



HipaccVX: wedding of OpenVX and DSL-based code generation

M. Akif Özkan¹ · Burak Ok¹ · Bo Qiao¹ · Jürgen Teich¹ · Frank Hannig¹

Received: 4 February 2020 / Accepted: 27 August 2020 / Published online: 22 September 2020
© The Author(s) 2020

Abstract

Writing programs for heterogeneous platforms optimized for high performance is hard since this requires the code to be tuned at a low level with architecture-specific optimizations that are most times based on fundamentally differing programming paradigms and languages. OpenVX promises to solve this issue for computer vision applications with a royalty-free industry standard that is based on a graph-execution model. Yet, the OpenVX ' algorithm space is constrained to a small set of vision functions. This hinders accelerating computations that are not included in the standard. In this paper, we analyze OpenVX vision functions to find an orthogonal set of computational abstractions. Based on these abstractions, we couple an existing domain-specific language (DSL) back end to the OpenVX environment and provide language constructs to the programmer for the definition of user-defined nodes. In this way, we enable optimizations that are not possible to detect with OpenVX graph implementations using the standard computer vision functions. These optimizations can double the throughput on an Nvidia GTX GPU and decrease the resource usage of a Xilinx Zynq FPGA by 50% for our benchmarks. Finally, we show that our proposed compiler framework, called HipaccVX, can achieve better results than the state-of-the-art approaches Nvidia VisionWorks and Halide-HLS.

Keywords OpenVX · Domain-specific language · Image processing · GPU · FPGA

1 Introduction

The emergence of cheap, low-power cameras and embedded platforms has boosted the use of smart systems with Computer Vision (CV) capabilities in a broad spectrum of markets, ranging from consumer electronics, such as mobile, to real-time automotive applications and industrial automation, e.g., semiconductors, pharmaceuticals, packaging. The global machine vision market size was valued at \$16.0

billion already in 2018, and yet is expected to reach a value of \$24.8 billion by 2023 [2]. A CV application might be implemented on a great variety of hardware architectures ranging from Graphics Processing Units (GPUs) to Field Programmable Gate Arrays (FPGAs) depending on the domain and the associated constraints (e.g., performance, power, energy, and cost). Yet, for sophisticated real-life applications, the best trade-off is often achieved by heterogeneous systems incorporating different computing components that are specialized for particular tasks.

Optimizing CV programs to achieve high performance on such heterogeneous systems usually goes along with sacrificing readability, portability, and modularity. The programs need to be tuned at a low level with architecture-specific optimizations that are typically based on drastically different programming paradigms and languages (e.g., parallel programming of multicore processors using C++ combined with OpenMP; vector data types, libraries, or intrinsics to utilize the SIMD¹ units of CPU; CUDA or OpenCL for programming GPU accelerators; hardware description

✉ M. Akif Özkan
akif.oezkan@fau.de

Burak Ok
ok.burak@web.de

Bo Qiao
bo.qiao@fau.de

Jürgen Teich
teich@fau.de

Frank Hannig
hannig@fau.de

¹ Hardware/Software Co-Design, Department of Computer Science, Friedrich-Alexander University Erlangen-Nürnberg (FAU), Erlangen, Germany

¹ Single Instruction, Multiple Data (SIMD) units are CPU components for vector processing, i.e., they execute the same operation on multiple data elements in parallel.

Table 1 Available features in OpenVX (VX), DSL compiler Hipacc (H), and our joint approach HipaccVX (HVX)

Features	VX	H	HVX
Industrial standard (open, royalty free)	✓	✗	✓
Community-driven open-source implementations	✗	✓	✓
Well-known CV functions (e.g., optical flow)	✓	✗	✓
High-level abstractions that adhere to distinct memory access patterns (e.g., local)	✗	✓	✓
Custom node execution on accelerator devices (i.e., OpenCL)	✓	✗	✓
Acceleration of the custom nodes that are based on high-level abstractions	✗	✓	✓

languages such as Verilog or VHDL for targeting FPGAs). Partitioning a program across different computing units, and accordingly, synchronizing the execution is difficult. To achieve these ambitious goals, high development effort and architecture expert knowledge are required.

In 2014, the Khronos Group released OpenVX as a C-based API to facilitate cross-platform portability not only of the code but also of the performance for CV applications [27]. This is momentous since OpenVX is the first (royalty free) standard for a graph-based specification of CV algorithms. Yet, the OpenVX ' algorithm space is constrained to a relatively small set of vision functions. Users are allowed to instantiate additional code in the form of custom nodes, but these cannot be analyzed at the system-level by the graph-based optimizations applied from an OpenVX back end. Additionally, this requires users to optimize their implementations, who supposedly should not consider the optimizations of the performance. Standard programming languages such as OpenCL do not offer performance portability across different computing platforms [4, 22]. Therefore, the user code, even optimized for one specific device, might not provide the expected high performance when compiled for another target device. These deficiencies are listed in Table 1.

A solution to the problems mentioned above is offered by the community working on Domain-Specific Languages (DSLs) for image processing. Recent works show that excellent results can be achieved when high-level image processing abstractions are specialized to a target device via modern metaprogramming, compiler, or code generation approaches [8, 10, 16]. These DSLs are able to generate code from a set of *algorithmic abstractions* that lead to high-performance execution for diverse types of computing platforms. However, existing DSLs lack formal verification; hence, they do not ensure the safe execution of a user application whereas OpenVX is an industrial standard.

In this paper, we couple the advantages of DSL-based code generation with OpenVX (summarized in Table 1). We present a set of abstractions that are used as basic building blocks for expressing OpenVX ' standard CV functions. These building blocks are suitable for generating optimized, device-specific code from the same functional description, and are systematically utilized for graph-based

optimizations. In this way, we achieve performance portability not only for OpenVX ' CV functions but also for user-defined kernels² that are expressed with these computational abstractions. The contributions of this paper are summarized as follows:

- We systematically categorize and specify OpenVX ' CV functions by high-level abstractions that adhere to distinct memory access patterns (see Sect. 4).
- We propose a framework called HipaccVX, which is an OpenVX implementation that achieves high performance for a wide variety of target platforms, namely GPUs, CPUs, and FPGAs (see Sect. 5).
- HipaccVX³ supports the definition of custom nodes (i.e., user-defined kernels) based on the proposed abstractions (see Sect. 5.1).
- To the best of our knowledge, our approach is the first one that allows for graph-based optimizations that incorporate not only standard OpenVX CV nodes but also user-defined custom nodes (see Sect. 5.2), i.e., optimizations across standard and custom nodes.

2 Related work

The OpenVX specification is not constrained to a certain memory model as OpenCL and OpenMP, therefore enables better performance portability than traditional libraries such as OpenCV [17]. It has been implemented by a few major vendors, including Nvidia, Intel, AMD, and Synopsys [28]. The authors of [5, 9, 25, 31, 32] focus on graph scheduling and design space exploration for heterogeneous systems consisting of GPUs, CPUs, and custom instruction-set architectures. Unlike the prior work, [24] suggests static OpenVX compilation for low-power embedded systems instead of runtime-library implementations. Our work is similar to this since we statically analyze a given OpenVX application and combine the benefits of domain-specific code generation approaches [3, 8, 10, 14, 16, 19].

² A kernel in OpenVX is the abstract representation of a computer vision function [30].

³ <https://github.com/akifoezkan/HipaccVX>




Fig. 1 Graph representation for the OpenVX code given in Listing 1. The output image (*img1*) contains solely the horizontal edges extracted from the input image (*img0*). The *virt2* image is defined only because OpenVX' Sobel function returns both horizontal and

vertical edges. This redundant computation is eliminated during the optimization passes of our HipaccVX compiler framework (see Sect. 5.2.2)

Halide [16], Hipacc [8], and PolyMage [10] are image processing DSLs that provide language constructs and scheduling primitives to generate code that is optimized for the target device, i.e., CPUs, GPUs. Halide [16] decouples the algorithm description from scheduling primitives, i.e., vectorization, tiling, while Hipacc [8] and PolyMage [10] implicitly apply these optimizations on a graph-based description similar to OpenVX. CAPH [20], RIPL [23], and Rigel [6] are image processing DSLs that generate optimized code for FPGAs. Hipacc-FPGA [19] supports HLS tools of both Xilinx and Intel, while Halide-HLS [14], PolyMage-HLS [3], and RIPL only target Xilinx devices. CAPH relies upon the actor/dataflow model of computation to generate VHDL or SystemC code. Our approach could also be used to implement OpenVX by these image processing DSLs.

There is no publicly available OpenVX implementation for Xilinx FPGAs to the best of our knowledge. Intel OpenVino [7] provides a few example applications that are specific to Arria-10 FPGAs. Taheri et al. [26] provide some initial results for FPGAs, where the main attention is the scheduling of statistical kernels (i.e., histogram). The image processing DSLs in [3, 19] use similar techniques to implement user applications as a streaming pipeline. Section 5.2.1 shows how to instrument these techniques for the OpenVX API. Omidian et al. [12] present a heuristic algorithm for the design space exploration of OpenVX graphs for FPGAs. This algorithm could be simplified by using HipaccVX' abstractions (see Sect. 4) instead of OpenVX' CV functions. Then it could be used in conjunction with HipaccVX to explore the design space of hardware/software platforms. Moreover, Omidian et al. [11] suggest an overlay architecture for FPGA implementations of OpenVX. The proposed overlay implementation requires the optimized implementation of OpenVX' CV functions, which could be generated by HipaccVX. Furthermore, an overlay architecture based on HipaccVX's abstractions, which is a smaller set of functions compared to OpenVX CV functions, could reduce resource usage in [11].

Intel's OpenVX implementation [1] is the first work extending the OpenVX standard with an interoperability API for OpenCL. This is supported in OpenVX v1.3 [30]. Yet, performance portability still cannot be assured for the custom nodes. An OpenCL code tuned for a specific CPU might perform very poorly on FPGAs and GPU architectures [4,

22]. Contrary to our approach, the performance of this approach relies on the user code.

3 OpenVX and image processing DSLs

In the following Sects. 3.1 and 3.2, we briefly explain the programming models of OpenVX and image processing DSLs, respectively. Then we discuss the complementary features of these approaches in Sect. 3.3, which are the motivation of this work.

3.1 OpenVX programming model

OpenVX is an open, royalty-free C-based standard for the cross-platform acceleration of computer vision applications. The specification does not mandate any optimizations or requirements on device execution; instead, it concentrates on software abstractions that are freed from low-level platform-specific declarations. The OpenVX API is totally opaque; that is, the memory hierarchy and device synchronization are hidden from the user. Typically, platform experts of the

```

1 vx_context context = vxCreateContext();
2 vx_graph graph = vxCreateGraph(context);
3
4 vx_image img[] = {
5     vxCreateImage(context, width, height, VX_DF_IMAGE_UYVY),
6     vxCreateImage(context, width, height, VX_DF_IMAGE_U8)};
7
8 vx_image virt[] = {
9     vxCreateVirtualImage(graph, 0, 0, VX_DF_IMAGE_VIRT),
10    vxCreateVirtualImage(graph, 0, 0, VX_DF_IMAGE_VIRT),
11    vxCreateVirtualImage(graph, 0, 0, VX_DF_IMAGE_VIRT),
12    vxCreateVirtualImage(graph, 0, 0, VX_DF_IMAGE_VIRT)};
13
14 vxChannelExtractNode(graph, img[0], VX_CHANNEL_Y, virt[0]);
15 vxGaussian3x3Node(graph, virt[0], virt[1]);
16 vxSobel3x3Node(graph, virt[1], virt[2], virt[3]);
17 vxMagnitudeNode(graph, virt[3], virt[3], virt[4]);
18 vxThresholdNode(graph, virt[4], thresh, img[1]);
19
20 status = vxVerifyGraph(graph);
21 if (status == VX_SUCCESS)
22     status = vxProcessGraph(graph);
23
24 vxReleaseContext(&context);
  
```

Listing 1: OpenVX code for an edge detection algorithm. The application graph derived for this OpenVX program is shown in Fig. 1

individual hardware vendors provide optimized implementations of the OpenVX API [28].

Listing 1 shows an example OpenVX code for a simple edge detection algorithm, for which the application graph is shown in Fig. 1. An application is described as a Directed Acyclic Graph (DAG), where nodes represent CV functions (see Lines 14–18) and data objects, i.e., images, scalars (see Lines 4–12), while edges show the dependencies between nodes. All OpenVX *objects* (i.e., *graph*, *node*, *image*) exist within a *context* (Line 1). A *context* keeps track of the allocated memory resources and promotes implicit freeing mechanisms at release calls (Line 24). A *graph* (Line 2) solely operates on the data objects attached to the same context.

The data objects that are used only for the intermediate steps of a calculation, which can be inaccessible for the rest of the application, should be specified as *virtual* by the users. Virtual data objects (i.e., *virtual images* defined in Lines 9–12) cannot be accessed via read/write operations. This paves the way for system-level optimizations applied in a platform-specific back end, i.e., host-device data transfers or memory allocations [17].

The execution is not eager; an OpenVX graph must be *verified* (Line 20) before it is *executed* (Line 22). The verification ensures the safe execution of the graph and resolves the implementation types of virtual data objects. The OpenVX standard mandates that a verification procedure must (i) validate the node parameters (i.e., presence, directions, data types, range checks), and (ii) assure the graph connectivity (detection of cycles), at the minimum [29]. Optimizations of an OpenVX back end should be performed during the verification phase. The verification is considered to be an initialization procedure and might restructure the application graph before the execution. A verified graph can be executed repeatedly for different input parameters (i.e., a new frame in video processing).

3.1.1 Deficiencies of OpenVX

As mentioned above, the OpenVX standard relieves an application programmer from low-level, implementation-specific descriptions, and thus enables portability across a variety of computing platforms. In OpenVX, the smallest component to express a computation is a graph node (e.g., `vxGaussian3x3Node`) from the set of base CV functions. However, these CV functions are restricted to a small set since OpenVX has a tight focus on cross-platform acceleration [30]. Custom nodes can be added to extend this

⁴ The support for the execution of a user code (custom node) as part of an application graph on an accelerator device was introduced only recently (August 2019) with the release of OpenVX v1.3 [30]. Previous versions [29] constraint the usage of the user-defined kernels to the host platform and required them to be implemented as C++ kernels.




Fig. 2 HipaccVX enables performance portability for user-defined code by representing OpenVX ' CV functions as well as custom nodes by a small set of computational abstractions

functionality⁴, but, they leave the following issues unresolved: (i) Users are responsible for the performance of a custom node, who supposedly should not consider performance optimizations. (ii) Portability of performance cannot be enabled for the cross-platform acceleration of user code. (iii) The graph optimization routines cannot analyze custom nodes.

For instance, consider Fig. 2 that depicts an OpenVX application graph with three CV function nodes (red) and a user-defined kernel node (blue). A GPU back end would offer optimized implementations of the `vxNodes` (e.g., Gauss), but the user code (custom node) is a black box for the graph optimizations.

Programming models such as OpenCL can be used to implement custom nodes. This enables functional portability across a great variety of computing platforms. However, the user should have *expertise* in the target architecture in order to optimize an implementation for high performance. Furthermore, OpenCL cannot assure the portability of

```

1 // filter mask for Gaussian blur filter
2 const float filter_mask[3][3] = {
3   { 0.057118f, 0.124758f, 0.057118f },
4   { 0.124758f, 0.272496f, 0.124758f },
5   { 0.057118f, 0.124758f, 0.057118f }
6 };
7 Mask<float> mask(filter_mask);
8
9 // input and output images
10 size_t width, height;
11 uchar *image = read_image(&width, &height, "input.pgm");
12 Image<uchar> in(width, height, image);
13 Image<uchar> out(width, height);
14
15 // reading from in with clamping as boundary condition
16 BoundaryCondition<uchar> cond(in, mask, Boundary::CLAMP);
17 Accessor<uchar> acc(cond);
18
19 // output image (region of interest is the whole image)
20 IterationSpace<uchar> iter(out);
21
22 // instantiate and launch the Gaussian blur filter
23 LinearFilter Gaussian(iter, acc, mask, 3);
24 Gaussian.execute();

```

Listing 2: Hipacc application code for a Gaussian filter. It instantiates the **LinearFilter** Kernel given in Listing 3

the performance since the code needs to be tuned according to the target device, i.e., usage of device-specific synchronization primitives, exploitation of texture memory if available, usage of vector operations, or different numbers of hardware threads [4, 22]. In fact, an OpenCL code optimized for an Instruction Set Architecture (ISA) has to be ultimately rewritten for an FPGA implementation to deliver high-performance [13].

3.2 Image processing DSLs

Recently proposed DSL compilers for image processing, such as Halide [16], Hipacc [8], and PolyMage [10], enable the portability of high-performance across varying computing platforms. All of them take as input a high-level, functional description of the algorithm and generate platform-specific code tuned for the target device. In this work, we use Hipacc to present our approach.

```

1 class LinearFilter: public Kernel <uchar> {
2   // ...
3   public:
4     LinearFilter(Accessor<uchar> &input, // input image
5                 IterationSpace<uchar> &out, // output image
6                 Mask<float> &mask)        // mask
7       : { /* ... */ }
8
9   void kernel() { // convolve -> local operator
10    output() = convolve(mask, Reduce::SUM, [&] () -> uchar {
11      return mask() * input(mask);
12    });
13  }
14 };

```

Listing 3: Hipacc kernel code for an FIR filter

Hipacc provides language constructs that are embedded into C++ for the concise description of computations. Applications are defined in a Single Program, Multiple Data (SPMD) context, similar to kernels in CUDA and OpenCL. For instance, Listing 2 shows the description of a discrete Gaussian blur filter application. First, a **Mask** is defined in Line 7 from a constant array. Then, input and output **Images** are defined as C++ objects in Line 12 and 13, respectively. *Clamping* is selected as the image boundary handling mode for the input image in Line 16. The whole input and output images are defined as Region of Interest (ROI) by the **Accessor** and **IterationSpace** objects that are specified in Line 17 and 20, respectively. Finally, the Gaussian kernel is instantiated in Line 23 and executed in Line 24.

Listing 3 describes the actual *operator kernel* for the Gaussian shown in Listing 2. The **LinearFilter** is a user-defined class that is derived from Hipacc's **Kernel** class, where the `kernel` method is overridden. There, a user describes a convolution as a lambda function using the **convolve()** construct, which computes an output pixel (**output()**) from an input window (**input(mask)**). Hipacc's compiler utilizes Clang's Abstract Syntax Tree (AST) to specialize the lambda function according to the selected platform and generates device-specific code that provides high-performance implementations when compiled with the target architecture compiler. We refer to [8, 19] for more detailed explanations, further programming language constructs of Hipacc as well as corresponding code generation techniques.

3.3 Combining OpenVX with image processing DSLs

Our solution to the posed challenges in Sect. 3.1.1 is introducing an orthogonal set of so-called *computational abstractions* that enables high-performance implementations for

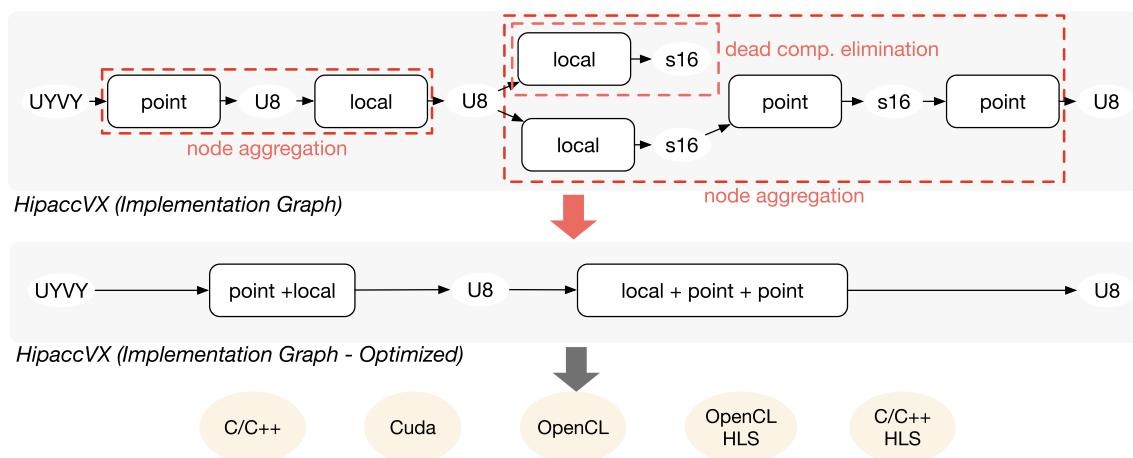


Fig. 3 The application graph in Fig. 1 is implemented using high-level abstractions called *point* and *local* (explained Sect. 4) instead of OpenVX vision function. This enables high-performance code gener-

ation for various targets when coupled with a DSL compiler and additional optimizations such as dead computation elimination and node aggregation (see Sects. 5.1.1 and 5.2)

a variety of computing platforms (such as CPUs, GPUs, FPGAs), similar to the DSLs discussed in Sect. 3.2. These abstractions should be used to implement OpenVX ' CV functions and, at the same time, be served to the user for the definition of custom nodes.

Assume that the geometric shapes in Fig. 2 represent the abstractions above. By implementing both the OpenVX CV functions and the custom node using the basic building block (different geometric shapes in the figure), a consistent graph is constructed for the implementation. Consequently, the problem of instantiating the user code as a black box is eliminated. Likewise, assume that all the CV functions of the OpenVX code in Listing 1 are implemented by using the computational abstractions called *point* and *local* (explained in Sect. 4). Then its application graph (Fig. 1) transforms into the implementation graph shown in Fig. 3. This implementation graph could be used for target-specific optimizations and code generation similar to the DSL compiler approaches for image processing.

In this paper, we implement the OpenVX standard by the computational abstractions explained in Sect. 4. We accomplish this task by developing a back end for OpenVX using Hipacc (as an existing image processing DSL) instead of standard programming languages. In this way, we get the best of both worlds (OpenVX and DSL works). Our approach relies on OpenVX ' industry-standard graph specification and enables DSL-based code generation. The user is offered well-known CV functions as well as DSL elements (i.e., programming constructs, abstractions) for the description of custom nodes. As a result of this, programmers can write functional descriptions for custom nodes without having concerns about the performance; and, as a consequence, allows writing performance-portable OpenVX programs for a larger algorithm space.

4 Computational abstractions

We have analyzed OpenVX ' CV functions and categorized them into the computational abstractions summarized in Table 2. The categorization is mainly based on three groups of operators: (i) *point operators* that compute an output from one input pixel, (ii) *local operators* depend on neighbor pixels over a certain region, and (iii) *global operators* where the

output might depend on the whole input image, (presented in Fig. 4). We have identified the following patterns for the global operators: (a) *reduction*: traverses an input image to compute one output (e.g., max, mean), (b) *histogram*: categorizes (maps) input pixels to bins according to a binning (reduce) function, (c) *scaling*: downsizes or expands input images by interpolation, (d) *scan*: each output pixel depends on the previous output pixel. *Warp*, *transpose*, and *matrix multiplication* are denoted as global operator blocks.

Through the introduction of the node-internal computational abstractions, our approach enables additional optimizations that manipulate the computation (see Sect. 5.2 and 5.1.1). This is also illustrated in Fig. 3, where redundant computations are eliminated, and nodes are aggregated for better exploitation of locality. Memory access patterns of our abstractions entail system-level optimization strategies motivated by the OpenVX standard, such as image tiling [25] and hardware-software partitioning [26]. An abstraction-based implementation allows expressing aggregated computations as part of the reconstructed graph. In this way, an implementation graph, as well as an application graph can be expressed using the same graph structure. Furthermore, using the proposed set of abstractions reduces code duplication compared to typical approaches, where the libraries are implemented using hand-written CV functions. For instance, 36 of OpenVX ' CV functions can be implemented solely with the description of point and local operators as shown in Table 2; that is, a few highly optimized building blocks for a single target platform (e.g., GPU) can be reused.

5 The HipaccVX framework

In this paper, we developed a framework, called HipaccVX, which is a DSL-based implementation of OpenVX. We extended OpenVX specification by Hipacc code interoperability (see Sect. 5.1) such that programmers are allowed to register Hipacc kernels as custom nodes to OpenVX programs. The HipaccVX framework consists of an OpenVX graph implementation and optimization routines that verify and optimize input OpenVX applications (see Sect. 5.2). Ultimately, it generates a device-specific code for the target platform using Hipacc ' s code generation. The tool flow is presented in Fig. 5.




Fig. 4 The considered computational abstractions (listed in Table 2) are based on three groups of operators

Table 2 Categorization of the OpenVX Kernels according to data access patterns

OpenVX Kernels	HipaccVX Abstractions	hipacc abstractions
AbsDiff, Copy, Add, Subtract, And, Xor, Or, Not, ChannelCombine, ChannelExtract, ColorConvert, ConvertDepth, Magnitude, Phase, Multiply, ScaleImage, Threshold, TensorAdd, TensorSubtract, TensorConvertDepth, TensorMultiply, ScalarOperation, Select, Remap	Point	Kernel
NonMaxSuppression, Dilate3x3, Erode3x3, NonLinearFilter, Median3x3, BilateralFilter, Sobel3x3, Box3x3, Convolve, Gaussian3x3, LBP, FastCorners	Local	Kernel
MinMaxLoc, MeanStdDev, Min, Max	Reduce (global)	Reduction
Histogram	Histogram (global)	Histogram
Scale-image	Scale (global)	Interpolation
GaussianPyramid, LaplacianPyramid, LaplacianReconstruct	Pyramid (global)	Pyramid
IntegralImage	Scan (global)	Software
WarpAffine, WarpPerspective	Warp (global)	Software
TensorTranspose, TensorMatrixMultiply	(Global) transpose, matrixMult	Software
HarrisCorners	Point + local + <i>custom</i>	Kernel, Software
EqualizeHist	Histogram + point	Kernel, Histogram
OpticalFlowPyrLK	Point + local + pyramid + <i>custom</i>	Kernel, Pyramid, Software
HOGCells	<i>Custom</i> + local + histogram	Kernel, Software
CannyEdge	Point + local + <i>custom</i>	Kernel, Software

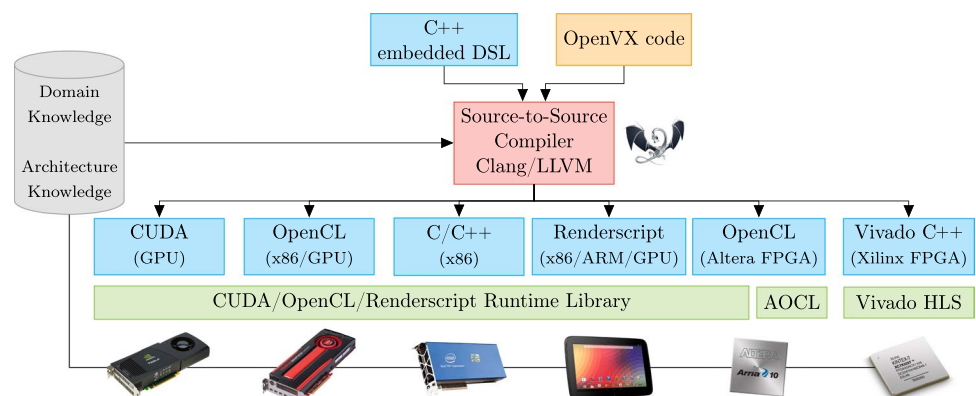
5.1 DSL back-end and user-defined kernels

OpenVX mandates the verification of parameters and the relationship between input and output and parameters as presented in Listing 4. There, first, a user kernel and all of its parameters should be defined (lines 6–26). Then a custom node should be created by `vxCreateGenericNode` (Line 30) after the user kernel is finalized by a `vxFinalizeKernel` call (Line 27). The kernel parameter types are defined, and the node parameters are set by `vxAddParameterToKernel` (lines 20–26) and `vxSetParameterByIndex` (lines 31–33), respectively.

We extended OpenVX by **`vxHipaccKernel`** function (Line 6) to instantiate a Hipacc kernel as an OpenVX kernel. The Hipacc kernels should be written in a separate file and added as a generic node according to the OpenVX

standard [30]. Programmers do not have to describe the dependency between Hipacc kernels as in Listing 2, instead, they write a regular OpenVX program to describe an application graph. This sustains the custom node definition procedure of OpenVX. Ultimately, the HipaccVX framework verifies and optimizes a given OpenVX application, generates the corresponding Hipacc code, and employs Hipacc for device-specific code generation.

OpenVX ' CV functions are implemented as a library using our extension for Hipacc code instantiation. For instance, the HipaccVX implementation of the **`vxGaussian3x3Node`** API is shown in Listing 4. Users can simply use these CV functions as in Listing 1. A minority of OpenVX functions are implemented as OpenCV kernels since they cannot be fully described in Hipacc. These are listed in Table 2 with a

Fig. 5 HipaccVX overview

```

1  vx_node vxGaussian3x3Node(vx_graph graph,
2                             vx_image arr,
3                             vx_image out) {
4
5  // Extension: An OpenVX kernel from a Hipacc kernel
6  vx_kernel cstmk = vxHipaccKernel("gaussian3x3.cpp");
7
8  /** The code below is the standard OpenVX API **/
9  // Create vx_matrix for mask
10 const float coeffs[3][3] = /* ... */;
11 vx_matrix mask = vxCreateMatrix(context,
12                                VX_TYPE_FLOAT32,
13                                3, 3);
14 vxCopyMatrix(mask, (void*)coeffs, VX_WRITE_ONLY,
15              VX_MEMORY_TYPE_HOST);
16
17 // Set input/output parameters for a kernel
18 vxAddParameterToKernel(cstmk, 0, VX_OUTPUT,
19                        VX_TYPE_IMAGE,
20                        VX_PARAMETER_STATE_REQUIRED);
21 vxAddParameterToKernel(cstmk, 1, VX_INPUT,
22                        VX_TYPE_IMAGE,
23                        VX_PARAMETER_STATE_REQUIRED);
24 vxAddParameterToKernel(cstmk, 2, VX_INPUT,
25                        VX_TYPE_MATRIX,
26                        VX_PARAMETER_STATE_REQUIRED);
27 vxFinalizeKernel(cstmk);
28
29 // Create generic node
30 vx_node node = vxCreateGenericNode(graph, cstmk);
31 vxSetParameterByIndex(node, 0, (vx_reference) out);
32 vxSetParameterByIndex(node, 1, (vx_reference) arr);
33 vxSetParameterByIndex(node, 2, (vx_reference) mask);
34
35 return node;
36 }

```

Listing 4: DSL code interoperability extension (onlyLine 6)

Software label instead of a Hipacc abstraction type. As future work, we can extend Hipacc to support these functions.

5.1.1 Optimizations based on code generation

We inherited many device-specific optimization techniques by implementing a Hipacc back end for OpenVX. Hipacc internally applies several optimizations for the code generation from its DSL abstractions. These include memory padding, constant propagation, utilization of textures, loop unrolling, kernel fusion, thread-coarsening, implicit use of unified CPU/GPU memory [8, 15, 18]. At the same time, Hipacc targets Intel and Xilinx FPGAs using their High-Level Synthesis (HLS) tools. There, an input application is implemented through application circuits derived from the DSL abstractions and optimized by hardware techniques such as pipelining and loop coarsening [19].

5.2 OpenVX graph and system-level optimizations

As mentioned before, an OpenVX application is represented by a DAG $G_{\text{app}} = (V, E)$, where V is a set of vertices, and E is a set of edges $E \subseteq V \times V$ denoting data dependencies

between nodes. The set of vertices V can further be divided into two disjoint sets D and N ($V = D \cup N$, $D \cap N = \emptyset$) denoting data objects and CV functions, respectively.

Both data (i.e., Image, Scalar, Array) and node (i.e., CV functions) objects are implemented as C++ classes that inherit the OpenVX `Object` class. Vertices $v \in V$ of our OpenVX graph implementation consist of OpenVX `Object` pointers. The verification phase first checks if an application graph G_{app} (derived from the user code, see, e.g., Listing 1) does not contain any cycles. Then it verifies that the description is a bipartite graph, i.e., $\forall (v, w) \in E : v \in D \wedge w \in N \vee v \in N \wedge w \in D$. Finally, the verification phase applies the following optimizations:

5.2.1 Reduction of data transfers

Data nodes of an application graph that are not virtual must be accessible to the host, while the intermediate (virtual) points of a computation should be stored in the device memory. We distinguish these two data node types by the set of non-virtual data nodes D_{nv} and the set of virtual data nodes D_v , where $D = D_{\text{nv}} \cup D_v$, $D_{\text{nv}} \cap D_v = \emptyset$. HipaccVX keeps this information in its graph implementation and determines the subgraphs between non-virtual data nodes, which can be kept in the device memory. In this way, data transfers between host and device are avoided.

5.2.2 Elimination of dead computations

An application graph may consist of nodes that do not affect the results. Inefficient user code or other compiler transformations might cause such dead code. A less apparent reason could be the usage of OpenVX compound CV functions for smaller tasks. Consider `Sobel13x3` as an example, which computes two images, one for the horizontal and one for the vertical derivative of a given image. As the OpenVX API does not offer these algorithms separately, programmers have to call `Sobel13x3`, even when they are only interested in one of the two resulting images. Our implementation is based on abstractions and allows a better analysis of the computation compared to OpenVX ' CV functions, i.e., the Sobel API is implemented by two parallel local operators as shown in Fig. 3. HipaccVX optimizes a given application graph using the procedure described in Algorithm 1. Conventional compilers do not analyze this redundancy if utilizing the host/device execution paradigm (e.g., OpenCL, CUDA); that means, when OpenVX kernels are offloaded to an accelerator device, and device kernels are executed by the host according to the application dependency (see Sect. 6.2).

Algorithm 1 assumes that the non-virtual data nodes whose input and output degrees are zero must be the inputs (D_{in}) and the results (D_{out}) of an application, respectively.

Other non-virtual data nodes could be input, output, or intermediate points of an application depending on the number of connected virtual data nodes. These are initialized in Line 2. Then, all of the nodes in the same component between the node v_{start} and the set V_{in} are traversed via the *depth-first visit* function (Line 18) and marked as alive (Lines 2–20). Finally, in Line 21, a filtered view of an application graph is created from the set of alive nodes.

The complexity of the functions *transpose* (Line 15) and *depth-first visit* (Line 18) are $\mathcal{O}(|V| + |E|)$ and $\mathcal{O}(|E|)$, correspondingly. The filter graph function (Line 21) is only an adaptor that requires no change in the application graph [21]. In the worst case, the graph has $|V| - 2$ output data nodes. That is, the complexity of Algorithm 1 becomes $\mathcal{O}(|V|^2 + |E|)$ in time and $\mathcal{O}(|V| + |E|)$ in space.

Algorithm 1: Graph Analysis for Dead Computation Elimination

```

input :  $G_{\text{app}}$  – application graph
          $D_{\text{nv}}$  – set of are non-virtual data nodes
output :  $G_{\text{filt}}$  – optimized application graph
1 function eliminate_death_nodes( $G_{\text{app}}, D_{\text{nv}}$ )
   /* Find candidate non-virtual roots and leaves */
2    $D_{\text{in}} \leftarrow \emptyset, D_{\text{out}} \leftarrow \emptyset$ 
3   forall  $v \in D_{\text{nv}}$  do
4     if  $\text{deg}^-(v) = 0$  then
5        $D_{\text{in}} \leftarrow D_{\text{in}} \cup v$  // input non-virtual data
        nodes
6     end
7     else if  $\text{deg}^+(v) = 0$  then
8        $D_{\text{out}} \leftarrow D_{\text{out}} \cup v$  // out non-virtual data
        nodes
9     end
10    else
11       $D_{\text{in}} \leftarrow D_{\text{in}} \cup v$ 
12       $D_{\text{out}} \leftarrow D_{\text{out}} \cup v$ 
13    end
14  end
   /* Mark the nodes between roots and leaves as alive */
15   $G_{\text{trans}} \leftarrow \text{transpose\_graph}(G_{\text{app}})$ 
16   $V_{\text{alive}} \leftarrow \emptyset$ 
17  forall  $v_{\text{start}} \in D_{\text{out}}$  do
18     $V_v \leftarrow \text{depth-first\_visit}(v_{\text{start}}, D_{\text{in}}, G_{\text{trans}})$ 
19     $V_{\text{alive}} \leftarrow V_{\text{alive}} \cup V_v$ 
20  end
   /* Filter, keep only the alive nodes and their edges */
21   $G_{\text{filt}} \leftarrow \text{filter\_graph}(G_{\text{app}}, \text{KEEP\_EDGES}, V_{\text{alive}})$ 
22  return  $G_{\text{filt}}$ 
23 end

```

6 Evaluation and results

We present results for a Xilinx Zynq ZYNQ-zc706 FPGA using Xilinx Vivado HLS 2019.1 and an Nvidia GeForce GTX 680 with CUDA driver 10.0. We evaluate the following applications: As image smoothers, we consider a Gaussian blur (*Gauss*) and a *Laplacian* filter with a 5×5 and 3×3 local node, respectively. The filter chain (*FChain*) is an image pre-processing algorithm consisting of three convolution (local) nodes. The *SobelX* determines the horizontal derivative of an input image using the OpenVX `vxSobel` function. The edge detector in Fig. 1 (*EdgFig2*) finds horizontal edges in an input image,

while *Sobel* computes both horizontal and vertical edges using three CV nodes. The *Unsharp* filter sharpens the edges of an input image using one Gauss node and three point operator nodes. Both *Harris* and *Tomasi* detect corners of a given image using 13 (4 local + 9 point) and 14 (4 local + 10 point) CV nodes, respectively. These applications are representative to show the optimization techniques discussed in this paper. The performance of a simple CV application (e.g., *Gauss*) solely depends on the quality of code generation, while graph-based optimizations can further optimize the performance of more complex applications (e.g., *Tomasi*). *Laplacian* uses the OpenVX ' custom convolution API and *EdgFig2* consists of redundant kernels.

6.1 Acceleration of user-defined nodes

User-defined nodes can be accelerated on a target platform (e.g., GPU accelerator) when they are expressed with HipaccVX ' abstractions (see Sect. 5.1). A C++ implementation of these custom nodes results in executing them on the host device. This is illustrated in Fig. 6 for a corner detection algorithm that consists of nine kernels. The CPU codes for these custom nodes are also acquired using hipacc. As can be seen in Fig. 6, HipaccVX provides the same performance invariant to the number of user-defined nodes, whereas using the OpenVX API decreases the throughput severely since each user-defined node has to be executed on the host CPU.




Fig. 6 Throughput for different versions of the same corner detection application (consisting of 9 kernels) on the Nvidia GTX680 (higher is better). The blue bars denote an increasing number of CV functions implemented as user-defined nodes using C++. In OpenVX, these user-defined functions have to be executed on the host CPU, which leads to a performance degradation; whereas, HipaccVX accelerates all user-defined nodes on the GPU




Fig. 7 Normalized execution time (lower is better) for 1024×1024 images. HipaccVX eliminates redundant transfers by analyzing OpenVX 'graph-based application code




Fig. 8 Normalized execution time (lower is better) for 1024×1024 images




Fig. 9 Post-Place and Route (PPnR) results for the Xilinx Zynq FPGA. Elimination of dead computation reduces the area, significantly

6.2 System-level optimizations based on OpenVX Graph

Reduction of data transfers HipaccVX eliminates the data transfers between the execution of subsequent functions on a target accelerator device, as explained in Sect. 5.2.1. This is disabled for *naïve* implementations. The improvements for the two applications are shown in Fig. 7. HipaccVX ' throughput optimizations reach a speedup of 13.5.

Elimination of dead computation HipaccVX eliminates the computations that do not affect the results of an application (see Sect. 5.2.2). This is illustrated in Fig. 8. HipaccVX improves the throughput by a factor of 2.1 on the GTX 680. The throughput improvement for the Zynq FPGA is only

slightly better since the applications fit into the target device; thus, run in parallel. Yet, HipaccVX ' FPGA implementation for the same application reduces the number of FPGA resources (elementary programmable logic blocks called *slices* and on-chip block RAMs, short *BRAMs*) significantly (around 50% for SobelX) on the Zynq (see Fig. 9).

6.3 Evaluation of the performance

In Fig. 10, we compare HipaccVX with the VisionWorks (v1.6) provided by Nvidia, which provides an optimized commercial implementation of OpenVX. HipaccVX, as well as typical library implementations, exploit the graph-based OpenVX API to apply system-level optimizations [17], such as reduction of data transfers (see Sect. 5.2). Additionally, HipaccVX generates code that is specific to target GPU architectures and applies optimizations such as constant propagation, thread coarsening, and multiple program multiple data (MPMD) [8]. As shown in Fig. 10, HipaccVX can generate implementations that provide higher throughput than VisionWorks. Here, the speedups for applications that are composed of multiple kernels (Harris, Tomasi, Sobel, Unsharp) are higher than the ones solely consisting of one OpenVX CV function (Gauss and Laplacian). This performance boost is, to a large extent, due to the locality optimization achieved by fusing consecutive kernels at the compiler level [15]. This requires code rewriting and the resource analysis of the target GPU architectures.

There was no publicly available FPGA implementation of OpenVX at the time this paper was written. Therefore, in Table 3, we compare HipaccVX with Halide-HLS [14], which is a state-of-the-art DSL targeting Xilinx FPGAs. As can be seen, HipaccVX uses fewer resources and achieves a higher throughput for the benchmark applications.

HipaccVX transforms a given OpenVX application into a streaming pipeline by replacing virtual images with FIFO


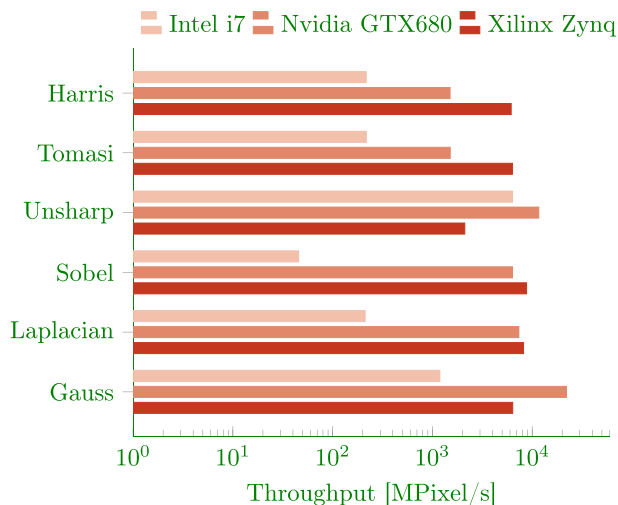


Fig. 10 Comparison of Nvidia VisionWorks v1.6 and HipaccVX on the Nvidia GTX 680. Image sizes are 2048×2048

Table 3 PPnR results for the Xilinx Zynq for images of 1020×1020 and $T_{target} = 5$ ns (corresponds to $f_{target} = 200$ MHz)

App	v		BRAM	SLICE	DSP	Latency [cyc.]
Gauss	1	HipaccVX	8	473	16	1044500
		Halide-HLS	8	1823	50	1052673
	4	HipaccVX	16	1519	64	261649
		Halide-HLS	16	4112	180	266241
Harris	1	HipaccVX	20	1457	34	1042466
		Halide-HLS	16	2688	35	1052673
	2	HipaccVX	20	2326	68	521756
		Halide-HLS	16	4011	70	528385

**Fig. 11** Comparison of throughput for the Nvidia GTX680, Xilinx Zynq, and Intel i7-4790 CPU. The same OpenVX application code is used to generate different accelerator implementations. The HipaccVX framework allows for both code and performance portability by generating optimized implementations for a diverse range of accelerators.

semantics. Thereby, it uses an internal representation in Static Single Assignment (SSA) form. Furthermore, it replicates the innermost kernel to achieve higher parallelism for a given factor v . For practical purposes, we present results only for Xilinx technology. Prior work [13, 19] shows that Hipacc can achieve a performance similar to handwritten examples provided by Intel for image processing. This also indicates that the memory abstractions given in Table 2 are suitable to generate optimized code for HLS tools.

Figure 11 compares the throughputs that were achieved from the same OpenVX application code for different accelerators. Here, we generated OpenCL, CUDA, and Vivado HLS (C++) code to implement a given application on an Intel i7-4790 CPU, an Nvidia GTX680 GPU, and a Xilinx Zynq FPGA, respectively. GPUs and FPGAs can exploit data-level parallelism by processing a significantly higher number of operations in parallel compared to CPUs. This makes them very suitable for computer vision applications. Modern GPUs

operate on a higher clock frequency compared to existing FPGAs; therefore, they could provide higher throughput for the abundantly parallel applications. This is the case for Gauss and Unsharp. Whereas FPGAs can exploit temporal locality using *pipelining* and eliminate unnecessary data transfers to global memory between consecutive kernels. Therefore, all the FPGA implementations in Fig. 11 achieve a similar throughput.

7 Conclusion

In this paper, we presented a set of computational abstractions that are used for expressing OpenVX ' CV functions as well as user-defined kernels. This enables the execution of user nodes on a target accelerator similar to the CV functions and additional optimizations that improve the performance. We presented HipaccVX, an implementation for OpenVX using the proposed abstractions to generate code for GPUs, CPUs, and FPGAs.

Funding Open Access funding provided by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Ashbaugh, B., et al.: OpenCL interoperability with OpenVX graphs. In: Proc. of the 5th Intern. Workshop on OpenCL, p. 26. ACM (2017)
2. BCC Research: Global markets for machine vision technologies. Tech. rep. (2018)

3. Chugh, N., et al.: A DSL compiler for accelerating image processing pipelines on FPGAs. In: Proc. of the Intern. Conf. on Parallel Architecture and Compilation Techniques (PACT), pp. 327–338. IEEE (2016)
4. Du, P., et al.: From CUDA to OpenCL: towards a performance-portable solution for multi-platform GPU programming. *Parallel Comput.* **38**(8), 391–407 (2012)
5. Elliott, G.A., et al.: Supporting real-time computer vision workloads using OpenVX on multicore+GPU platforms. In: Proc. of the Real-Time Systems Symp. (RTSS), pp. 273–284. IEEE (2015)
6. Hegarty, J., et al.: Rigel: flexible multi-rate image processing hardware. *ACM Trans. Graph. (TOG)* **35**(4), 85:1–85:11 (2016)
7. Intel: Intel's OpenVX developer guide
8. Membarth, R., et al.: Hipacc: a domain-specific language and compiler for image processing. *Trans. Parallel Distrib. Syst. (TPDS)* **27**(1), 210–224 (2016)
9. Mori, J.Y., et al.: A design methodology for the next generation real-time vision processors. In: Proc. of the Intern. Symp. on Applied Reconfigurable Computing (ARC), pp. 14–25. Springer (2016)
10. Mullapudi, R.T., et al.: Polymage: Automatic optimization for image processing pipelines. In: Proc. of the Intern. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS), pp. 429–443. ACM (2015)
11. Omidian, H., et al.: An accelerated OpenVX overlay for pure software programmers. In: Proc. of the Intern. Conf. on Field Programmable Technology (FPT) (2018)
12. Omidian, H., et al.: JANUS: a compilation system for balancing parallelism and performance in OpenVX. *J. Phys. Conf. Ser.* **1004**(1), 012011 (2018)
13. Özkan, M.A., et al.: FPGA-based accelerator design from a domain-specific language. In: Proc. of the 26th Intern. Conf. on Field-Programmable Logic and Applications (FPL). IEEE
14. Pu, J., et al.: Programming heterogeneous systems from an image processing DSL. *ACM Trans. Arch. Code Optim. (TACO)* **14**(3), 26:1–26:25 (2017)
15. Qiao, B., et al.: From loop fusion to kernel fusion: a domain-specific approach to locality optimization. In: Proc. of the Intern. Symp. on Code Generation and Optimization (CGO) (2019)
16. Ragan-Kelley, J., et al.: Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines. In: Proc. of the Conf. on Programming Language Design and Implementation (PLDI), pp. 519–530. ACM (2013)
17. Rainey, E., et al.: Addressing system-level optimization with OpenVX graphs. In: Proc. of the Conf. on Computer Vision and Pattern Recognition Workshops, pp. 644–649. IEEE (2014)
18. Reiche, O., et al.: Auto-vectorization for image processing DSLs. In: ACM SIGPLAN Notices, vol. 52, pp. 21–30. ACM (2017)
19. Reiche, O., et al.: Generating FPGA-based image processing accelerators with Hipacc. In: Proc. of the Intern. Conf. on Computer Aided Design (ICCAD), pp. 1026–1033. IEEE (2017)
20. Sérot, J., et al.: CAPH: a language for implementing stream-processing applications on FPGAs. In: *Embedded Systems Design with FPGAs*, pp. 201–224. Springer (2013)
21. Siek, J., et al.: *The Boost Graph Library: User Guide and Reference Manual*. Addison-Wesley, Boston (2002)
22. Steuwer, M., et al.: Generating performance portable code using rewrite rules: from high-level functional expressions to high-performance OpenCL code. *ACM SIGPLAN Not.* **50**(9), 205–217 (2015)
23. Stewart, R., et al.: A dataflow IR for memory efficient RIPL compilation to FPGAs. In: Proc. of the Intern. Conf. on Algorithms and Architectures for Parallel Processing (ICA3PP), pp. 174–188. Springer
24. Tagliavini, G., et al.: Enabling OpenVX support in mW-scale parallel accelerators. In: Proc. of the Intern. Conf. on Compilers, Architectures and Synthesis for Embedded Systems (CASES), pp. 1–10. IEEE (2016)
25. Tagliavini, G., et al.: Optimizing memory bandwidth exploitation for OpenVX applications on embedded many-core accelerators. *J. Real-Time Image Process.* **15**(1), 73–92 (2018)
26. Taheri, S., et al.: Acceleration framework for FPGA implementation of OpenVX graph pipelines. In: Proc. of the Intern. Symp. on Field-Programmable Custom Computing Machines (FCCM), pp. 227–227. IEEE (2018)
27. The Khronos Group: Khronos finalizes and releases OpenVX 1.0 specification for computer vision acceleration. Press Release (2014)
28. The Khronos Group: OpenVX resources (2018)
29. The Khronos Vision Working Group and others: The OpenVX specification v1.2.1 (2018)
30. The Khronos Vision Working Group and others: The OpenVX specification v1.3 (2019)
31. Yang, M., et al.: Making OpenVX really “real-time”. In: Proc. of the Real-Time Systems Symp. (RTSS) (2018)
32. Zhang, J., et al.: DS-DSE: Domain-specific design space exploration for streaming applications. In: Proc. of the Conf. on Design, Automation and Test in Europe (DATE), pp. 165–170. IEEE (2018)


Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



M. Akif Özkan received B.S. and M.S. degrees in Electronics Engineering from Istanbul Technical University, Turkey, in 2011 and 2014, respectively. Since 2015, he is a Ph.D. student in the Department of Computer Science, Friedrich-Alexander University Erlangen-Nürnberg. His primary research interests are SoC ASIC and FPGA design, efficient mapping strategies of image algorithms to heterogeneous architectures and domain-specific languages.



Burak Ok holds a Bachelor's and a Master's degree in Computer Science, both from the Department of Computer Science at the Friedrich-Alexander University Erlangen-Nürnberg. His research interests are embedded systems, parallelization of algorithms, and efficient mapping strategies of image algorithms to heterogeneous architectures.



Bo Qiao received his B.Sc. degree from Shanghai Maritime University, China, in 2013, and his M.Sc. degree from Eindhoven University of Technology, the Netherlands, in 2015. Since 2017, he is a Ph.D. student at the chair for Hardware/Software Co-Design in the Department of Computer Science at the Friedrich-Alexander University Erlangen-Nürnberg. His research interests are compiler optimizations, domain-specific languages in image processing, and general purpose GPU programming.



Jürgen Teich received the M.Sc. degree (Dipl.-Ing. with honors) from the University of Kaiserslautern, Kaiserslautern, Germany, in 1989, and the Ph.D. degree (summa cum laude) from the University of Saarland, Saarbrücken, Germany, in 1993. In 1994, he joined the DSP design group of Prof. E. A. Lee in the Department of Electrical Engineering and Computer Sciences (EECS), University of California at Berkeley, Berkeley, CA, USA (postdoctoral work). From 1995 to 1998, he held a position at the

Institute of Computer Engineering and Communications Networks Laboratory (TIK), ETH Zurich, Zurich, Switzerland (habilitation). From 1998 to 2002, he was Full Professor in the Electrical Engineering

and Information Technology Department, University of Paderborn, Paderborn, Germany. Since 2003, he has been Full Professor in the Department of Computer Science, Friedrich-Alexander University Erlangen-Nürnberg (FAU), Erlangen, Germany, holding a Chair in Hardware/Software Co-Design. Prof. Teich is a Fellow of the IEEE and member of the Academia Europaea. Since 2010, he is the coordinator of the Transregional Research Center 89 on Invasive Computing funded by the German Research Foundation (DFG).



Frank Hannig received the Diploma degree in an interdisciplinary course of study in electrical engineering and computer science from the University of Paderborn, Germany, in 2000; the Ph.D. degree (Dr.-Ing.) and Habilitation degree in computer science from Friedrich-Alexander University Erlangen-Nürnberg (FAU), Germany, in 2009 and 2018, respectively. He has led the Architecture and Compiler Design Group in the Computer Science Department, FAU, since 2004. His primary research

interests are the design of massively parallel architectures, ranging from dedicated hardware to multicore architectures, mapping methodologies for domain-specific computing, and architecture/compiler code-sign. He has authored or coauthored more than 160 peer-reviewed publications. Dr. Hannig serves in the program committees of several international conferences (ARC, ASAP, CODES+ISSS, DATE, DASIP, SAC). He is a Senior Member of the IEEE and an affiliate member of the European Network of Excellence on High Performance and Embedded Architecture and Compilation (HiPEAC).