

# SYCL in the Edge: Performance Evaluation for Heterogeneous Acceleration

Youssef Faqir-Rhazoui (✉ [yelfaqir@ucm.es](mailto:yelfaqir@ucm.es))

Universidad Complutense de Madrid

Carlos García

Universidad Complutense de Madrid

---

## Research Article

**Keywords:** SYCL, CUDA, Edge Computing, Polybench, Jetson, Optic Flow

**Posted Date:** October 16th, 2023

**DOI:** <https://doi.org/10.21203/rs.3.rs-3439288/v1>

**License:**   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

**Additional Declarations:** No competing interests reported.

---

# SYCL in the Edge: Performance Evaluation for Heterogeneous Acceleration

Youssef Faqir-Rhazoui<sup>1\*</sup> and Carlos García<sup>1</sup>

<sup>1\*</sup>Department of Computer Architecture and Automatics, Universidad Complutense de Madrid, Madrid, Spain.

\*Corresponding author(s). E-mail(s): [yelfaqir@ucm.es](mailto:yelfaqir@ucm.es);  
Contributing authors: [garsanca@ucm.es](mailto:garsanca@ucm.es);

## Abstract

Edge computing is essential to handle increasing data volumes and processing capacities. It provides real-time, secure data processing near data sources, like smart devices, alleviating cloud computing energy use and saving network bandwidth. Specialized accelerators, like GPUs and FPGAs, are vital for low-latency edge computing but the requirements to customized code for different hardware and vendors supposes important compatibility issues.

This paper evaluates the potential of SYCL in addressing code portability issues encountered in edge computing. We employed the Polybench suite to compare various SYCL implementations, specifically DPC++ and AdaptiveCpp, with the native solution, CUDA. The disparity between SYCL implementations was negligible, at just 5%.

Furthermore, we evaluated SYCL in the context of specific edge computing applications such as video processing using three different optical flow algorithms. The results exposed a potential performance gap of 19% between CUDA and SYCL. This performance differential is the price one may need to pay when achieving the ability to successfully run the same code on two distinct edge boards with four different architectures, including x86/64 CPU, ARM CPU, NVIDIA GPU, and Intel GPU. These findings underscore SYCL's capacity to increase productivity in term of development costs and facilitate the IoT deployment without being locked into a particular platform or manufacturer.

**Keywords:** SYCL, CUDA, Edge Computing, Polybench, Jetson, Optic Flow.

# 1 Introduction

Edge computing has emerged as a crucial technology due to the growing volume of data and processing demands. Edge computing operates in close proximity to data sources [1, 2], like smart devices, storing and processing data at the network’s edge. It offers fast, real-time, and secure data processing [3], addressing issues like energy consumption in cloud computing, cost reduction, and network bandwidth relief. The increasing prominence of IoT [4] has transformed edge computing into a highly discussed subject, presenting ongoing challenges [5–7] such as selecting the most suitable platform to achieve among others real-time data processing near the data source and ensuring robust data privacy. Nevertheless, one of the foremost challenges persisting in the deployment of IoT systems is the imperative of achieving reduced energy consumption [8] while concurrently upholding robust computational capabilities essential for supporting real-time AI or ML applications.

To address these handicaps, there is a growing trend towards the adoption of accelerators, collectively referred to as xPU (including GPUs, FPGAs, SoCs, and more), which substantially reduce power footprint [9] when compared to general-purpose CPUs. However, employing accelerator languages designed for specific hardware architectures introduces compatibility obstacles meanwhile a custom code for each device (e.g., CUDA, VHDL, etc.) is an imperative. The industry’s motivation to progress in this direction is compounded by two significant challenges: firstly, to select the most suitable system from a huge plethora of devices with notably architectural differences, and secondly, the absence of a universally accepted programming standard. Under this premise we can highlight recent advances with the creation of the Unified Acceleration (UXL) Foundation<sup>1</sup>, announced by the Linux Foundation last September, which proposes oneAPI [10] and SYCL [11] programming as an open-source specification to support a common code base capable of running across multiple architectures.

Until now, native accelerator languages have empowered programmers to deploy code tailored for specialized hardware devices like GPUs, FPGAs, or ASICs. These languages mostly proprietary APIs are engineered to enhance the performance and efficiency of compute-intensive applications. Nevertheless, a common challenge faced by most accelerator languages is their propensity to disrupt compatibility among different hardware architectures. For instance, CUDA [12] is tailored for NVIDIA GPUs, HIP [13] for AMD GPUs, or VHDL for FPGAs.

In contrast, SYCL [11], is a versatile programming model and standard that empowers developers to create heterogeneous parallel code based on ISO C++. SYCL streamlines the process by allowing programmers to write code once, which can then seamlessly execute across multiple vendor CPUs, GPUs, and FPGAs via OpenCL. What sets SYCL apart is its compatibility with modern C++ features like templates, lambdas, and exceptions, which facilitate the expression of parallelism and data movement. SYCL’s remarkable versatility not only facilitates the development of portable applications for diverse heterogeneous edge computing systems [14], including CPUs, GPUs, and FPGAs but also serves as a foundational tool for implementing

---

<sup>1</sup>Unified Acceleration (UXL) Foundation: <https://uxlfoundation.org/>

cost-effective exploration methodologies aimed at reducing development complexity. By employing a unified development approach across multiple edge computing platforms, it becomes possible to discern the architecture that best suits specific problem domains, especially those reliant on critical factors such as power efficiency, cost-effectiveness, and real-time performance requirements.

Moreover, SYCL has been extensively tested on HPC environments and compared with other programming languages such as CUDA, OpenMP or OpenCL [15–17]. While the utilization of SYCL in the realm of edge computing remains relatively unexplored [18] apart from preliminary experiments of porting CUDA codes [19]. We believe that its adoption holds significant potential for achieving performance portability. In this paper, we assess the effectiveness of SYCL on two edge computing boards. We employ a suite of benchmarks to verify SYCL’s compatibility across different architectures. Furthermore, we explore the portability of various motion estimation-based vision algorithms, incorporating accelerators from different vendors.

The following paper is organised as follows. Section 2 introduces the SYCL language and program architecture. In Section 3, the benchmarks use in this study are we discussed. Section 4 focuses on the environment configuration and experiment methodology used. In Section 5, the experiments and results achieved are presented. In section 6 an experiment discussion is performed. And finally, the Section 7 concludes with the main remarks.

## 2 The SYCL Paradigm in a Nutshell

SYCL is a standard (SYCL 2020) developed and maintained by the Khronos Group, similar to other standards such as OpenMP (e.g. 4.5, 5.1, etc) or OpenCL (e.g. 2.1, 3.0, etc) [20–22]. Its main purpose is to enable developers to use any ISO C++ compiler (e.g. GCC, Clang, NVCC, ICC, etc), at least C++ 17, and utilize C++ lambdas to encapsulate device kernels execution. SYCL doesn’t aim to replace other parallel models or backends (e.g. CUDA, HIP, OpenCL, etc.) but rather to complement them. Since all these models are C++-compatible, SYCL uses C++ lambdas to extend the native API of different backends. For instance, when allocating memory on an NVIDIA GPU, a SYCL memory allocation automatically triggers a native CUDA allocation at background. Then, you can consider SYCL as the facade design pattern, which serves as a front-facing interface to other backends [23].

Up to this point, we have solely addressed the SYCL standard, however, it is crucial to recognize that SYCL does not have a singular implementation. The most feature-rich implementation is Intel Data Parallel C++ (DPC++) [14], which not only conforms to the SYCL 2020 standard but also includes other custom features<sup>2</sup>. The Intel oneAPI DPC++/C++ compiler known as DPC++ is a compiler-based implementation, meaning that it is integrated with the new Intel compiler, which is based on and forked from the Clang/LLVM project<sup>3</sup>. It is important to remark that DPC++ compiler is open source, although Intel also offers a commercial alternative available on the oneAPI toolkits. The oneAPI includes additional tools such as profiler

---

<sup>2</sup>These features are typically incorporated into new SYCL releases. However, we did not use them in this paper.

<sup>3</sup><https://github.com/intel/llvm/blob/sycl/sycl/doc/GetStartedGuide.md>

and optimized libraries. While custom features are present in DPC++, we opted not to employ them in our study, with the aim of ensuring the portability of our developments.

The other noteworthy implementation is AdaptiveCpp previously known as hipSYCL, which is a library-based implementation. This means that they have developed a C++ library and rely on third-party compilers. It is generally recommended to use it with the regular Clang/LLVM compiler, which was designed to support CUDA, HIP, OpenMP, and OpenCL source codes [24, 25].

The Clang/LLVM compiler is responsible for compiling SYCL code, and it performs different steps such as front-end, middle-end, and back-end. In the front-end phase, the compiler separates the host code from the device code, while in the middle-end phase, it transforms the device code into an intermediate representation known as LLVM IR. The back-end stage then compiles the LLVM IR representation into the device’s native code and combines everything into a final file called as “fat binary”. This compiler can generate a final binary that can run on multiple devices, including multi-vendor GPUs or even FPGAs, as described in [10, 26].

Figure 1 illustrates a basic SYCL program that sums two vectors into a third one. The usual SYCL scheme begins by creating a queue associated with the target device. The queue receives the following kernels and is responsible for placing them on the device based on a policy. Additionally, we can allocate the program memory, which in this particular example is shared between the host and the device. This feature allows to the host to be able to initialize the memory data on its side without any other restrictions, and later being used by the device without any explicit data movement.

The next step is to invoke the kernel execution on the device. SYCL supports multiple parallel patterns, but in this code example, the *parallel\_for* scheme is performed, specifying the problem size (length) so the kernel is launched length instances or threads. In SYCL by default, the kernel launch is asynchronous so it is mandatory to add the *wait()* clause to maintain the execution coherence. Finally, the memory is freed with the corresponding call because it is tied to the queue.

### 3 Benchmarking SYCL in Edge Platforms

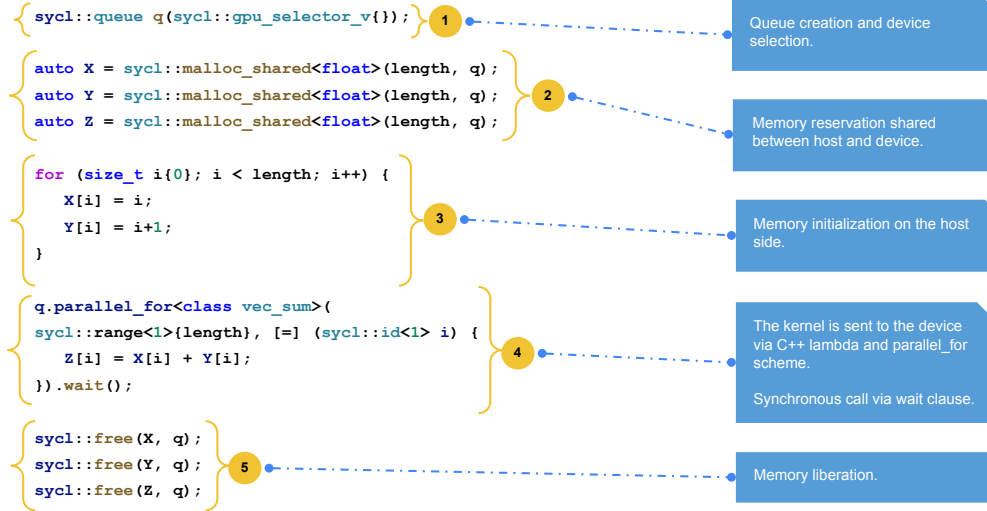
SYCL was tested on both conventional and HPC systems, as the next subsection highlights. However, there is a dearth of literature regarding the potential use of SYCL on the edge. We considered the use of benchmark suites developed for or adapted to SYCL, as they would favor direct comparisons with other programming models like CUDA, and also permit the compilation with various SYCL implementations, including DPC++ and AdaptiveCpp.

#### 3.1 SYCL Benchmark Suites

When it comes to benchmarks, there are several suites available for SYCL. The Rodinia<sup>4</sup> benchmarks are implemented in multiple languages, including SYCL, and encompass a wide range of field benchmarks, such as medical imaging or image compression [15, 27].

---

<sup>4</sup>Rodinia code migrated to SYCL through oneAPI: <https://github.com/artecs-group/rodinia-dpct-dpcpp>



**Fig. 1** SYCL piece of code performing a vector addition.

XSbench<sup>5</sup> is a benchmark suite designed to evaluate the performance of Monte Carlo neutron transport codes used in the field of nuclear engineering and reactor physics. The benchmark suite provides a set of representative problems that simulate the behavior of neutrons in a nuclear reactor. These problems cover a range of materials, geometries, and physics phenomena to assess the performance of different Monte Carlo codes accurately [28].

On its side, HeCBench<sup>6</sup> is a large collection of heterogeneous programming models such as (SYCL, OpenCL, CUDA, etc.). Since HeCBench includes many self-made benchmarks, many others were integrated from Rodinia or XSbench [29].

Polybench<sup>7</sup> consists of a set of computationally intensive kernels that represent common algorithmic patterns found in scientific and engineering applications, such as linear algebra computations, image processing, stencil computations, and more. These kernels are implemented in C, CUDA, OpenMP among other programming languages, and are designed to be representative of real-world HPC workloads [30].

SYCL-Bench<sup>8</sup> provides a set of benchmark kernels and applications that cover a range of common parallel computing patterns and algorithms. These benchmarks are implemented using SYCL and are designed to evaluate the performance of SYCL compilers, runtime systems, and underlying hardware architectures. SYCL-Bench also integrates fifteen kernels/applications from Polybench. This suite also has the possibility to execute on different SYCL implementations, such as DPC++, ComputeCpp, triSYCL, and AdaptiveCpp [31].

<sup>5</sup><https://github.com/ANL-CESAR/XSBench>

<sup>6</sup><https://github.com/zjin-lcf/HeCBench/tree/master>

<sup>7</sup><https://github.com/sgrauerg/polybenchGpu>

<sup>8</sup><https://github.com/unisa-hpc/sycl-bench>

### 3.2 Image Processing for Optic Flow

Optical flow, a crucial component in machine vision systems, calculates a dense field of displacement vector which represents the pixel motion [32] of adjacent frames in consecutive image frames. It holds a pivotal significance in applications of image processing such as video coding, tracking, autonomous driving, or biomedical imaging. It is based on finding the apparent motion of objects in a sequence of images from a camera, extracting a two-dimensional vector related to the object’s motion.

In recent decades, significant advancements in optical flow estimation have been fueled by two main factors. First, the emergence of advanced-level datasets [33–35] has led to continuous improvements in optical flow algorithms. Second, the growing computational resources available in modern microchips such as GPUs accelerators have pushed the development of novel strategies rooted in deep learning approaches.

Horn and Schunck (HS)[36] pioneered the initial optical flow estimation proposal, employing a variational method that leveraged both brightness constancy and spatial smoothness assumptions. It is based on applying spatial and temporal derivatives [37] to the intensity of the image to extract the optical flow vector by solving a multi-dimensional system of equations. To speedup the convergence, hierarchy processing techniques can also be applied [37, 38]. An implementation of the CUDA Horn-Schunck method can be found in the CUDA Toolkit examples<sup>9</sup>, and it has recently been ported to SYCL using an automatic compatibility tool available on the Intel’s oneAPI suite<sup>10</sup>.

Subsequently, the Lucas and Kanade (LK) method [39], proposed by Bruce D. Lucas and Takeo Kanade, is based on the premise that optical flow remains largely consistent within the immediate vicinity of the analyzed pixel. This technique involves solving the core optical flow equations for all pixels within this local neighbourhood through the application of the least squares criterion. It is important to mention that the well-known computer vision library OpenCV implements the Optical Flow functionality based on the LK algorithm<sup>11</sup>.

While HS and LK represent the current state-of-the-art in optical flow techniques and have been used as benchmarks to evaluate ad-hoc implementations in several platforms based on GPUs, FPGAs or DSPs [40–42], they still are pertinent in the embedded system scope. However, it is worth noting that numerous research endeavours have since addressed issues such as high-speed object detection, occlusion handling, illumination changes, and noise reduction. This underscores the community’s commitment to enhancing these techniques [43]. A recent notable proposal that has garnered significant attention from researchers is the TV-L<sup>1</sup> method by Zach et al. [44–46], which employs a variational approach to tackle challenges such as illumination changes, outliers, and flow discontinuities. Other studies, such as those cited in references [47, 48], provide evidence of its advantageous trade-offs on embedded hardware.

---

<sup>9</sup>HSOpticalFlow - Optical Flow: [https://github.com/NVIDIA/cuda-samples/tree/master/Samples/5\\_Domain-Specific/HSOpticalFlow](https://github.com/NVIDIA/cuda-samples/tree/master/Samples/5_Domain-Specific/HSOpticalFlow)

<sup>10</sup>Migrating the HSOpticalFlow Estimation from CUDA to SYCL: <https://www.intel.com/content/www/us/en/developer/articles/technical/migrating-hsopticalflow-from-cuda-to-sycl.html>

<sup>11</sup>OpenCV Optical Flow: [https://docs.opencv.org/3.4/d4/dee/tutorial\\_optical\\_flow.html](https://docs.opencv.org/3.4/d4/dee/tutorial_optical_flow.html)

## 4 Methods

This section briefly describes the configuration and methodology used for the experimentation.

### 4.1 Environment Configuration

Table 1 summarizes the main characteristics of the boards used in this research: the Nvidia Jetson Orin Nano<sup>12</sup> and the UP Squared Pro 7000 Edge.<sup>13</sup> While the first system is based on a SoC equipped with an ARM CPU (Cortex-A78AE) and an NVIDIA Ampere GPU, the second one is based on a SoC equipped with an Intel Atom X7425E and a UHD Graphics Gen 12 GPU.

Despite the Nvidia Jetson Orin Nano can be configured to operate at either seven or fifteen watts, it has been set to fifteen watts. In contrast, the power consumption is not configurable at the UP Squared Pro 7000 Edge board, and it works at twelve watts.

Regarding the software configuration, we utilized two SYCL flavours: DPC++ and AdaptiveCpp. DPC++ can be built from scratch following the instructions from the Intel public repository<sup>14</sup>. As oneAPI is primarily designed for x86/64 architecture, we underwent the process of compiling the DPC++ compiler for ARM architecture, which was then applied to the Orin Nano board. From using AdaptiveCpp<sup>15</sup>, it is necessary to build from the sources on both boards.

Two more aspects worth to mention regarding SYCL implementations. Meanwhile SYCL implementations prioritize the portability of the developed codes for running on various devices, SYCL implementation accomplishes this task in different manners. For instance, DPC++ utilizes OpenCL to run on multicore CPUs, while AdaptiveCpp exploits parallel facilities by means of OpenMP. It is noteworthy to remark that although OpenMP enhances compatibility with non-x86 architectures, it may lead to reduced performance compared to OpenCL [49, 50]. In contrast, DPC++ restricts execution on ARM-based CPUs due to the lack of official OpenCL support. Lastly, the regardless of OpenCL or Intel Level0 backends in the current state of AdaptiveCpp makes impossible its support on Intel GPUs,

### 4.2 Benchmarking Methodology

To assess the performance portability of SYCL, we evaluate both CPU and GPU performance, as the same code can run on both devices. Additionally, we compare the performance of SYCL against the native CUDA code in the Jetson Orin GPU.

Since SYCL-Bench has a specific SYCL benchmarks suite, we have just selected a subset known as the Polybench benchmarks to perform the comparison between SYCL and CUDA. In particular, we choose the Polybench suite available on<sup>16</sup> for CUDA evaluation.

---

<sup>12</sup>Orin Nano specs: <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/>

<sup>13</sup>UP Pro 7000 specs: <https://www.mouser.es/new/aaeon-up/aaeon-up-pro-7000-boards>

<sup>14</sup>DPC++ compiler: <https://github.com/intel/llvm/blob/sycl/sycl/doc/GetStartedGuide.md>

<sup>15</sup><https://github.com/AdaptiveCpp/AdaptiveCpp/blob/develop/doc/installing.md>

<sup>16</sup><https://github.com/sgrauerg/polybenchGpu>



	NVIDIA Jetson Orin Nano	UP Squared Pro 7000 Edge
<b>CPU</b>	<b>Name</b>	ARM Cortex-A78AE
	<b>Frecuency</b>	1.50 GHz (Base) 3.40 GHz (Boost)
	<b>Cores</b>	6
<b>GPU</b>	<b>Name</b>	Intel Atom X7425E
	<b>Frecuency</b>	1 GHz
	<b>Cores</b>	24 execution units
	<b>Performance (FP32)</b>	460.8 GFLOPS
	<b>Driver</b>	23.22.26516.18 (OpenCL)
<b>Memory</b>	8 GB	8 GB
<b>Storage</b>	SD Card Slot & external NVMe via M.2 Key M (not included)	64 GB eMMC
<b>Power Consumption (TDP)</b>	7W / 15W	12W
<b>Price</b>	499\$	399\$

**Table 1** Technical specifications of the Jetson Orin Nano and Up Squared Pro 7000 Edge.

In order to make as fair a comparison as possible, we kept all the default parameter configurations for the tests. Table 2 provides an overview of the benchmark descriptions and the parameters established for the benchmarking.

With the purpose of evaluating modern embedded systems in a more realistic scenario, we choose a workload associated with computer vision as a case study. This experimentation is based on the evaluation of the performance of relevant motion estimation algorithms such as Luka-Kanade (LK), Horn-Schunck (Hs), and TV-L<sup>1</sup>.

The LK algorithm was developed from scratch. Although HS CUDA and SYCL implementations were inspired from the previously mentioned sources, it was necessary to update them to comply with the SYCL 2020 standard. In the case of the TV-L<sup>1</sup> algorithm, no sources were found except for the OpenMP implementation<sup>17</sup> in this work [46]. Both the LK and TV-L<sup>1</sup> algorithms were ported to SYCL using the SYCLomatic tool and later fine-tuned to enhance performance and readability.

To assess this study, we selected a recognized suite of benchmarks widely used in the field of optical flow. The key characteristics of the datasets used are outlined below:

- Schoolgirls with an image resolution of  $432 \times 240$  pixels found at<sup>18</sup>.
- Middlebury dataset [33] includes a twelve scenes with images of  $640 \times 480$ .
- MPI-Sintel [34] is a synthetic dataset based on an animation film which contains frames of  $1024 \times 436$  size.

<sup>17</sup><https://www.ipol.im/pub/art/2013/26/>

<sup>18</sup><https://github.com/hitachinsk/FGT>

Area	Benchmark	Size	Description
<b>Convolution</b>	2DConv	4,096	2D convolution
	3DConv	512	3D convolution
<b>Linear Algebra</b>	2mm	1,024	2 Matrix Multiplications ( $D=A.B$ ; $E=C.D$ )
	3mm	512	3 Matrix Multiplications ( $E=A.B$ ; $F=C.D$ ; $G=E.F$ )
	atax	4,096	Matrix Transpose and Vector Multiplication
	bicg	16,384	BiCG Sub Kernel of BiCGStab Linear Solver
	gemm	1,024	Matrix-multiply $C=\alpha.A.B+\beta.C$
	gesummv	16,384	Scalar, Vector and Matrix Multiplication
	gramschmidt	1,024	Gram-Schmidt decomposition
	mvt	16,384	Matrix Vector Product and Transpose
	syr2k	1,024	Symmetric rank-2k operations
	syrk	1,024	Symmetric rank-k operations
<b>Datamining</b>	correlation	1,024	Correlation Computation
	covariance	1,024	Covariance Computation
<b>Stencil</b>	fdtd-2d	1,024	2-D Finite Different Time Domain Kernel

**Table 2** Polybench suite description and the input parameter size used.

## 5 Experimental Results


This section presents the results achieved from the Polybench suite and optical flow methods. We have divided this section into two parts, each dedicated to one of the experiments.

### 5.1 Polybench Experiments

The Figure 2 illustrates the execution times obtained from Polybench on the Jetson Orin Nano board. It includes the execution on GPU devices using CUDA programming model, AdaptiveCpp and DPC++ for SYCL as well as on the ARM A78AE CPU through AdaptiveCpp. It is worth noting that the DPC++ cannot be used on the CPU due to the absence support of OpenCL on ARM processors. In more detail, as expected benchmarks such as `2DConvolution`, `3DConvolution`, `Atax`, `Fdtd2d`, and `Mvt` get better performance rates using the CUDA implementation, while `Covariance`, `Syr2k`, and `Syrk` for the SYCL version. Furthermore, other benchmarks such as `2mm`, `3mm`, `Bicg`, `Correlation`, `Gemm`, `Gesummv` or `Gramschmidt` achieved almost equivalent execution time using different programming models.

The Table 3 displays the overall performance improvement achieved by each compiler, i.e. the benchmarks based on CUDA are on average  $\times 1.17$  faster in comparison to their AdaptiveCpp counterpart. In line with expectations, the CUDA version outperforms the SYCL versions in all cases, achieving an overall speedup of  $\times 1.17$  and  $\times 1.22$  compared to AdaptiveCpp and DPC++. Shifting our focus to the SYCL implementations, AdaptiveCpp and DPC++ exhibit similar performance metrics. Even in

cases where they diverge, differences are small enough to be encompassed by the standard deviation. Consequently, the disparities between the versions are not relevant. Although, we have also include the execution on ARM A78AE CPU by means of AdaptiveCpp implementation, it is worth noting that the performance is not particularly favorable when compared to GPU times.



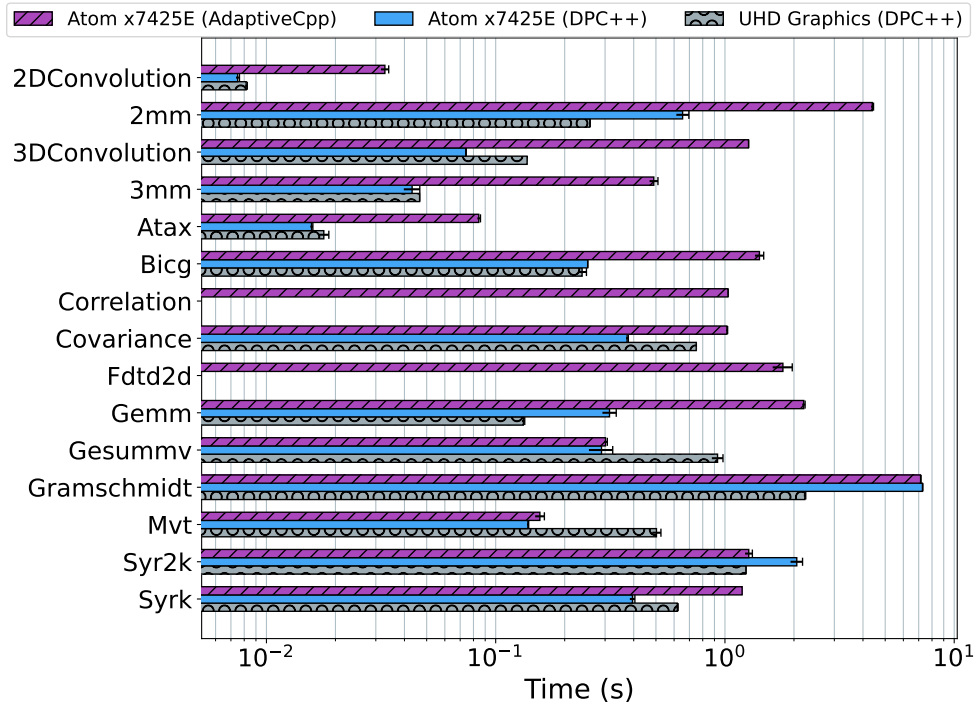
**Fig. 2** Execution time recorded for Polybench suite tests on the Jetson Orin Nano using CUDA and SYCL.

Focusing on the UP Squared Pro 7000 Edge board, the Figure 3 depicts the execution times. It is noteworthy to mention that the Intel UHD GPU couldn't be utilized in conjunction with AdaptiveCpp due to the absence of OpenCL or native baremetal

Speedup	CUDA	AdaptiveCpp	DPC++	AdaptiveCpp (ARM)
<b>CUDA</b>	1	1.17	1.22	5.75
<b>AdaptiveCpp</b>	-	1	1.07	5.46
<b>DPC++</b>	-	-	1	4.99
<b>AdaptiveCpp (ARM)</b>	-	-	-	1

**Table 3** Average speedup obtained from Polybench suite in the Jetson Orin Nano.

support. Moreover, the `Correlation` and `Fdtd2d` couldn't run on the DPC++ implementation due to the requirement for double-precision computations. This issue is motivated from the lack of hardware support of double precision on Intel UHD GPU, and for the Atom CPU the reason is associated to the current OpenCL driver which does not provide support for double-precision.



**Fig. 3** Execution time recorded for Polybench suite tests on the UP Squared Pro 7000 Edge.

Table 4 summarizes the speedups obtained by each device and SYCL implementation. The Atom CPU with the AdaptiveCpp compiler obtains the worst performance because SYCL code is translated to OpenMP, while DPC++ is conducted by the OpenCL backend. This point makes the difference between both implementations

[31, 50]. When comparing DPC++ performance on both CPU and GPU (UHD Graphics), it’s noteworthy that the Atom processor even outperforms the GPU. Given the utilization of default-sized problem parameters, it doesn’t appear to be worthwhile to use the GPU.

Speedup	AdaptiveCpp (Atom)	DPC++(Atom)	DPC++(UHD)
<b>AdaptiveCpp (Atom)</b>	1	0.55	0.79
<b>DPC++ (Atom)</b>	-	1	1.35
<b>DPC++ (UHD)</b>	-	-	1

**Table 4** Average speedup obtained from Polybench suit in the UP Squared Pro 7000 Edge.

## 5.2 Optic Flow Experiments

The Table 5 collects the performance (measured in Frame Per Second-FPS) achieved varying video resolution, GPU device, accelerator implementation, and algorithm. We have also highlighted in bold type the best result fulfilled in each dataset and algorithm. In order to avoid execution variability test have run 10 times, so the table show the average and standard deviation. We would like to clarify that we decided to omit the results on CPU devices since either the ARM Cortex-A78E or the Intel Atom X7425E are far away from the GPU counterpart execution times. Furthermore, for the sake of clarity, we have also removed from the final results from AdaptiveCpp compiler due to the similar times achieved with DPC++.

Regarding the LK algorithm, it is noteworthy that the Intel UHD Graphics is the most suitable device when resolution increases. For the HS algorithm, the Ampere GPU is prominent, but distinguishing between CUDA and DPC++ implementations in terms of performance is challenging, as in most instances, both implementations yield nearly identical fps. Using the TV-L<sup>1</sup> algorithm as benchmark, once again, we observe that the Ampere GPU reports the best performance rates. Diving deeper into the comparison of implementations on the Ampere GPU, an overall difference of approximately 2.9% is observed between CUDA and DPC++. When examining each algorithm individually, we find that LK exhibits a 5.4% improvement with DPC++, HS favours DPC++ by 5%, and the TV-L<sup>1</sup> implementation performs 19% better with CUDA.

On the UHD Graphics side, a direct comparison is made with the Ampere GPU along DPC++. The overall difference is 60.9% in favour of the Orin GPU. When looking at each algorithm individually, we observe a 20.2% improvement for the UHD Graphics in the LK algorithm, an 80.4% advantage for the Ampere GPU in the HS algorithm, and 122% for the Ampere GPU in the TV-L<sup>1</sup> algorithm.

## 6 Discussion

SYCL showed the benefits of using as programming model in the edge market segment. In fact, we could successfully run the same code on different devices without

Dataset	Device	Lukas-Kanade	Horn-Schunck	TV-L <sup>1</sup>
<b>Schoolgirls</b> <b>(432x240)</b>	Ampere GPU	$x \sim = 458$ FPS	$x \sim = 19.8$ FPS	$x \sim = \mathbf{36.5}$ FPS
	(CUDA)	$\sigma_X = 119.4$	$\sigma_X = 0.13$	$\sigma_X = \mathbf{0.44}$
	Ampere GPU	$x \sim = \mathbf{583}$ FPS	$x \sim = \mathbf{21.8}$ FPS	$x \sim = 31.2$ FPS
	(DPC++)	$\sigma_X = \mathbf{49.1}$	$\sigma_X = \mathbf{0.11}$	$\sigma_X = 0.37$
	UHD Graphics	$x \sim = 528$ FPS	$x \sim = 16.24$ FPS	$x \sim = 12.1$ FPS
<b>Middlebury</b> <b>(640x480)</b>	(DPC++)	$\sigma_X = 46.9$	$\sigma_X = 0.19$	$\sigma_X = 0.15$
	Ampere GPU	$x \sim = 168$ FPS	$x \sim = \mathbf{8.92}$ FPS	$x \sim = \mathbf{19.2}$ FPS
	(CUDA)	$\sigma_X = 2.42$	$\sigma_X = \mathbf{0.10}$	$\sigma_X = \mathbf{0.47}$
	Ampere GPU	$x \sim = 147$ FPS	$x \sim = 8.66$ FPS	$x \sim = 15.17$ FPS
	(DPC++)	$\sigma_X = 9.44$	$\sigma_X = 0.10$	$\sigma_X = 0.37$
<b>MPI-Sintel</b> <b>(1024x436)</b>	UHD Graphics	$x \sim = \mathbf{250}$ FPS	$x \sim = 5.01$ FPS	$x \sim = 7.61$ FPS
	(DPC++)	$\sigma_X = \mathbf{2.79}$	$\sigma_X = 0.08$	$\sigma_X = 0.32$
	Ampere GPU	$x \sim = 136$ FPS	$x \sim = 6.36$ FPS	$x \sim = \mathbf{14.6}$ FPS
	(CUDA)	$\sigma_X = 1.27$	$\sigma_X = 0.03$	$\sigma_X = \mathbf{0.12}$
	Ampere GPU	$x \sim = 150$ FPS	$x \sim = \mathbf{6.98}$ FPS	$x \sim = 12.8$ FPS
	(DPC++)	$\sigma_X = 6.61$	$\sigma_X = \mathbf{0.07}$	$\sigma_X = 0.1$
	UHD Graphics	$x \sim = \mathbf{214}$ FPS	$x \sim = 2.98$ FPS	$x \sim = 6.09$ FPS
	(DPC++)	$\sigma_X = \mathbf{2.42}$	$\sigma_X = 0.03$	$\sigma_X = 0.17$

**Table 5** Frames Per Second (FPS) achieved during the execution of optic flow algorithms on various datasets and devices. The table shows the median and the standard deviation of each measure.

an important performance degradation: two edge boards from different vendors were employed for the same task.

In the initial phase, we tested SYCL along with the aforementioned boards using the Polybench suite. This part of the experiment aimed to demonstrate the ability to run the same SYCL code on various architectures, the differences between the most commonly used SYCL implementations, and the minimal performance differences when compared to native implementations, such as CUDA. Polybench results help shed light on these objectives. First and foremost, thanks to SYCL, the Polybench suite was able to run on x86/64 CPU, ARM CPU, NVIDIA GPU, and Intel GPU. This is the primary advantage of SYCL when compared to native implementations, as it simplifies development across various architectures. It is evident that employing the SYCL language to articulate an application’s parallelism not only ensures portability across various architectures and vendors but also enhances productivity.

On one hand, it is also important to mention that regarding to DPC++ and AdaptiveCpp, both encountered difficulties in running on all the architectures tested. DPC++ failed to operate on ARM CPUs, while AdaptiveCpp encountered issues with Intel GPUs. Nonetheless, these problems could be addressed through improved documentation on how to compile AdaptiveCpp for Intel GPUs using OpenCL or Level0

backends, or by employing open-source OpenCL implementations for ARM CPUs such as *pocl*<sup>19</sup>. Regarding their performance, the CUDA GPU architectures exhibited minimal variation, approximately 7%, which depended on the specific benchmark being observed. Conversely, on x86/64 CPUs, AdaptiveCpp failed to achieve comparable results to DPC++ with a notable 45% drop in performance based on the underlying OpenMP conversion for CPU architectures. However, it’s important to note that OpenMP compatibility can be advantageous for emerging architectures like the promising RISC-V.

On the other hand, it is important to consider the comparison between SYCL implementations and CUDA. The overall metric indicates that CUDA outperforms SYCL by approximately 17-22%, depending on the specific implementation. Nevertheless, when examined on a benchmark-by-benchmark basis, the superiority of CUDA is not consistently clear-cut. Out of the five tests performed better with CUDA, while SYCL excelled in the other three, and the remaining eight showed similar performance. In light of these results, it can be inferred that utilizing SYCL does not significantly degrade performance at all.

The second phase of the experiment aimed to test SYCL in real-world scenarios. One common application in the edge computing sphere is computer vision, with a specific focus on optical flow in this case. We evaluated three different datasets—Schoolgirls, Middlebury, and MPI-Sintel—using three optical flow algorithms: LK, HS, and TV-L<sup>1</sup>. There are two points that need to be addressed: the performance difference between CUDA-SYCL and the SYCL comparison between boards.

The primary distinction among the implementations is evident in the TV-L<sup>1</sup> algorithm. HS demonstrates similar processing times in both versions, while LK stands out as an exceptionally lightweight algorithm worth considering. The variances in TV-L<sup>1</sup> processing times—16% for Schoolgirls, 26% for Middlebury, and 14% for MPI-Sintel—should be viewed as the price for achieving code portability. Maintaining different versions of the same algorithm, even if it outperforms, inevitably raises development costs. Therefore, SYCL for edge computing, like in other platforms such as HPC, is no exception and also incurs a "minor" cost of 19% to ensure portability.

Lastly, employing SYCL could also be a sensible choice when comparing across various boards. This is because it helps narrow the gap between the software and the algorithm’s implementation, which can vary depending on the programming language used. To demonstrate this premise, the theoretical performance of the Jetson Orin Nano GPU is 1,280 GFLOPS, meanwhile the UP Squared Pro 7000 Edge GPU offers 460 GFLOPS (a 178% difference). The respective board prices are \$499 and \$399 (a 25% variance). Therefore, given the performance results and the actual performance achieved in optical flow, a comparison based on cost is warranted. Table 6 presents an assessment of the monetary cost (in dollars) per performance unit (FPS).

When analyzing the algorithms, we observed the following cost differences: for LK, the UHD GPU is 4.5% less expensive, while for HS is 27% more cost-effective than the Orin GPU. In the case of TV-L<sup>1</sup>, the Orin board’s cost is 43% lower. In conclusion, these comparative analyses allow us to select the most suitable board based on specific

---

<sup>19</sup><https://github.com/pocl/pocl>

priorities like cost, power consumption, performance, or real-time demands. The ability to use a single, portable code greatly enhances decision-making efficiency and promotes the widespread deployment of IoT applications on various architectures and vendors, eliminating the need for maintaining multiple development efforts or dependency on the commercial policies of a specific manufacturer.

		Lukas-Kanade	Horn-Schunck	TV-L <sup>1</sup>
<b>Shoolgirls</b>	Ampere GPU	\$0.01 per frame	<b>\$0.38 per frame</b>	<b>\$0.27 per frame</b>
<b>(432x240)</b>	UHD Graphics	<b>\$0.01 per frame</b>	\$0.41 per frame	\$0.55 per frame
<b>Basketball</b>	Ampere GPU	\$0.06 per frame	<b>\$0.96 per frame</b>	<b>\$0.55 per frame</b>
<b>(640x480)</b>	UHD Graphics	<b>\$0.03 per frame</b>	\$1.33 per frame	\$0.87 per frame
<b>MPI-Sintel</b>	Ampere GPU	\$0.06 per frame	<b>\$1.19 per frame</b>	<b>\$0.64 per frame</b>
<b>(1024x436)</b>	UHD Graphics	<b>\$0.03 per frame</b>	\$2.23 per frame	\$1.09 per frame

**Table 6** USD per minute and frame computed by the GPUs and datasets.

## 7 Conclusion

The rapid growth of edge computing has introduced various solutions, many of which incorporate low-power accelerators to enhance performance. Accelerators are typically designed to work with specific custom languages such as CUDA, HIP, VHDL, and others. However, this approach creates compatibility issues, as it necessitates customizing the code for each architecture.

This work demonstrated the ability of edge computing to execute and leverage SYCL code on different boards and custom accelerators. We employed the Polybench suite to evaluate various SYCL implementations on the same hardware, and the performance gap was found to be negligible.

Additionally, we utilized a realistic computer vision application based on optical flow algorithms to assess the practical application of SYCL in edge computing scenarios. The experiments revealed a performance disparity between native solutions like CUDA and SYCL. Nevertheless, we deliberated on the significance of SYCL’s portability in development tasks and the trade-off in performance that developers may encounter. Utilizing a single, portable code streamlines decision-making and enables broad IoT deployment across different architectures and vendors, reducing the reliance on multiple development efforts and specific manufacturer policies. To the best of the author’s knowledge, this work represents one of the earliest efforts focused on edge computing and code portability utilizing SYCL.

Future work should focus on incorporating performance portability metrics to facilitate a comparison with the native version. Given the prevalent use of edge computing in image processing and real-time applications, further investigations could explore the advantages of employing SYCL in image processing frameworks such as OpenCV.



Moreover, extending the research to encompass other edge devices and evaluating their performance and power consumption would provide valuable insights.

## Declarations

### Ethical Approval

Not applicable for this item.

### Conflict of interest/Competing interests

Do not have any conflicts of interest with your journal and no mutual conflicts of interest among the authors.

### Authors' contributions

All authors contributed to the research in the main concepts and design. The software was developed by Y. FR. Y. FR also performed experiments. C. G. analyzed the results and proposed methodology in the experimentation phase. All authors write and approve the final manuscript.

### Funding

This paper has been supported by the EU (FEDER), the Spanish MINECO under grants PID2021-126576NB-I00 and TED2021-130123B-I00 funded by MCIN/AEI/10.13039/501100011033 and by European Union “ERDF A way of making Europe” and the NextGenerationEU/PRT.

### Availability of data and materials

Some datasets employed for the current study are available in the *artecs-group/sycl-optic-flow* repository, <https://github.com/artecs-group/sycl-optic-flow/tree/main/dataset>. The full datasets can be found in [43].

### Code availability

The code supporting the results of this article is available in the *artecs-group/sycl-optic-flow* repository, <https://github.com/artecs-group/sycl-optic-flow>.

## References

- [1] Cao, K., Liu, Y., Meng, G., Sun, Q.: An overview on edge computing research. *IEEE Access* **8**, 85714–85728 (2020). <https://doi.org/10.1109/ACCESS.2020.2991734>
- [2] Mansouri, Y., Babar, M.A.: A review of edge computing: Features and resource virtualization. *Journal of Parallel and Distributed Computing* **150**, 155–183 (2021). <https://doi.org/10.1016/j.jpdc.2020.12.015>

- [3] Satyanarayanan, M.: The emergence of edge computing. *Computer* **50**(1), 30–39 (2017). <https://doi.org/10.1109/MC.2017.9>
- [4] Kong, X., Wu, Y., Wang, H., Xia, F.: Edge computing for internet of everything: A survey. *IEEE Internet of Things Journal* **9**(23), 23472–23485 (2022). <https://doi.org/10.1109/JIOT.2022.3200431>
- [5] Tripathy, B., Anuradha, J.: *Internet of Things (IoT): Technologies, Applications, Challenges and Solutions*, p. 358. CRC press, USA (2018). <https://www.routledge.com/Internet-of-Things-IoT-Technologies-Applications-Challenges-and-Solutions/Tripathy-Anuradha/p/book/9780367572921>
- [6] Afzal, B., Umair, M., Shah, G.A., Ahmed, E.: Enabling iot platforms for social iot applications: Vision, feature mapping, and challenges. *Future Generation Computer Systems* **92**, 718–731 (2019)
- [7] Tavana, M., Hajipour, V., Oveisi, S.: Iot-based enterprise resource planning: Challenges, open issues, applications, architecture, and future research directions. *Internet of Things* **11**, 100262 (2020)
- [8] Himeur, Y., Alsalemi, A., Al-Kababji, A., Bensaali, F., Amira, A., Sardianos, C., Dimitrakopoulos, G., Varlamis, I.: A survey of recommender systems for energy efficiency in buildings: Principles, challenges and prospects. *Information Fusion* **72**, 1–21 (2021). <https://doi.org/10.1016/j.inffus.2021.02.002>
- [9] Ramachandran, P., Ranganath, S., Bhandaru, M.K., Tibrewala, S.: A survey of ai enabled edge computing for future networks. 2021 IEEE 4th 5G World Forum (5GWF), 459–463 (2021)
- [10] Intel: oneAPI DPC++ Compiler and Runtime architecture design. <https://intel.github.io/llvm-docs/design/CompilerAndRuntimeDesign.html> (2023)
- [11] Keryell, R., Reyes, R., Howes, L.: Khronos sycl for opencl: a tutorial. In: *Proceedings of the 3rd International Workshop on OpenCL*, pp. 1–1 (2015)
- [12] Buck, I.: Gpu computing with nvidia cuda. In: *ACM SIGGRAPH 2007 Courses*, p. 6 (2007)
- [13] Bauman, P., Chalmers, N., Curtis, N., Freitag, C., Greathouse, J., Malaya, N., McDougall, D., Moe, S., van Oostrum, R., Wolfe, N., et al.: *Introduction to amd gpu programming with hip*. Presentation at Oak Ridge National Laboratory. Online at: <https://www.olcf.ornl.gov/calendar/intro-to-amd-gpu-programming-with-hip> (2019)
- [14] Reinders, J., Ashbaugh, B., Brodman, J., Kinsner, M., Pennycook, J., Tian, X.:

- Data Parallel C++: Mastering DPC++ for Programming of Heterogeneous Systems Using C++ and SYCL. Second Edition, p. 558. Springer, USA (2023). <https://doi.org/10.1007/978-1-4842-9691-2>
- [15] Castaño, G., Faqir-Rhazoui, Y., García, C., Prieto-Matías, M.: Evaluation of intel's dpc++ compatibility tool in heterogeneous computing. *Journal of Parallel and Distributed Computing* **165**, 120–129 (2022). <https://doi.org/10.1016/j.jpdc.2022.03.017>
  - [16] Deakin, T., McIntosh-Smith, S.: Evaluating the performance of hpc-style sycl applications. In: *Proceedings of the International Workshop on OpenCL. IWOCL '20*. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3388333.3388643>. <https://doi.org/10.1145/3388333.3388643>
  - [17] Breyer, M., Van Craen, A., Pflüger, D.: A comparison of sycl, opencl, cuda, and openmp for massively parallel support vector machine classification on multi-vendor hardware. In: *International Workshop on OpenCL. IWOCL'22*. Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3529538.3529980>. <https://doi.org/10.1145/3529538.3529980>
  - [18] Kang, P.: Programming for high-performance computing on edge accelerators. *Mathematics* **11**(4) (2023). <https://doi.org/10.3390/math11041055>
  - [19] Angus, D., Georgiev, S., Arroyo Gonzalez, H., Riordan, J., Keir, P., Goli, M.: Porting sycl accelerated neural network frameworks to edge devices. In: *Proceedings of the 2023 International Workshop on OpenCL. IWOCL '23*. Association for Computing Machinery, New York, NY, USA (2023). <https://doi.org/10.1145/3585341.3585346>. <https://doi.org/10.1145/3585341.3585346>
  - [20] Khronos SYCL working group: SYCL Specification. <https://registry.khronos.org/SYCL/> (2023)
  - [21] OpenMP: The OpenMP Specification. <https://www.openmp.org/> (2023)
  - [22] Khronos SYCL working group: The OpenCL Specification. <https://registry.khronos.org/OpenCL/> (2023)
  - [23] Ludwig, K.: Performance portability and evaluation of heterogeneous components of seissol targeted to upcoming intel hpc gpus (2021)
  - [24] LLVM-Project: User Guide for AMDGPU Backend. <https://www.llvm.org/docs/AMDGPUUsage.html> (2023)
  - [25] Marangoni, M., Wischgoll, T.: Togpu: Automatic source transformation from c++ to cuda using clang/llvm. *Electronic Imaging* **2016**(1), 1–9 (2016)
  - [26] illuhad: AdaptiveCpp design and architecture. <https://github.com/OpenSYCL/>

- [OpenSYCL/blob/develop/doc/architecture.md](#) (2021)
- [27] Jin, Z.: The rodinia benchmark suite in sycl. Technical report, Argonne National Lab.(ANL), Argonne, IL (United States). Argonne Leadership ... (2020)
  - [28] Tramm, J.R., Siegel, A.R., Islam, T., Schulz, M.: Xsbench-the development and verification of a performance abstraction for monte carlo reactor analysis. The Role of Reactor Physics toward a Sustainable Future (PHYSOR) (2014)
  - [29] Alpay, A., Soproni, B., Wünsche, H., Heuveline, V.: Exploring the possibility of a hipsycl-based implementation of oneapi. In: International Workshop on OpenCL. IWOCL'22. Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3529538.3530005>. <https://doi.org/10.1145/3529538.3530005>
  - [30] Grauer-Gray, S., Xu, L., Searles, R., Ayalasomayajula, S., Cavazos, J.: Auto-tuning a high-level language targeted to gpu codes. In: 2012 Innovative Parallel Computing (InPar), pp. 1–10 (2012). <https://doi.org/10.1109/InPar.2012.6339595>
  - [31] Lal, S., Alpay, A., Salzmann, P., Cosenza, B., Hirsch, A., Stawinoga, N., Thoman, P., Fahringer, T., Heuveline, V.: Sycl-bench: a versatile cross-platform benchmark suite for heterogeneous computing. In: Euro-Par 2020: Parallel Processing: 26th International Conference on Parallel and Distributed Computing, Warsaw, Poland, August 24–28, 2020, Proceedings 26, pp. 629–644 (2020). [https://doi.org/10.1007/978-3-030-57675-2\\_39](https://doi.org/10.1007/978-3-030-57675-2_39). Springer
  - [32] Stiller, C., Konrad, J.: Estimating motion in image sequences. IEEE Signal Processing Magazine **16**(4), 70–91 (1999). <https://doi.org/10.1109/79.774934>
  - [33] Baker, S., Roth, S., Scharstein, D., Black, M.J., Lewis, J.P., Szeliski, R.: A database and evaluation methodology for optical flow. In: 2007 IEEE 11th International Conference on Computer Vision, pp. 1–8 (2007). <https://doi.org/10.1109/ICCV.2007.4408903>
  - [34] Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part VI 12, pp. 611–625 (2012). Springer
  - [35] Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The KITTI dataset. The International Journal of Robotics Research **32**(11), 1231–1237 (2013). <https://doi.org/10.1177/0278364913491297>
  - [36] Horn, B.K.P., Schunck, B.G.: Determining optical flow. Artificial Intelligence **17**(1), 185–203 (1981). [https://doi.org/10.1016/0004-3702\(81\)90024-2](https://doi.org/10.1016/0004-3702(81)90024-2)

- [37] Sun, D., Roth, S., Black, M.J.: Secrets of optical flow estimation and their principles. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 2432–2439 (2010). <https://doi.org/10.1109/CVPR.2010.5539939>
- [38] Borzi, A., Schulz, V.: Multigrid methods for pde optimization. *SIAM Review* **51**(2), 361–395 (2009) <https://doi.org/10.1137/060671590>. <https://doi.org/10.1137/060671590>
- [39] Lucas, B.D., Kanade, T.: An iterative image registration technique with an application to stereo vision. In: Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2. IJCAI’81, pp. 674–679. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1981)
- [40] Botella, G., Garcia, A., Rodriguez-Alvarez, M., Ros, E., Meyer-Baese, U., Molina, M.C.: Robust bioinspired architecture for optical-flow computation. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* **18**(4), 616–629 (2010). <https://doi.org/10.1109/TVLSI.2009.2013957>
- [41] Gong, Y., Zhang, J., Liu, X., Li, J., Lei, Y., Zhang, Z., Yang, C., Geng, L.: A real-time and efficient optical flow tracking accelerator on fpga platform. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 1–14 (2023). <https://doi.org/10.1109/TCSI.2023.3298969>
- [42] Jaiswal, D., Kumar, P.: A survey on parallel computing for traditional computer vision. *Concurrency and Computation: Practice and Experience* **34**(4), 6638 (2022)
- [43] Zhai, M., Xiang, X., Lv, N., Kong, X.: Optical flow and scene flow estimation: A survey. *Pattern Recognition* **114**, 107861 (2021). <https://doi.org/10.1016/j.patcog.2021.107861>
- [44] Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime tv-l1 optical flow. In: Proceedings of the 29th DAGM Conference on Pattern Recognition, pp. 214–223. Springer, Berlin, Heidelberg (2007)
- [45] Wedel, A., Pock, T., Zach, C., Bischof, H., Cremers, D.: An improved algorithm for tv-l1 optical flow. In: Statistical and Geometrical Approaches to Visual Motion Analysis: International Dagstuhl Seminar, Dagstuhl Castle, Germany, July 13–18, 2008. Revised Papers, pp. 23–45. Springer, Berlin, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-03061-1\\_2](https://doi.org/10.1007/978-3-642-03061-1_2)
- [46] Sánchez Pérez, J., Meinhardt-Llopis, E., Facciolo, G.: TV-L1 Optical Flow Estimation. *Image Processing On Line* **3**, 137–150 (2013). <https://doi.org/10.5201/ipol.2013.26>
- [47] Romera, T., Petreto, A., Lemaitre, F., Bouyer, M., Meunier, Q., Lacassagne, L.,

- Etiemble, D.: Optical flow algorithms optimized for speed, energy and accuracy on embedded gpus. *Journal of Real-Time Image Processing* **20**(2), 32 (2023). <https://doi.org/10.1007/s11554-023-01288-6>
- [48] Romera, T., Petreto, A., Lemaitre, F., Bouyer, M., Meunier, Q., Lacassagne, L.: Implementations impact on iterative image processing for embedded gpu. In: 2021 29th European Signal Processing Conference (EUSIPCO), pp. 736–740 (2021). <https://doi.org/10.23919/EUSIPCO54536.2021.9615947>
- [49] Alpay, A., Heuveline, V.: Sycl beyond opencl: The architecture, current state and future direction of hipsycl. In: Proceedings of the International Workshop on OpenCL. IWOCL '20. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3388333.3388658>
- [50] Alpay, A.: hipSYCL 0.9.2 - compiler-accelerated CPU backend, nvc++ support and more. <https://adaptivecpp.github.io/hipsycl/release/cpu/extension/nvc++/hipsycl-0.9.2/>