of a verification failure may be that
 - The HALIDE compiler produced incorrect code w.r.t. the specifications.
 - More *auxiliary* annotations from step 1 are needed to guide VERCORS.

- A limitation has been encountered of the tools HALIVER relies on, e.g., VERCORS or the underlying SMT solver.

In the remainder of this section we explain how to write annotations, and address front-end and back-end verification approaches. We also discuss the soundness and current limitations of the technique.

3.1 HALIDE annotations

HALIVER makes it possible to add annotations when writing a HALIDE algorithm. Intuitively, these annotations are added as a Hoare triple. We consider three types of annotation: *pipeline*, *intermediate* and *reduction invariant* annotations.

In Listing 1 annotations have been added. The lines 1–2 are *pipeline annotations*: they specify the pre- and postconditions of the whole function and can only contain references to input buffers or output functions. Note that the results are stored directly in the `blur_y` function. Line 1 specifies how the input and output bounds should be related. Line 2 indicates what the output values are. One can add *intermediate annotations* after any (update) function call to specify state predicates for particular locations in the pipeline. Examples are the `blur_x.ensures` and `blur_y.ensures` state predicates of Listing 1 (lines 6 and 8).

HALIDE functions map coordinates to values pointwise. To achieve a one-to-one relationship between function and annotations, the intermediate annotations for a function should also specify how coordinates relate to values pointwise. However, input buffers can be used freely with any point. For example, `blur_x.ensures(blur_x(x,y) ≥ inp(x+1,y))` is valid, but `blur_x.ensures(blur_x(x+1,y) ≥ 0)` is not, because the latter refers to `blur_x(x+1,y)` as opposed to `blur_x(x,y)`. HALIVER requires this because each point of the function may be computed in parallel in the back-end, so it must be possible to reason about the points individually.

For ease of annotation, HALIVER automatically generates a pipeline postcondition. This postcondition is derived from the intermediate annotation of the last pipeline function in the algorithm. For Listing 1, HALIVER can generate line 2, which is included here for completeness, based on line 8.

To prove that a *reduction* is correct, *reduction invariant* annotations must be provided for reduction domains. In Listing 2, an example is given of a reduction (line 5) together with its reduction invariant (line 6) and post-state predicate (line 7). Intuitively, a reduction invariant is similar to a loop invariant. First, it must hold before the reduction starts. In our example this means that `count(x)` has the value 0, which is ensured by the previous definition of `count` (line 4). Second, it must be preserved by each step of the reduction. In our example, `count` is bounded by the reduction variable. Finally, after each reduction variable has reached its maximum value, the reduction invariant should imply the post-state predicate of the function. For the example, note that the invariant implies the post-state predicate when `r` has reached the value 10. The actual used value goes to 9, and `r==10` indicates that the reduction is done.

Listing 3. The front-end PVL code for the blur example (Listing 1). We omitted the `decreases` clauses for brevity.

```

1  pure int inp(int x, int y);
2  pure int inp_x_min(); pure int inp_x_max(); pure int inp_y_min(); pure int inp_y_max();
3  pure int blur_y_x_min(); pure int blur_y_x_max();
4  pure int blur_y_y_min(); pure int blur_y_y_max();
5
6  ensures \result ≡ (inp(x, y) + inp(x+1, y) + inp(x+2, y))/3;
7  pure int blur_x(int x, int y) = (inp(x, y) + inp(x+1, y) + inp(x+2, y))/3;
8
9  ensures \result ≡ ((inp(x, y) + inp(x+1, y) + inp(x+2, y))/3
10 + (inp(x, y+1) + inp(x+1, y+1) + inp(x+2, y+1))/3
11 + (inp(x, y+2) + inp(x+1, y+2) + inp(x+2, y+2))/3)/3;
12 pure int blur_y(int x, int y) = (blur_x(x, y) + blur_x(x, y+1) + blur_x(x, y+2))/3;
13
14 requires inp_x_min() ≡ blur_y_x_min() ∧ inp_x_max() ≡ blur_y_x_max()+2
15 ∧ inp_y_min() ≡ blur_y_y_min() ∧ inp_y_max() ≡ blur_y_y_max()+2;
16 ensures (∀ x, y; blur_y_x_min() ≤ x ∧ x < blur_y_x_max() ∧ blur_y_y_min() ≤ y ∧ y <
17 blur_y_y_max());
18 blur_y(x, y) ≡ ((inp(x, y) + inp(x+1, y) + inp(x+2, y))/3
19 + (inp(x, y+1) + inp(x+1, y+1) + inp(x+2, y+1))/3
20 + (inp(x, y+2) + inp(x+1, y+2) + inp(x+2, y+2))/3)/3;
void pipeline() { }

```

3.2 Front-end verification approach

For verifying the algorithm part of a HALIDE program, an annotated HALIDE algorithm is encoded into annotated PVL code. Listings 3 and 4 show how HALIVER translates the examples of Listings 1 and 2, respectively. Input buffers are translated into abstract functions to verify the pipeline w.r.t. arbitrary input. The bounds of input buffers and functions are modelled via functions that are abstract if the bound is unknown or otherwise return a concrete value. For example, the `inp` buffer of the blur example is translated to a function `inp` in Listing 3 (line 1), with its bounds represented by the pure functions on line 2.

Update-free HALIDE functions are translated directly into `pure` PVL functions, and post-state predicates are translated into postconditions of these functions. In the example, `blur_x` and `blur_y` are translated to the functions on lines 6–7 and 9–12 of Listing 3, respectively, and the `ensures` lines express the postconditions of those functions, using `\result` to refer to the expected result.

The pre- and postconditions of a HALIDE algorithm are translated into a PVL lemma to be checked by VERCORS. In the example, lines 14–19 of Listing 3 address the pre- and postconditions on lines 1–2 of Listing 1. On line 20, a method called `pipeline` is given, which represents the HALIDE pipeline.

For an update definition, references to itself are replaced by references to the previous definition, thus the output of one definition is the input of the next.

For a reduction, the initialisation and update parts are translated into separate functions, and reduction domain variables are explicitly added as function parameters. Listing 4 illustrates this for the `cnt` example. The function `count0` on line 8 corresponds to the initialisation (line 3 of Listing 2), with the translated post-state predicate on line 6. The function `count1r` (lines 13–14) corresponds to the update function (line 5 of Listing 2). Note that the annotation on line

Listing 4. The front-end PVL code for the reduction example of Listing 2.

```

1  decreases;
2  pure int inp(int x, int y);
3  decreases;
4  pure int inp_x_min(); pure int inp_x_max(); pure int inp_y_min(); pure int inp_y_max();
5
6  ensures \result  $\equiv$  0;
7  decreases;
8  pure int count0(int x) = 0;
9
10 requires  $0 \leq r \wedge r \leq 10$ ;
11 ensures  $(0 \leq \text{\result} \wedge \text{\result} \leq r)$ ;
12 decreases r;
13 pure int count1r(int x, int r) = r  $\equiv$  0 ? count0(x)
14   : inp(x, r-1) > 0 ? count1r(x, r-1) + 1 : count1r(x, r-1);
15
16 ensures  $(0 \leq \text{\result} \wedge \text{\result} \leq 10)$ ;
17 decreases;
18 pure int count(int x) = count1r(x, 10);

```

10 refers to the reduction domain. The reason for using references to $r-1$ on line 14 is that the result of the whole computation corresponds to r with its maximum value 10 (see line 18). This is computed by recursively decrementing r . The invariant on line 6 of Listing 2 is translated into the postcondition of `count1r` (line 11), reflecting that the invariant should hold after each reduction iteration. For the `decreases r` annotation added on line 12, `VERCORS` will try to prove that this recursive function terminates. The reduction postcondition is represented by the `ensures` annotation on line 16.

Guarantees. For the front-end verification approach, `HALIVER` straightforwardly encodes a `HALIDE` function without reductions, as it defines the function pointwise in PVL. For reductions, `HALIVER` mimics the iterative updates with recursion, as shown in the `cnt` example of Listings 2 and 4. `HALIVER` adds `decreases` clauses to check that the recursive functions terminate.

With `HALIVER`'s approach, functional correctness of the algorithm part can be proven. Since memory safety depends on how a `HALIDE` algorithm is compiled into actual code according to a schedule, this is checked using the back-end verification approach.

3.3 Back-end verification approach

For verifying a `HALIDE` algorithm with a schedule, `HALIVER` adds annotations to the generated C code that can be checked by `VERCORS`. First, `HALIVER` generates read and write permissions and preconditions for functions used in definitions. This generation of permissions makes it possible to keep the annotations of `HALIDE` algorithms concise, since the user does not have to specify permissions. Second, `HALIVER` transforms the annotations and adds them to the intermediate representation used by the `HALIDE` compiler. Finally, `HALIVER` adds the annotations to the code, during the code generation of the `HALIDE` compiler.

Listing 5. A schedule for the blur example (Listing 1), together with the loop nest the HALIDE compiler produces, given in the intermediate representation of HALIDE. The `blur_y` bounds are assumed to be from 0 up to 1,024 for dimensions `x` and `y`.

```

1 blur_y.split(y, yo, yi, 8).parallel(yo).split(x, xo, xi, 2).unroll(xi);
2 blur_x.store_at(blur_y, yo).compute_at(blur_y, yi).split(x, xo, xi, 2).unroll(xi);
3 // Below is the loop nest produced (not part of the schedule)
4 produce blur_y:
5   parallel y.yo in [0, 127]:
6     store blur_x:
7       for y.yi in [0, 7]:
8         produce blur_x:
9           for y:
10            for x.xo in [0, 511]:
11              unrolled x.xi in [0, 2]:
12                blur_x(...) = ...
13            consume blur_x:
14              for x.xo in [0, 511]:
15                unrolled x.xi in [0, 2]:
16                  blur_y(...) = ...

```

Annotation generation. Since HALIDE algorithms consist of pure point-wise functions, permissions are relatively straightforward: for a function $f(x, \dots)$, HALIVER generates the write permission $\text{Perm}(f(x, \dots), 1\setminus 1)$. For the blur example from Listing 1, HALIVER generates `blur_x.context(Perm(blur_x(x,y), 1\setminus 1))` and `blur_y.context(Perm(blur_y(x,y), 1\setminus 1))` for function `blur_x` and `blur_y`, respectively.

For update functions and reductions, HALIVER generates (1) read permissions for function values that are not being updated, and (2) a pre-state predicate, using the post-state predicate of the previous update step.

Once a function is fully defined, read permission is given to all values wherever the function is used, along with a context predicate containing any intermediate annotations of the function.

Transformation of annotations. Next, HALIVER transforms the annotations according to the schedule given by the user and associates them with the corresponding parts of the optimised HALIDE program expressed in HALIDE’s intermediate language.

HALIVER supports the `split`, `fuse`, `parallel`, `unroll`, `store_at`, `reorder` and `compute_at` scheduling directives. Of the most commonly used directives in the HALIDE example apps⁵, only `vectorize` is not supported because VERCORS does not yet support verification of vectorised code as produced by HALIDE.⁶ With these directives, HALIVER provides the means to verify optimised programs w.r.t. memory locality, parallelism and recomputation. This is the optimisation space in which HALIDE resides [24]. We illustrate the meaning of these directives

⁵ <https://github.com/halide/Halide/tree/main/apps>

⁶ The `vectorize` scheduling directive is the same as the `unroll` directive from the perspective of transforming annotations. So they can be treated exactly the same and already are in HaliVer.

with an example. Listing 5 shows a schedule for `blur` on lines 1–2, and below that the *loop nest* structure of the resulting program. Loop nests are program statements of nested for loops. The loops can be sequentially executed or be parallelized, unrolled or vectorized. The allocation of space for a function result is indicated by `store`, and `produce` and `consume` refer to writing and reading function results, respectively. This loop nesting corresponds to the actual code produced by the HALIDE compiler.

Assuming that the output dimensions in the example are both of size 1,024, the directive `split(y, yo, yi, 8)` (line 1 of Listing 5) splits the dimension `y` into two nested dimensions `y.yo` (line 5) and `y.yi` (line 7) of sizes 128 and 8, respectively. HALIVER similarly renames references to `y` in annotations. The `parallel(yo)` directive (line 1) expresses that `y.yo` should be executed in parallel (line 5). The `store_at(blur_y, yo)` directive (line 2) expresses that `blur_x` must be stored at the start of the `y.yo` loop (line 6). The directive `compute_at(blur_y, yi)` (line 2) defines that the values for `blur_x` should be produced at `y.yi` (line 8). The directive `unroll(xi)` (line 1 and 2) expresses that the dimension `xi` should be completely unrolled.

The `for` loops are sequential. In this example, `fuse` and `reorder` are not used; they express that two dimensions should be fused into one and the nesting order of the loops should be changed, respectively.

HALIVER moves bottom-up through the program, constructing loop invariants for each loop by taking the constructed state predicates from the loop body and extending them with quantifications over the loop variables. Below, we give an example of this exact process for the `blur` example of Listing 1. Table 3 in the appendix explains the approach in a more general way.

Encoding of HALIDE program. Finally, HALIVER adds annotations to the C code during the code generation of the HALIDE compiler. As an example, we show how HALIVER adds annotations of the `blur_y` function of Listing 1 with the schedule of Listing 5. The result of this can be found in Listing 6. It shows the structure of the whole program, but is focussed on the code below the `consume blur_x` node (line 13 of Listing 5). The complete C code can be found in the appendix in Listings 7–9.

First, HALIVER updates its pipeline annotations (lines 1–2 of Listing 1), to match the flattened array structure the HALIDE back-end uses, and adds them to the function contract (lines 8–15 of Listing 6). HALIVER also uses the HALIDE definition of division (`hdiv`), i.e., Euclidean⁷ [14] with $x/0 \equiv 0$.

Next, HALIVER transforms the annotations added to the `blur_y` function, before it adds them to any loop nest. The HALIDE compiler flattens the two-dimensional function `blur_y(x, y)` into a one-dimensional array `blur_y[y*1024 + x]`, so HALIVER does the same for all function references in the annotations. Next, from the schedule, the directive `split(x, xo, xi, 2)` splits `x` into `xo` and `xi` of sizes 512 and 2, respectively. A similar split is performed for `y`. The

⁷ The HALIDE compiler uses bit operators to define euclidean division. However, bit operators are not supported in VERCORS, so HALIVER uses an equivalent definition.

Listing 6. The C code and annotations the HALIDE compiler produces together with HALIVER for the function `blur_y`, focussing on the `consume blur_x` node (see line 13 of Listing 5). Listings 7–9 from the appendix give the complete encoding for the `blur_y` pipeline.

```

1  struct halide_dimension_t {int32_t min, max;};
2  struct buffer {int32_t dimensions; struct halide_dimension_t *dim; int32_t *host;};
3  int div_eucl(int x, int y);
4  //@ pure int hdiv(int x, int y) = y == 0 ? 0 : div_eucl(x, y);
5  //@ pure int p_i(int x);
6  /*@ ... // Buffers annotations
7  context (forall int x, int y; 0 <= x < 1024 & 0 <= y < 1026; inpb -> host[y * 1026 + x] == p_i(y * 1026 + x));
8  // Pipeline preconditions
9  requires inpb -> dim[0].min == blur_yb -> dim[0].min & inpb -> dim[0].max == blur_yb -> dim[0].max + 2;
10 requires inpb -> dim[1].min == blur_yb -> dim[1].min & inpb -> dim[1].max == blur_yb -> dim[1].max + 2;
11 // Pipeline postconditions
12 ensures (forall int x, int y; 0 <= x < 1024 & 0 <= y < 1024; blur_yb -> host[y * 1024 + x] == hdiv(
13   hdiv(inpb -> host[y * 1026 + x + 1027] + inpb -> host[y * 1026 + x + 1028] + inpb -> host[y * 1026 + x + 1026], 3) +
14   hdiv(inpb -> host[y * 1026 + x + 2053] + inpb -> host[y * 1026 + x + 2054] + inpb -> host[y * 1026 + x + 2052], 3) +
15   hdiv(inpb -> host[y * 1026 + x + 1] + inpb -> host[y * 1026 + x + 2] + inpb -> host[y * 1026 + x], 3), 3)); @*/
16 int blur_3(struct buffer *inpb, struct buffer *blur_yb) {
17   int32_t * blur_y = blur_yb -> host;
18   int32_t * inp = inpb -> host;
19   // produce blur_y
20   #pragma omp parallel for
21   for (int yo = 0; yo < 0 + 128; yo++)
22     ... // Annotations blur_y.yo
23     {
24       int64_t _2 = 10240;
25       int32_t *blur_x = (int32_t *)malloc(sizeof(int32_t) * _2);
26       int32_t _t11 = (yo * 8);
27       ... // Annotations blur_y.yi
28       for (int yi = 0; yi < 0 + 8; yi++)
29         {... // produce blur_x
30          // consume blur_x
31          int32_t _t16 = (yi + _t11) * 512;
32          int32_t _t15 = yi * 512;
33          /*@ loop_invariant 0 <= xo & xo <= 0 + 512;
34           loop_invariant (forall int x, int y; 0 <= x & x < 1024 & yo * 8 <= y & y < yo * 8 + 10;
35            Perm(&blur_x[(y - yo * 8) * 1024 + x], 1 \ 2));
36           loop_invariant (forall int xo, int y; 0 <= xo & xo < 1024 & yo * 8 + yi <= y & y <= yo * 8 + yi + 2;
37            blur_x[(y - yo * 8) * 1024 + xo] == hdiv(p_i(y * 1026 + xo) + p_i(y * 1026 + xo + 1) + p_i(y * 1026 + xo
38              + 2), 3));
39           loop_invariant (forall int xif, int xof; 0 <= xof & xof < 512 & 0 <= xif & xif < 2;
40            Perm(&blur_y[(yo * 8 + yi) * 1024 + xof * 2 + xif], 1 \ 1));
41           loop_invariant (forall int xof, int xif; 0 <= xof & xof < xo & 0 <= xif & xif < 2; blur_y[(
42             yo * 8 + yi) * 1024 + xof * 2 + xif] ==
43             hdiv(hdiv(p_i((yo * 8 + yi) * 1026 + xof * 2 + xif) + p_i((yo * 8 + yi) * 1026 + xof * 2 + xif + 1) + p_i((
44               yo * 8 + yi) * 1026 + xof * 2 + xif + 2), 3) +
45             hdiv(p_i((yo * 8 + yi) * 1026 + xof * 2 + xif + 1026) + p_i((yo * 8 + yi) * 1026 + xof * 2 + xif + 1027) + p_i
46               ((yo * 8 + yi) * 1026 + xof * 2 + xif + 1028), 3) +
47             hdiv(p_i((yo * 8 + yi) * 1026 + xof * 2 + xif + 2052) + p_i((yo * 8 + yi) * 1026 + xof * 2 + xif + 2053) + p_i
48               ((yo * 8 + yi) * 1026 + xof * 2 + xif + 2054), 3), 3)); @*/
49         for (int xo = 0; xo < 0 + 512; xo++)
50           {
51             int32_t _t9 = (xo + _t15);
52             blur_y[(xo + _t16) * 2] = div_eucl(blur_x[_t9 * 2] + blur_x[_t9 * 2 + 1024] +
53               blur_x[_t9 * 2 + 2048], 3);
54             blur_y[(xo + _t16) * 2 + 1] = div_eucl(blur_x[_t9 * 2 + 1] + blur_x[_t9 * 2 + 1025]
55               + blur_x[_t9 * 2 + 2049], 3);
56           } // for xo
57         } // for yi
58       free(blur_x);
59     } // for yo
60   return 0;
61 }

```

generated annotation `context (Perm(blur_y(x,y), 1\1))` becomes `context Perm(&blur_y[(yo*8+yi)+ xo*2+ xi], 1\1)`.

For the annotation `ensures(blur_y(x,y)==((inp(x,y)+... , HALIVER` replaces the calls to `inp(x,y)` with calls to an abstract pure function `p_i`. This is done because quantification instantiation in `VERCORS` can become unstable if `inp` is used frequently. Where `inp` is used in the code, `HALIVER` adds annotations stating that `inp` and `p_i` have the same value (line 7 of Listing 6).

`HALIVER` adds these annotations to the first loop nest, starting bottom up. In Listing 5, this is `xi`, but since this loop is unrolled, additional annotations are not needed. After passing this loop nest, anything for `xi=0` and `xi=1` now holds. `HALIVER` changes the annotations by quantifying over `xi`'s domain. It uses `xif` as variable and changes any references to `xi` towards `xif`. The resulting permissions are $(\forall xif; 0 \leq xif \wedge xif < 2; \text{Perm}(\text{blur_y}[(yo*8+yi)+ xo*2+xif], 1 \setminus 1))$. The other annotations are processed in a similar way.

Next, `HALIVER` arrives at the loop nest for `xo`, which needs loop invariants. First, the tool adds the bounds of the `xo` dimension (line 33 of Listing 6). The annotation is transformed depending on whether it was a `requires`, `ensures` or `context` annotation. The write permission (`context`), should hold before the loop starts and after the loop ends. Therefore, `HALIVER` adds the permission, but quantifies over dimension `xo`, which results in a loop invariant (lines 38–39 of Listing 6). The `ensure` annotation does not hold at the start of the loop, but after each iteration of the loop, one more value for `xo` holds. Therefore, `HALIVER` quantifies over `xof` bounded by zero and the iteration variable `xo`, and replaces occurrences to `xo` with `xof`, which leads to a loop invariant (lines 40–43 of Listing 6). For loops above this loop nest, the `ensure` annotations hold for the whole domain of `xo`, resulting in `ensures` $(\forall xof, xif; 0 \leq xof \wedge xof < 512 \wedge 0 \leq xif \wedge xif < 2; \text{blur_y}[(yo*8+yi)*1024+xof*2+xif] \equiv \dots$. This annotation is added to the parallel for loop (line 66 of Listing 8).

After constructing the `produce` node for `blur_y`, the `produce` node for `blur_x` is constructed in a similar way. The bound inferencer of `HALIDE` detects it only needs to calculate for `y` values of `8*y0+yi` up to `8*y0+yi+2`. The annotations are transformed respecting that fact. After the `produce` node, the `blur_x` is consumed (line 30 of Listing 6). So for each loop below the `consume` statement, `HALIVER` adds read permission (lines 34–35 of Listing 6s) and the post-state predicate of `blur_x` (lines 36–37 of Listing 6) as context annotations. For the loop of `xo`, this means they are valid for any value of `xo`.

Guarantees. With the back-end verification approach, `HALIVER` can prove that the optimised code produced by the `HALIDE` compiler is correct w.r.t. specifications. Memory safety is proven without any additional effort, as the permission annotations for this are generated automatically. For functional correctness, a specification needs to be provided. For any non-inlined function, an intermediate annotation is required to guide `VERCORS` in correct functional verification.

The approach is sound, but not necessarily complete. One concern is that, since we have not formally proved the correctness of the transformation, our implementation could in principle be wrong. `HALIVER` addresses this by keeping

Table 1. Number of lines of code and annotations for different HALIDE algorithms, schedules and resulting programs, and the verification times required by VERCORS to prove memory safety, given that no annotations are provided by the user. The letters after each schedule denote the used directives: `compute_at`, `fuse`, `parallel`, `reorder`, `split`, `store_at` and `unroll`. F stands for verification failed. Times with † are inconsistent, i.e. they are successfully verified, but can also sometimes fail or timeout.

Name		HALIDE	Sched.	C			
		LoC	Dir.	LoC	LoA.	Loops	T. (s).
blur	V0	38	0	178	60	2	18
	V1-{f,p}	"	2	172	56	1	19
	V2-{c,p,r,s}	"	6	212	74	6	29
	V3-{c,p,s,st,u}	"	8	211	72	5	24
hist	V0	71	2	299	98	11	30
	V1-{c,p,r,u}	"	4	308	99	11	38
	V2-{c,p,r,u}	"	6	311	105	13	48
	V3-{c,p,r,u}	"	13	312	101	13	48
conv	V0	44	0	273	148	7	30
	V1-{c,f,p,u}	"	4	281	145	8	37
	V2-{p,r,s,u}	"	6	302	166	10	209
	V3-{c,p,r,s,u}	"	15	279	148	7	168
gemm	V0	70	0	218	105	3	41
	V1-{c,p,r,s}	"	8	274	136	10	89
	V2-{c,p,r,s}	"	16	342	173	19	196 †
	V3-{c,f,p,r,s,u}	"	24	451	221	31	F
auto	V0	112	0	443	118	19	35
	V1-{c}	"	9	402	139	23	180
	V2-{c,p}	"	12	440	156	27	170
	V3-{c,p,r,s}	"	27	443	152	25	105
camera_pipe	{c,p,r,s,st}	345	27	701	236	25	F
bilateral_grid	{c,p,r,u}	88	18	562	180	39	140
depthwise_separable_conv	{c,p,r,s}	94	13	562	315	44	480

the *pipeline* annotations very close to what the user has written as annotations. These pipeline annotations act as the formal contract that will be verified, and the user can inspect these at any time. If an intermediate annotation is not correctly transformed, the verification will fail, thus remaining sound but not complete. Of course we have not constructed any transformations to be wrong, but even if there is an oversight, we will remain sound. Moreover, in Section 4, we show that our approach works for real world examples.

4 Evaluation

The goal of the evaluation of HALIVER is four-fold. (1) We evaluate that the front-end verification approach of HALIVER can verify functional correctness properties for a representative set of HALIDE algorithms. (2) For the back-end verification approach, the annotations that HALIVER generates and transforms should lead to successful verification for a representative set of HALIDE programs, with schedules that use the most important scheduling directives in different combinations. (3) We evaluate the verification speed for front-end and back-end verification. (4) Lastly, we evaluate how many annotations are needed in HALIVER compared to manually annotating the generated C programs.⁸

Set-up: We used a machine with an 11th Gen Intel(R) Core(TM) i7-11800H @ 2.30GHz with 32GB running Ubuntu 23.04.

⁸ The experiments can be found at <https://github.com/sakehl/HaliVerExperiments>.

Table 2. Number of lines of code and annotations for different HALIDE algorithms, schedules and resulting programs, and the verification times required by VERCORS.

Name	HALIDE		Front-end T. (s)	Sched LoC	C				LoA incr.
	LoC	LoA			LoC	LoA	Loops	T. (s)	
blur V0	38	2	8	0	178	63	2	21	31.5x
V1-{f,p}	"	"	"	2	172	58	1	23	29.0x
V2-{c,p,r,s}	"	"	"	6	212	83	6	52	41.5x
V3-{c,p,s,st,u}	"	"	"	8	211	79	5	97 [†]	39.5x
hist V0	71	10	8	0	299	118	11	34	11.8x
V1-{c,p,r,u}	"	"	"	4	308	118	11	47	11.8x
V2-{c,p,r,u}	"	"	"	6	311	123	13	56	12.3x
V3-{c,p,r,u}	"	"	"	13	312	125	13	64	12.5x
conv V0	44	7	8	0	273	177	7	111	25.3x
layer V1-{c,f,p,u}	"	"	"	4	281	174	8	108	24.9x
V2-{p,r,s,u}	"	"	"	6	302	204	10	283	29.1x
V3-{c,p,r,s,u}	"	"	"	15	279	177	7	207	25.3x
gemm V0	70	12	7	0	218	120	3	43	10.0x
V1-{c,p,r,s}	"	"	"	8	274	169	10	133	14.1x
V2-{c,p,r,s}	"	"	"	16	342	230	19	368	19.2x
V3-{c,f,p,r,s,u}	"	"	"	24	451	310	31	F	25.8x
auto V0	112	15	8	0	443	158	19	152 [†]	10.5x
viz V1-{c}	"	"	"	9	402	210	23	216	14.0x
V2-{c,p}	"	"	"	12	440	235	27	230 [†]	15.7x
V3-{c,p,r,s}	"	"	"	27	443	229	25	192 [†]	15.3x

We used eight characteristic programs from the HALIDE repository.⁹ These are representative HALIDE algorithm examples. They cover all scheduling directives supported by HALIVER, in commonly-used combinations. We removed any scheduling directives that we do not support. As we discuss in Section A of the appendix, VERCORS is unable to deal with large dimensions that are unrolled, thus we removed some `unroll` directives as well.

The original schedule, as found in the HALIDE repository, is indicated with V3 if there are multiple schedules present. For five of these programs we defined annotations that express functional properties. These five programs are also evaluated with the standard schedule (V0), which tries to inline functions as much as possible, and two additional schedules (V1 and V2) we constructed.

Memory safety results: We evaluate 8 HALIDE programs, with in total 23 schedules, and prove data race freedom and memory safety for 21 of them. No user provided annotations are needed. The results can be found in Table 1.

For each case, we provide: the number of lines of code (LoC)¹⁰ for the HALIDE algorithm, without the schedule and number of scheduling directives (Sched. Dir.). For the generated programs (C) we list: lines of code (LoC), lines of annotations (LoA.), number of (parallel) loops (Loops). These numbers indicate how large programs tend to become w.r.t. HALIDE algorithms, and how much annotation effort would be required to manually annotate the programs. Verification running times (T. (s)) are given in seconds, averaged over five runs.

For `camera_pipe`, VERCORS gives a verification failure. It could not prove a `loop_invariant`, but after simplifying parts of the generated C program not related to this specific invariant, it leads to a successful verification. This indicates that the program is too complex for the underlying solvers. We also coded this

⁹ <https://github.com/halide/Halide/tree/main/apps> `gemm` is part of `linear_algebra`.

¹⁰ These lines are counted automatically and indicate the size of the programs.

example in similar PVL code instead of C, which verifies in 193s. We suspect the failure is caused by quantifier instantiation, which instantiates too many quantifiers, resulting in the SMT solver on which VERCORS relies stopping the exploration of quantifiers that are needed for successful verification.

For `gemm V3`, verification fails due to VERCORS not sufficiently rewriting annotations of the `fuse` directive. This is further explained in section A.

Functional correctness results: Next, we evaluate five¹¹ algorithms with annotations and 20 schedules, both for the front-end and back-end. HALIVER proves functional correctness for the front-end, and both functional correctness and data race freedom and memory safety for the back-end for 19 of the 20 schedules. These results are given in Table 2. The table additionally has the amount of user provided annotations (LoA.) and the last column (Ann. incr.) indicates the growth of the annotations. The annotations of the C file (LoC) contain both the generated annotations, which are already present in Table 1 and the transformed user annotations.

For optimised programs, the annotation size is strongly related to the number of loops, as each loop needs its own loop invariants. Front-end verification is successful for all examples and is relatively fast compared to back-end verification. In verification of the C files produced by the back-end verification approach, time increases as the number of scheduling directives increases. Here, `gemm V3` also fails for the same reason as outlined above.

Inconsistent results: For `gemm V2` for the memory benchmarks and for `blur V3` and `auto_viz V0, V2` and `V3`, VERCORS does not always succeed with the verification. In the case of `gemm V2`, the verification sometimes hangs, which is timed out after 10 minutes. In the other cases, VERCORS sometimes gave a verification failure. This inconsistency is due to the non-deterministic nature of the underlying SMT solvers.

Conclusions: With the front-end verification approach of HALIVER we are able to prove functional correctness properties for representative HALIDE algorithms. Using HALIVER’s back-end verification approach, the tool provides correct annotations for the generated C programs. VERCORS successfully verifies all but two programs. However, in the unsuccessful cases, HALIVER runs into limitations of the underlying tools. The verified programs are all verified within ten minutes. Finally, the manual annotation effort required is an order of magnitude larger than the effort required for HALIVER’s approach.

5 Related Work

There is much work on optimising program transformations, either applied automatically or manually [2, 13], sometimes using scheduling languages [3, 7, 8, 10, 24, 25, 30]. The vast majority of this does not address functional correctness.

¹¹ The other three algorithms from the memory safety results are typical image processing pipelines. They are therefore less suitable for checking functional correctness and are not used here.

Work on functional correctness consists of techniques that apply verification every time a program is transformed, and techniques that verify the compiler.

Liu *et al.* [17] propose an approach inspired by scheduling languages, with proof obligations generated when a program is optimised, for automatic verification using Coq. The COGENT language [22] uses refinement proofs, to be verified in ISABELLE/HOL. However, it does not separate algorithms from schedules. In [19, 20] an integer constraint solver and a proof checker are used, respectively, to verify the transformation of a program. In all these approaches, semantics-preservation is the focus, as opposed to specifying the intended behaviour. Model-to-model transformations can be verified w.r.t. the preservation of functional properties [23]. However, that work targets models, not code.

Regarding the verification of compilers, COMPCERT [16] is a framework involving a formally verified C compiler. In [21], HALIDE’s Term Rewriting System, used to reason about the applicability of schedules, is verified using Z3 and Coq. These approaches do not require verification every time an optimisation is applied, but verifying the compiler is time-consuming and complex, and has to be redone whenever the compiler is updated. Furthermore, they focus on semantics-preservation, not the intended behaviour of individual programs.

ALPINIST [29] is most closely related. This tool automatically optimises PVL code, along with its annotations, for verification with VERCORS. It allows the specification of intended behaviour, but it does not separate algorithms from schedules, forcing the user to reason about the technical details of parallelisation.

6 Conclusions & Future Work

We presented HALIVER, a tool for verifying optimised code by exploiting the strengths of scheduling languages and deductive verification. It allows focussing on functionality when annotating programs, keeping annotations succinct.

For future work, we want to extend the HALIVER tool with aspects not directly supported by VERCORS, such as vectorisation. The master thesis of [26] defines a natural semantics for HALIDE. We want to formalise our front-end PVL encoding with an axiomatic semantics to match this semantics. We also want to investigate the inconsistent results and see whether annotations with quantifiers can be rephrased to allow VERCORS to be more consistent. In this work we have focussed on parallel CPU code, but we have designed our approach to be extendable to GPU code produced by HALIDE.

With the current expressiveness of the annotations, when reduction domains are present, HALIVER proves functional correctness for specific inputs. For example, in Listing 2 we can prove that `count(x)==9` if we require that `input(x,y)==x`. This can also be done for any input if the reduction domain is of known size, but then many annotations are needed. To make the annotations concise, a user needs to be able to use axiomatic data types¹² and pure functions

¹² <https://vercors.ewi.utwente.nl/wiki/#axiomatic-data-types>

in their annotations. We expect that these annotations can be similarly transformed by our approach, and that is thus orthogonal to this contribution, but this is planned as future work.

Most HALIDE programs use floating point numbers. These are currently modelled as reals in VERCORS. How to efficiently verify programs with floats using deductive verifiers is still an open research question. Once this is addressed, HALIVER will be able to give better guarantees.

We require that the bounds of a HALIDE program are set to concrete values for our back-end verification approach. HALIVER transforms the annotations the same way for not know bounds, but the underlying tools have difficulty verifying these programs. With unknown bounds, we end up with nonlinear arithmetic due to the flattening of multi-dimensional functions on one-dimensional arrays. This is generally undecidable, so the SMT solvers that VERCORS rely on cannot handle it. We will investigate if there are ways to tackle this in our domain-specific case.

Acknowledgements This work is carried out in the context of the NWO TTW ChEOPS project 17249. We want to thank Jan Martens for their discussions and feedback on this work.

References

1. Amighi, A., Haack, C., Huisman, M., Hurlin, C.: Permission-based separation logic for multithreaded Java programs. *LMCS* **11**(1) (2015)
2. Bacon, D., Graham, S., Sharp, O.: Compiler Transformations for High-Performance Computing. *ACM Computing Surveys* **26**(4), 345–420 (1994)
3. Baghdadi, R., Ray, J., Romdhane, M.B., Sozzo, E.D., Akkas, A., Zhang, Y., Suriana, P., Kamil, S., Amarasinghe, S.P.: Tiramisu: A Polyhedral Compiler for Expressing Fast and Portable Code. In: *CGO*. pp. 193–205. IEEE (2019). <https://doi.org/10.1109/CGO.2019.8661197>
4. Becker, N., Müller, P., Summers, A.J.: The axiom profiler: Understanding and debugging smt quantifier instantiations. In: *Tools and Algorithms for the Construction and Analysis of Systems: 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings, Part I* 25. pp. 99–116. Springer (2019)
5. Blom, S., Darabi, S., Huisman, M., Oortwijn, W.: The VerCors Tool Set: Verification of Parallel and Concurrent Software. In: Polikarpova, N., Schneider, S. (eds.) *Integr. Form. Methods*. pp. 102–110. *Lecture Notes in Computer Science*, Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-66845-1_7
6. Bornat, R., Calcagno, C., O’Hearn, P., Parkinson, M.: Permission accounting in separation logic. In: *POPL*. pp. 259–270 (2005)
7. Chame, C.C.J., Hall, M.: CHILL: A framework for composing high-level loop transformations. 08-897, University of Southern California (2008)

8. Chen, T., Moreau, T., Jiang, Z., Zheng, L., Yan, E., Cowan, M., Shen, H., Wang, L., Hu, Y., Ceze, L., Guestrin, C., Krishnamurthy, A.: TVM: An Automated End-to-End Optimizing Compiler for Deep Learning. In: 13th USENIX Symp. Oper. Syst. Des. Implement. OSDI 18. pp. 578–594. USENIX Association, USA (2018)
9. Dross, C., Conchon, S., Paskevich, A.: Reasoning with Triggers. Research Report RR-7986, INRIA (Jun 2012), <https://inria.hal.science/hal-00703207>
10. Hagedorn, B., Elliott, A.S., Barthels, H., Bodik, R., Grover, V.: Fireiron: A Data-Movement-Aware Scheduling Language for GPUs. In: Proc. ACM Int. Conf. Parallel Archit. Compil. Tech. pp. 71–82. ACM, Virtual Event GA USA (Sep 2020). <https://doi.org/10.1145/3410463.3414632>
11. Hähnle, R., Huisman, M.: Deductive Software Verification: From Pen-and-Paper Proofs to Industrial Tools. In: Computing and Software Science - State of the Art and Perspectives. LNCS, vol. 10000, pp. 345–373. Springer (2019)
12. Hijma, P., Heldens, S., Selocco, A., van Werkhoven, B., Bal, H.: Optimization Techniques for GPU Programming. ACM Computing Surveys **55**(11), 239:1–239:81 (2023)
13. Kowarschik, M., Weiß, C.: An Overview of Cache Optimization Techniques and Cache-Aware Numerical Algorithms. In: Algorithms for Memory Hierarchies. LNCS, vol. 2625, pp. 213–232. Springer (2003)
14. Leijen, D.: Division and modulus for computer scientists (July 2003), <https://www.microsoft.com/en-us/research/publication/division-and-modulus-for-computer-scientists/>, short note about division definitions in programming languages
15. Leiserson, C.E., Thompson, N.C., Emer, J.S., Kuszmaul, B.C., Lampson, B.W., Sanchez, D., Schardl, T.B.: There’s plenty of room at the top: What will drive computer performance after Moore’s law? Science **368**(6495) (2020). <https://doi.org/10.1126/science.aam9744>
16. Leroy, X.: A formally verified compiler back-end. Journal of Automated Reasoning **43**(4), 363–446 (2009)
17. Liu, A., Bernstein, G.L., Chlipala, A., Ragan-Kelley, J.: Verified tensor-program optimization via high-level scheduling rewrites. Proc. ACM Program. Lang. **6**(POPL), 55:1–55:28 (Jan 2022). <https://doi.org/10.1145/3498717>
18. Müller, P., Schwerhoff, M., Summers, A.: Viper - a verification infrastructure for permission-based reasoning. In: VMCAI (2016)
19. Namjoshi, K.S., Singhania, N.: Loopy: Programmable and formally verified loop transformations. In: International Static Analysis Symposium. pp. 383–402. Springer (2016)
20. Namjoshi, K.S., Xue, A.: A Self-certifying Compilation Framework for WebAssembly. In: International Conference on Verification, Model Checking, and Abstract Interpretation. pp. 127–148. Springer (2021)
21. Newcomb, J.L., Adams, A., Johnson, S., Bodik, R., Kamil, S.: Verifying and Improving Halide’s Term Rewriting System with Program Synthesis. Proc. ACM Program. Lang. **4**(OOPSLA), 166:1–166:28 (Nov 2020). <https://doi.org/10.1145/3428234>
22. O’Connor, L., Chen, Z., Rizkallah, C., Jackson, V., Amani, S., Klein, G., Murray, T., Sewell, T., Keller, G.: Cogent: Uniqueness Types and Certifying Compilation. Journal of Functional Programming **31**(e25), 1–66 (2021)
23. de Putter, S., Wijs, A.: Verifying a verifier: on the formal correctness of an LTS transformation verification technique. In: FASE. pp. 383–400. Springer (2016)
24. Ragan-Kelley, J., Adams, A., Sharlet, D., Barnes, C., Paris, S., Levoy, M., Amarasinghe, S., Durand, F.: Halide: Decoupling algorithms from schedules for

- high-performance image processing. *Commun. ACM* **61**(1), 106–115 (Dec 2017). <https://doi.org/10.1145/3150211>
25. Ragan-Kelley, J., Barnes, C., Adams, A., Paris, S., Durand, F., Amarasinghe, S.: Halide: A Language and Compiler for Optimizing Parallelism, Locality, and Recomputation in Image Processing Pipelines. *SIGPLAN Not.* **48**(6), 519–530 (Jun 2013). <https://doi.org/10.1145/2499370.2462176>
 26. Reinking, A., Bernstein, G., Ragan-Kelley, J.: Formal Semantics for the Halide Language. Master’s thesis, EECS Department, University of California, Berkeley (2020, May)
 27. Safari, M., Huisman, M.: Formal verification of parallel stream compaction and summed-area table algorithms. In: *International Colloquium on Theoretical Aspects of Computing*. pp. 181–199. Springer (2020)
 28. Safari, M., Oortwijn, W., Joosten, S., Huisman, M.: Formal Verification of Parallel Prefix Sum. In: Lee, R., Jha, S., Mavridou, A. (eds.) *NASA Form. Methods*. pp. 170–186. *Lecture Notes in Computer Science*, Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-55754-6_10
 29. Sakar, Ö., Safari, M., Huisman, M., Wijs, A.: Alpinist: An Annotation-Aware GPU Program Optimizer. In: *TACAS, LNCS*, vol. 13244, pp. 332–352. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-99527-0_18
 30. Zhang, Y., Yang, M., Baghdadi, R., Kamil, S., Shun, J., Amarasinghe, S.P.: GraphIt: a high-performance graph DSL. *Proc. ACM Program. Lang.* **2**(OOPSLA), 1–30 (2018)

A Quantification & Triggers

The verification approach of HALIVER relies heavily on modelling the specification with **forall** quantifiers. For a quantifier to be realised, a *trigger* [9] must be instantiated, which determines when exactly VERCORS can use the quantifier. For example, $\forall i; 0 \leq i < 10; a[i] \geq 0$ gets the trigger $a[i]$ and when VERCORS sees the value $a[5]$ it can test the quantifier for $i=5$ to see if it is valid. A condition like $\forall i, j; 0 \leq i < 10 \wedge 0 \leq j < 5; a[5*i+j]$ has no trigger, because when VERCORS sees $a[5]$ it cannot know which values i and j should have. This can be overcome in this case because the quantifier is equivalent to $\forall i_j; 0 \leq i_j < 10*5; a[i_j]$, for which $a[i_j]$ can be used as a trigger. VERCORS rewrites the conditions for some of these cases. Also, since VERCORS introduces more quantifiers when encoding parallel blocks, this rewriting *must* be done by VERCORS.

In some cases, quantifiers are introduced for which no rewriting rules exist yet. For example, when the **fuse** directive is used, it can introduce expressions such as $\forall i, j; 0 \leq i < 10 \wedge 3 \leq j < 6; a[5*i+(j \% 5)]$. We plan to investigate which additional safe rewrites can be performed for these triggers, either when transforming annotations, or inside VERCORS.

Additionally, with trigger instantiation we may run into the problem of matching loops [4]. For example, the quantifier $\forall i; \dots f[i] \equiv x[i] + y[i+1]$, we could use f as a trigger, but if a quantifier with $\forall i; \dots x[i] \equiv f[i+1]$ has a trigger $x[i]$, the first instantiated trigger could trigger the second, which loops back to the first trigger and continues indefinitely. We believe that we run into this problem when verifying HALIDE files that unroll a large dimension with the **unroll** directive. For example, if we want to prove that $i; 0 \leq x \wedge x < 10; f[x] \equiv x$ and we have unrolled the function f like $f[0] = 0; \dots f[9] = 9$, the verifier needs to realise that these 10 lines of code *together* help prove the quantifier. In future work we intend to investigate how to deal with this problem. For now, we have chosen to do our evaluation in section 4 only for small unrolled dimensions, as VERCORS gives unreliable results otherwise.

B Code Examples

Listings 7–9 present an example of code produced by HALIVER. This code is a translation of Listing 1, together with the schedule as given in Listing 5.

C Transformation of annotations for back-end verification approach

In Table 3 we explain how the annotations are transformed in a general sense for back-end verification.

Table 3. Here we show how to transform a function definition’s annotations. When we write $E \mapsto E'$ we mean that we replace the expression (or annotation) E with the expression E' . By $E[x \mapsto e]$ we mean that we replace every occurrence of the variable x in E with e . If a dimension x is not scheduled, it is marked as **serial**(x), which represents a sequential loop.

Loop nests	A loop nest is created for each function definition. We store the annotations which were attached to the function definition. We traverse the loop nest from bottom to top, transforming the annotations at each loop. The loops in the loop nest that we have not yet visited are called the outer loops. We always do two things when we pass through a loop: 1. append annotations based on the current stored annotations; 2. modify the stored annotations for use in the outer loops.
Without reductions	When passing a loop for the dimension x , HALIVER transforms the annotation P for the remaining outer loops by quantifying over x : $P \mapsto (\forall x f < x.\max; P[x \mapsto xf])$
<code>parallel</code> (x)	The parallel block contract uses the stored annotations as is.
<code>unroll</code> (x)	The code is unrolled and requires no further annotations
<code>serial</code> (x)	The annotation P becomes loop invariant and we quantify over x . The limits depend on whether it is a pre- or postcondition.
<code>ensures</code>	$P \mapsto \text{loop_invariant } (\forall x f; x.\min \leq x f < x; P[x \mapsto xf])$
<code>requires</code>	$P \mapsto \text{loop_invariant } (\forall x f; x \leq x f < x.\max; P[x \mapsto xf])$
<code>context</code>	$P \mapsto \text{loop_invariant } (\forall x f; x.\min \leq x f < x.\max; P[x \mapsto xf])$
With reductions	For loop nests with reductions, we only use their associated reduction invariant (denoted by I).
Reduction dimensions r	If we passed a previous reduction dimension (r_{prev}), replace that reduction variable with its min: $I \mapsto I[r_{\text{prev}} \mapsto r_{\text{prev}}.\min]$ and store that we passed a new reduction invariant. We use the, possibly updated, reduction invariant both here as a loop invariant and store it for the outer loops.
Non-reduction dimension x	First we transform the reduction invariant I into a precondition P and a postcondition Q . If we have not yet passed a reduction dimension, we look up the next reduction invariant r_{next} and construct $P = I$ and $Q = I[r_{\text{next}} \mapsto r_{\text{next}} + 1]$. When we have passed reduction dimensions, get the last reduction invariant r_{prev} and construct: $P = I[r_{\text{prev}} \mapsto r_{\text{prev}}.\min]$ and $Q = I[r_{\text{prev}} \mapsto r_{\text{prev}}.\max]$. Similar to the case without reductions, we quantify the invariant over x and store it for the remaining outer loops: $I \mapsto (\forall x f; x.\min \leq x f < x.\max; I[x \mapsto xf])$
<code>parallel</code> (x)	Add the constructed P as a precondition and Q as a postcondition to the parallel block.
<code>serial</code> (x)	Similar to the case without reductions, but for P and Q . <code>loop_invariant</code> ($\forall x f; x \leq x f < x.\max; P[x \mapsto xf]$) <code>loop_invariant</code> ($\forall x f; x.\min \leq x f < x.\max; Q[x \mapsto xf]$)
Manipulating dimensions	
<code>split</code> ($x, x_0 x_1, f$)	Splits x , we replace x and place a guard: $P \mapsto x_0 * f + x_1 < x.\max \Rightarrow P[x \mapsto x_0 * f + x_1]$
<code>fuse</code> (f, x, y)	Fuse dimension x and y together: $P \mapsto P[x \mapsto f \% x.\text{extent}, y \mapsto f / x.\text{extent}]$ The extent is size of dimensions x .
<code>reorder</code> (x, y)	Reorders the dimensions x and y . Sets x below y in the loop nest. A dimension can be a reduction.
Order of calculation and storage	
<code>f.compute_at</code> (g, x)	Computes f at the loop of dimension x of the loop nest for function g . The dimensions for f are changed, HALIVER makes sure that the annotations respect this.
<code>f.store_at</code> (g, x)	Stores f at the loop of dimension x of the loop nest for function g .

Listing 7. Back-end C code with annotations provided by HALIVER for the blur example of Listing 1 (1/3).

```

1  #include <stdint.h>
2  #include <stdlib.h>
3
4
5  // Euclidean division is defined internally in VerCors
6  //@ pure int hdiv(int x, int y) = y == 0 ? 0 : \euclidean_div(x, y);
7  /*@
8   requires y != 0;
9   ensures \result == \euclidean_div(x, y);
10 /*/
11 inline int div_eucl(int x, int y)
12 {
13     int q = x/y;
14     int r = x%y;
15     return r<0 ? q + (y > 0 ? -1 : 1) : q;
16 }
17
18 struct halide_dimension_t {int32_t min, max;};
19 struct buffer {int32_t dimensions; struct halide_dimension_t *dim; int32_t *host;};
20 pure int p_i(int x);
21 /*@
22 // Buffer Annotations
23 context inpb != NULL ** Perm(inpb, 1\2);
24 context Perm(inpb->dim, 1\2) ** inpb->dim != NULL;
25 context \pointer_length(inpb->dim) == 2;
26 context Perm(inpb->host, 1\2) ** inpb->host != NULL;
27 context Perm(&inpb->dim[0], 1\2);
28 context Perm(inpb->dim[0].min, 1\2) ** Perm(inpb->dim[0].max, 1\2);
29 context Perm(&inpb->dim[1], 1\2);
30 context Perm(inpb->dim[1].min, 1\2) ** Perm(inpb->dim[1].max, 1\2);
31 context \pointer_length(inpb->host) == 1026*1026;
32 context blur_yb != NULL ** Perm(blur_yb, 1\2);
33 context Perm(blur_yb->dim, 1\2) ** blur_yb->dim != NULL;
34 context \pointer_length(blur_yb->dim) == 2;
35 context Perm(blur_yb->host, 1\2) ** blur_yb->host != NULL;
36 context Perm(&blur_yb->dim[0], 1\2);
37 context Perm(blur_yb->dim[0].min, 1\2) ** Perm(blur_yb->dim[0].max, 1\2);
38 context Perm(&blur_yb->dim[1], 1\2);
39 context Perm(blur_yb->dim[1].min, 1\2) ** Perm(blur_yb->dim[1].max, 1\2);
40 context \pointer_length(blur_yb->host) == 1024*1024;
41 context blur_yb->host != inpb->host;
42 context inpb->dim[0].min == 0 ^ inpb->dim[0].max == 1026;
43 context inpb->dim[1].min == 0 ^ inpb->dim[1].max == 1026;
44 context (\forall int x, int y; 0<=x ^ x<1026 ^ 0<=y ^ y<1026; Perm(&inpb->host[y*1026 + x
45 ], 1\2));
46 context (\forall int x, int y; 0<=x ^ x<1026 ^ 0<=y ^ y<1026; inpb->host[y*1026 + x] == p_i
47 (y*1026 + x));
48 context blur_yb->dim[0].min == 0 ^ blur_yb->dim[0].max == 1024;
49 context blur_yb->dim[1].min == 0 ^ blur_yb->dim[1].max == 1024;
50 context (\forall int x, int y; 0<=x ^ x<1024 ^ 0<=y ^ y<1024; Perm(&blur_yb->host[y*1024
51 + x], 1\1));
52 // Pipeline preconditions
53 requires inpb->dim[0].min == blur_yb->dim[0].min ^ inpb->dim[0].max == blur_yb->dim[0].
54 max+2;
55 requires inpb->dim[1].min == blur_yb->dim[1].min ^ inpb->dim[1].max == blur_yb->dim[1].
56 max+2;
57 // Pipeline postconditions
58 ensures (\forall int x, int y; 0<=x ^ x<1024 ^ 0<=y ^ y<1024; blur_yb->host[y*1024 + x] ==
59 hdiv(hdiv(inpb->host[y*1026 + x + 1027] + inpb->host[y*1026 + x + 1028] + inpb->
60 host[y*1026 + x + 1026], 3) + (hdiv(inpb->host[y*1026 + x + 2053] + inpb->host[y*1
61 026 + x + 2054] + inpb->host[y*1026 + x + 2052], 3) + hdiv(inpb->host[y*1026 + x +
62 1] + inpb->host[y*1026 + x + 2] + inpb->host[y*1026 + x], 3)), 3));
63 /*/
64 int blur_3(struct buffer *inpb, struct buffer *blur_yb) {

```

Listing 8. Back-end C code with annotations provided by HALIVER for the blur example of Listing 1 (2/3).

```

56 int32_t* _blur_y = blur_yb→host;
57 int32_t* _inp = inpb→host;
58 // produce blur_y
59 #pragma omp parallel for
60 for (int yo = 0; yo<0 + 128; yo++)
61 /*@
62 context 0≤yo ∧ yo<0 + 128;
63 context (∀* int x, int y; 0≤x ∧ x<1026 ∧ 0≤y ∧ y<1026; Perm(&_inp[y*1026 + x], 1
  \ (2*128)));
64 context (∀ int x, int y; 0≤x ∧ x<1026 ∧ 0≤y ∧ y<1026; _inp[y*1026 + x] ≡ p_i(y*1
  026 + x));
65 context (∀* int xif, int xof, int yif; (((0≤yif ∧ yif<8) ∧ 0≤xof) ∧ xof<512) ∧ 0
  ≤xif) ∧ xif<2; Perm(&_blur_y[(yo*8 + yif)*1024 + xof*2 + xif], 1\1));
66 ensures (∀ int yif, int xof, int xif; (((0≤yif ∧ yif<8) ∧ 0≤xof) ∧ xof<512) ∧ 0
  ≤xif) ∧ xif<2; _blur_y[(yo*8 + yif)*1024 + xof*2 + xif] ≡ hdiv(hdiv(p_i((yo*
  8 + yif)*1026 + xof*2 + xif) + (p_i((yo*8 + yif)*1026 + xof*2 + xif + 1) + p_i((
  yo*8 + yif)*1026 + xof*2 + xif + 2)), 3) + (hdiv(p_i((yo*8 + yif)*1026 + xof*2 +
  xif + 1026) + (p_i((yo*8 + yif)*1026 + xof*2 + xif + 1027) + p_i((yo*8 + yif)*1
  026 + xof*2 + xif + 1028)), 3) + hdiv(p_i((yo*8 + yif)*1026 + xof*2 + xif + 2052
  ) + (p_i((yo*8 + yif)*1026 + xof*2 + xif + 2053) + p_i((yo*8 + yif)*1026 + xof*2
  + xif + 2054)), 3)), 3));
67 @*/
68 {
69 {
70 int64_t _2 = 10240;
71 int32_t *_blur_x = (int32_t *)malloc(sizeof(int32_t)*_2);
72 int32_t _t11 = (yo * 8);
73 /*@
74 loop_invariant 0≤yi ∧ yi≤0 + 8;
75 loop_invariant (∀* int x, int y; 0≤x ∧ x<1026 ∧ 0≤y ∧ y<1026; Perm(&_inp[y*1026
  + x], 1\ (2*128)));
76 loop_invariant (∀ int x, int y; 0≤x ∧ x<1026 ∧ 0≤y ∧ y<1026; _inp[y*1026 + x] ≡
  p_i(y*1026 + x));
77 loop_invariant (∀* int x, int y; 0≤x ∧ x<1024 ∧ yo*8≤y ∧ y<yo*8 + 10; Perm(&
  _blur_x[(y - yo*8)*1024 + x], 1\1));
78 loop_invariant (∀* int xif, int xof, int yif; 0≤yif ∧ yif<8 ∧ 0≤xof ∧ xof<512 ∧
  0≤xif ∧ xif<2; Perm(&_blur_y[(yo*8 + yif)*1024 + xof*2 + xif], 1\1));
79 loop_invariant (∀ int yif, int xof, int xif; 0≤yif ∧ yif<8 ∧ 0≤xof ∧ xof<512 ∧
  0≤xif ∧ xif<2; _blur_y[(yo*8 + yif)*1024 + xof*2 + xif] ≡ hdiv(hdiv(p_i((
  yo*8 + yif)*1026 + xof*2 + xif) + (p_i((yo*8 + yif)*1026 + xof*2 + xif + 1) +
  p_i((yo*8 + yif)*1026 + xof*2 + xif + 2)), 3) + (hdiv(p_i((yo*8 + yif)*1026 +
  xof*2 + xif + 1026) + p_i((yo*8 + yif)*1026 + xof*2 + xif + 1027) + p_i((yo*8
  + yif)*1026 + xof*2 + xif + 1028)), 3) + hdiv(p_i((yo*8 + yif)*1026 + xof*2 +
  xif + 2052) + p_i((yo*8 + yif)*1026 + xof*2 + xif + 2053) + p_i((yo*8 + yif)*1
  026 + xof*2 + xif + 2054)), 3)), 3));
80 @*/
81 for (int yi = 0; yi<0 + 8; yi++)
82 {
83 // produce blur_x
84 int32_t _t12 = (yi + _t11);
85 /*@
86 loop_invariant _t12≤y ∧ y≤_t12 + 3;
87 loop_invariant (∀* int x, int y; 0≤x ∧ x<1026 ∧ 0≤y ∧ y<1026; Perm(&_inp[y*102
  6 + x], 1\ (2*128)));
88 loop_invariant (∀ int x, int y; 0≤x ∧ x<1026 ∧ 0≤y ∧ y<1026; _inp[y*1026 + x]
  ≡ p_i(y*1026 + x));
89 loop_invariant (∀* int xif, int xof, int yf; yo*8 + yi≤yf ∧ yf<yo*8 + yi + 3 ∧ 0
  ≤xof ∧ xof<512 ∧ 0≤xif ∧ xif<2; Perm(&_blur_x[(yf - yo*8)*1024 + xof*2 +
  xif], 1\1));
90 loop_invariant (∀ int yf, int xof, int xif; yo*8 + yi≤yf ∧ yf<y ∧ 0≤xof ∧ xof<
  512 ∧ 0≤xif ∧ xif<2; _blur_x[(yf - yo*8)*1024 + xof*2 + xif] ≡ hdiv(p_i((
  yf*513 + xof)*2 + xif) + p_i((yf*513 + xof)*2 + xif + 1) + p_i((yf*513 + xof)
  *2 + xif + 2), 3));
91 @*/
92 for (int y = _t12; y<_t12 + 3; y++)
93 {

```

Listing 9. Back-end C code with annotations provided by HALIVER for the blur example of Listing 1 (3/3).

```

94     int32_t _t14 = (y - _t11) * 512;
95     int32_t _t13 = (y * 513);
96     /*@
97     loop_invariant 0 ≤ xo ∧ xo ≤ 0 + 512;
98     loop_invariant (∀* int x, int y; 0 ≤ x ∧ x < 1026 ∧ 0 ≤ y ∧ y < 1026; Perm(&_inp[y*102
99     6 + x], 1\2*128));
100    loop_invariant (∀ int x, int y; 0 ≤ x ∧ x < 1026 ∧ 0 ≤ y ∧ y < 1026; _inp[y*1026 + x]
101    ≡ p_i(y*1026 + x));
102    loop_invariant (∀* int xif, int xof; 0 ≤ xof ∧ xof < 512 ∧ 0 ≤ xif ∧ xif < 2; Perm(&
103    _blur_x[(y - yo*8)*1024 + xof*2 + xif], 1\1));
104    loop_invariant (∀ int xof, int xif; 0 ≤ xof ∧ xof < xo ∧ 0 ≤ xif ∧ xif < 2; _blur_x[(y
105    - yo*8)*1024 + xof*2 + xif] ≡ hdiv(p_i(y*1026 + xof*2 + xif) + p_i(y*1026 +
106    xof*2 + xif + 1) + p_i(y*1026 + xof*2 + xif + 2), 3));
107    */
108    for (int xo = 0; xo < 0 + 512; xo++)
109    {
110        int32_t _t7 = xo + _t13;
111        _blur_x[(xo + _t14) * 2] = div_eucl(_inp[_t7 * 2] + _inp[_t7 * 2 + 1] + _inp[_t7
112        * 2 + 2], 3);
113        _blur_x[((xo + _t14) * 2) + 1] = div_eucl(_inp[_t7 * 2 + 1] + _inp[_t7 * 2 + 2]
114        + _inp[_t7 * 2 + 3], 3);
115    } // for xo
116    } // for y
117    // consume blur_x
118    int32_t _t16 = (yi + _t11) * 512;
119    int32_t _t15 = yi * 512;
120    /*@
121    loop_invariant 0 ≤ xo ∧ xo ≤ 0 + 512;
122    loop_invariant (∀* int x, int y; 0 ≤ x ∧ x < 1024 ∧ yo*8 ≤ y ∧ y < yo*8 + 10; Perm(&
123    _blur_x[(y - yo*8)*1024 + x], 1\2));
124    loop_invariant (∀ int xo, int y; 0 ≤ xo ∧ xo < 1024 ∧ yo*8 + yi ≤ y ∧ y ≤ yo*8 + yi +
125    2; _blur_x[(y - yo*8)*1024 + xo] ≡ hdiv(p_i(y*1026 + xo) + p_i(y*1026 + xo +
126    1) + p_i(y*1026 + xo + 2), 3));
127    loop_invariant (∀* int xif, int xof; 0 ≤ xof ∧ xof < 512 ∧ 0 ≤ xif ∧ xif < 2; Perm(&
128    _blur_y[(yo*8 + yi)*1024 + xof*2 + xif], 1\1));
129    loop_invariant (∀ int xof, int xif; 0 ≤ xof ∧ xof < xo ∧ 0 ≤ xif ∧ xif < 2; _blur_y[(
130    yo*8 + yi)*1024 + xof*2 + xif] ≡ hdiv(hdiv(p_i((yo*8 + yi)*1026 + xof*2 + xif
131    ) + p_i((yo*8 + yi)*1026 + xof*2 + xif + 1) + p_i((yo*8 + yi)*1026 + xof*2 +
132    xif + 2), 3) + hdiv(p_i((yo*8 + yi)*1026 + xof*2 + xif + 1026) + p_i((yo*8 +
133    yi)*1026 + xof*2 + xif + 1027) + p_i((yo*8 + yi)*1026 + xof*2 + xif + 1028),
134    3) + hdiv(p_i((yo*8 + yi)*1026 + xof*2 + xif + 2052) + p_i((yo*8 + yi)*1026 +
135    xof*2 + xif + 2053) + p_i((yo*8 + yi)*1026 + xof*2 + xif + 2054), 3), 3));
136    */
137    for (int xo = 0; xo < 0 + 512; xo++)
138    {
139        int32_t _t9 = (xo + _t15);
140        _blur_y[(xo + _t16) * 2] = div_eucl(_blur_x[_t9 * 2] + _blur_x[_t9 * 2 + 1024] +
141        _blur_x[_t9 * 2 + 2048], 3);
142        _blur_y[(xo + _t16) * 2 + 1] = div_eucl(_blur_x[_t9 * 2 + 1] + _blur_x[_t9 * 2 + 10
143        25] + _blur_x[_t9 * 2 + 2049], 3);
144    } // for xo
145    } // for yi
146    free(_blur_x);
147    } // alloc _blur_x
148    } // for yo
149    return 0;
150    }

```