



Beehive SPIR-V Toolkit

Document Version

Final published version

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Fumero Alfonso, J., Rethy, G., Stratikopoulos, A., Foutris, N., & Kotselidis, C.-E. (2023). *Beehive SPIR-V Toolkit: A Composable and Functional API for Runtime SPIR-V Code Generation*.

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.



Beehive SPIR-V Toolkit

A Composable and Functional API for Runtime SPIR-V Code Generation

Juan Fumero
juan.fumero@manchester.ac.uk
The University of Manchester
United Kingdom

György Rethy*
grethy@student.ethz.ch
ETH Zurich
Switzerland

Athanasios Stratikopoulos
{first}.{last}@manchester.ac.uk
The University of Manchester
United Kingdom

Nikos Foutiris
nikos.foutiris@manchester.ac.uk
The University of Manchester
United Kingdom

Christos Kotselidis
christos.kotselidis@manchester.ac.uk
The University of Manchester
United Kingdom

Abstract

The Standard Portable Intermediate Representation (SPIR-V) is a low-level binary format designed for representing shaders and compute kernels that can be consumed by OpenCL for computing kernels, and Vulkan for graphics rendering. As a binary representation, SPIR-V is meant to be used by compilers and runtime systems, and is usually performed by C/C++ programs and the LLVM software and compiler ecosystem. However, not all programming environments, runtime systems, and language implementations are C/C++ or based on LLVM.

This paper presents the Beehive SPIR-V Toolkit; a framework that can automatically generate a Java composable and functional library for dynamically building SPIR-V binary modules. The Beehive SPIR-V Toolkit can be used by optimizing compilers and runtime systems to generate and validate SPIR-V binary modules from managed runtime systems. Furthermore, our framework is architected to accommodate new SPIR-V releases in an easy-to-maintain manner, and it facilitates the automatic generation of Java libraries for other standards, besides SPIR-V. The Beehive SPIR-V Toolkit also includes an assembler that emits SPIR-V binary modules from disassembled SPIR-V text files, and a disassembler that converts the SPIR-V binary code into a text file. To the best of our knowledge, the Beehive SPIR-V Toolkit is the first Java programming framework that can dynamically generate SPIR-V binary modules.

*Work done while he was associated at The University of Manchester.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

VMIL '23, October 23, 2023, Cascais, Portugal

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0401-7/23/10.

<https://doi.org/10.1145/3623507.3623555>

To demonstrate the use of our framework, we showcase the integration of the SPIR-V Beehive Toolkit in the context of the TornadoVM, a Java framework for automatically offloading and running Java programs on heterogeneous hardware. We show that, via the SPIR-V Beehive Toolkit, TornadoVM is able to compile code 3x faster than its existing OpenCL C JIT compiler, and it performs up to 1.52x faster than the existing OpenCL C backend in TornadoVM.

CCS Concepts: • Software and its engineering → API languages; Just-in-time compilers; Runtime environments.

Keywords: API, Library, Java, Metaprogramming, Runtime Code Generation, SPIR-V

ACM Reference Format:

Juan Fumero, György Rethy, Athanasios Stratikopoulos, Nikos Foutiris, and Christos Kotselidis. 2023. Beehive SPIR-V Toolkit: A Composable and Functional API for Runtime SPIR-V Code Generation. In *Proceedings of the 15th ACM SIGPLAN International Workshop on Virtual Machines and Intermediate Languages (VMIL '23)*, October 23, 2023, Cascais, Portugal. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3623507.3623555>

1 Introduction

The Standard Portable Intermediate Representation (SPIR-V) [15], maintained by the Khronos group [14], is an intermediate binary format for representing graphics and parallel computation that exploit parallel execution on heterogeneous hardware, such as GPUs and FPGAs. SPIR-V was proposed in March 2015 [15], and it can be used as an extension from OpenCL 2.1 to compute kernels that consume SPIR-V binaries instead of OpenCL C source code. Several companies, including Intel, AMD, NVIDIA, Codeplay and Google offer their own implementations¹, tools and compilers for generating and consuming SPIR-V binary code. Besides those tools, the Khronos Group has created a set of tools and utilities to compile, disassemble, validate and

¹<https://github.com/KhronosGroup/SPIRV-Headers/blob/main/include/spirv/spir-v.xml>

optimize SPIR-V binary code [16]. However, all these tools are only available for LLVM-based programming languages implementations, such as C/C++ (e.g., `clang` [21]), as they operate under the LLVM [22] ecosystem and compiler infrastructure².

This feature hinders the exploitation of SPIR-V tools from programming languages, such as Java, R, Ruby, and Scala, which are built on top of managed runtime systems, such as the Java Virtual Machine (JVM). Applications and systems' software, such as optimizing compilers and runtime systems (e.g., GraalVM [10]), that are implemented in those programming languages, cannot use existing standard SPIR-V tools in a direct way.

To enable the utilization of LLVM-based tools from managed programming languages, it is necessary to invoke the tools using native interfaces such as the Java Native Interface (JNI). For instance, Java programs could invoke existing LLVM tools by providing native methods that can be used as library calls via JNI.

Another way is to invoke LLVM utilities as an invocation of a new process from a guest programming language (e.g., by creating a subprocess that invokes the LLVM tools from Java). However, such integrations pose the following challenges: i) high complexity due to the interaction between different programming languages and the runtime system, ii) high difficulty to maintain, and iii) the necessity to recompile the Java JNI dispatch code for every new release of SPIR-V.

In this paper, we present the Beehive SPIR-V Toolkit, a Java framework that enables the automatic generation of a composable and functional Java library based on standardized grammar files (e.g., SPIR-V grammar). The generated library can be used to enhance software written in JVM programming languages with functionality (i.e., generation, validation, optimization) that is currently offered by LLVM-based SPIR-V tools. Besides, the proposed framework is architected to accommodate new SPIR-V releases in an easy-to-maintain manner and facilitate the automatic generation of Java libraries for other standards, besides SPIR-V.

To enable this functionality, the Beehive SPIR-V Toolkit architecture encompasses a template system engine that demonstrates a well-known technique based on model-driven engineering [12] for automatically generating Java libraries and composable APIs based on standardized specifications. Note that this is a common software engineering technique, and we have enabled it to provide new APIs that can be used by compilers and runtime systems in a transparent manner without the need to reconfigure new rules for every new version of the SPIR-V standard.

Lastly, the Beehive SPIR-V Toolkit provides a client utility that can: i) assemble a SPIR-V binary code from a kernel

description that is stored in a text file, and ii) disassemble a binary code to a kernel description stored in a file.

The Beehive SPIR-V Toolkit has been developed to be utilized by existing Just-In-Time (JIT) compilers and managed runtime systems, as a means to facilitate SPIR-V code generation and analysis from high-level JVM programming languages. The source code of the Beehive SPIR-V Toolkit is available on GitHub³ as an open-source project. As a use case, we have extended the TornadoVM [7, 13] (a heterogeneous programming framework for Java) JIT Compiler and runtime system, with a new backend for generating SPIR-V using the Beehive SPIR-V Toolkit as a library.

In a nutshell, this paper makes the following contributions:

- It presents a template-based technique to automatically generate composable and functional Java programming libraries from standard grammar files that are defined using the JSON format, showcasing the technique in the context of SPIR-V.
- It presents the Beehive SPIR-V Toolkit, a framework that automatically generates a composable and functional Java library for building SPIR-V binary modules at runtime. The proposed framework includes an assembler that emits SPIR-V binary modules from disassembled SPIR-V text files, as well as a disassembler that converts the SPIR-V binary code to a kernel description stored in a text file.
- It presents an extension of the TornadoVM JIT compiler and runtime system for automatically generating SPIR-V from Java bytecode.
- It presents a performance evaluation of our framework against existing OpenCL C backend of TornadoVM, showing end-to-end speedups of up to 3x for the code generation and up to 1.52x speedup for the execution of the generated code.

2 Background and Motivation

SPIR-V was created to address the need for a universal intermediate language for parallel computing and graphics processing applications. Similarly to OpenCL, SPIR-V language's tools generate code once and then compile it to different target architectures, such as CPUs, GPUs, and FPGAs, among others, but in binary format. However, with SPIR-V, developers do not need to distribute the source code of the compute kernels, but rather the binary representation of those kernels. Additionally, SPIR-V supports extensions from multiple vendors and parties (e.g., SPIR-V Math Extended Instruction Set [3]).

Since SPIR-V is a binary format representation, it might be more convenient for compilers to generate this Intermediate Representation (IR) rather than, for example, OpenCL C source code. In addition, SPIR-V kernels can be consumed

²Note that, although SPIR-V was created to be independent of the LLVM-IR, many of the tools and utilities are still under the LLVM and C++ compiler infrastructure.

³<https://github.com/beehive-lab/beehive-spirv-toolkit>

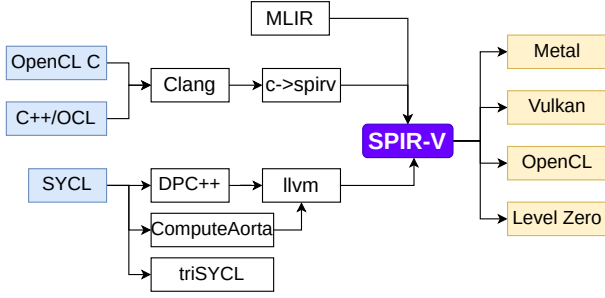


Figure 1. Language description and tooling ecosystem for SPIR-V as described in the SPIR-V SPEC [15].

by OpenCL programs (from OpenCL 2.1) and Intel Level Zero [27] applications.

Besides, SPIR-V allows high-level language front-ends to produce programs in a standardized intermediate format that can be consumed by driver implementations for Vulkan [18], OpenGL [30] or OpenCL [33], thus eliminating the need for high-level language front-end compilers in device drivers. This not only simplifies driver architecture, but also supports a broad range of language and framework front-ends for different hardware architectures, promoting an active ecosystem of open-source tools for analysis, porting, debugging, and optimization. Lastly, SPIR-V can speedup the final compilation stage by the GPU drivers, because the kernels are expressed in an intermediate low-level representation.

But, what is missing? Figure 1 shows the SPIR-V Language ecosystem and some of the tools associated with SPIR-V as they are specified in the SPIR-V Standard [15]. From the figure, we can see that SPIR-V represents the intermediate step between the high-level languages and the low-level programming models for computing and graphics processing. From the high-level language descriptions, we see that all tools and languages represent conversions from the C and C++ programming languages, with a heavy focus on the LLVM compiler ecosystem.

While SPIR-V was designed to be independent of the LLVM software and compiler ecosystem, current tools and translators are centered around LLVM. This paper presents a framework to generate SPIR-V code from managed runtime programming languages and managed runtime environments, such as the Java Virtual Machine and the Java programming language. As far as we know, this paper presents the first Java library to dynamically generate and validate SPIR-V code.

3 Beehive SPIR-V Toolkit

This section presents the Beehive SPIR-V Toolkit. First, it shows an overview of the overall software architecture, and it shows how the Beehive SPIR-V Toolkit works (Section 3.1). Then, Sections 3.2-3.3.4 describe each component in detail.

3.1 System Overview

To generate the Java library and the API, we have implemented a **Template System Engine (TSE)** (Section 3.2.2). The TSE is also fully implemented in Java, and it generates new Java types that will be used to compose SPIR-V binary modules. The generated library is then distributed as a standard Java JAR file to be imported by Java client applications, runtime systems and compilers that are also implemented in Java. Figure 2 shows the three main components, as follows:

SPIR-V Library Generator: This component generates the Java SPIR-V Library and the APIs for composing SPIR-V binary modules. The SPIR-V Library Generator takes, as inputs, a set of grammar files described in the JSON format that represents the SPIR-V grammar specified by the Khronos Group, and a set of templates implemented in Java. The Java templates provide the initial scaffolding to facilitate the automatic generation of the functional and composable APIs, as well as helper Java classes.

The core of the template generator is the TSE component, which is a fully automated system that allows the generation of composable APIs for every new version of the SPIR-V standard based on the new grammar files. We refer to the generated API as *composable* because, in order for client applications to build new SPIR-V modules, it makes use of function composition to create and build new constructs for the resulting SPIR-V binary modules.

SPIR-V Library: This component is the result of the SPIR-V library generator, and it is meant to be distributed to Java clients as a standard Java library in a JAR format. This library allows Java developers to implement JVM-based programs (e.g., those implemented in Java, Scala, etc.) that dynamically create SPIR-V modules. Furthermore, the SPIR-V library includes functionality to check and validate basic rules for SPIR-V modules.

SPIR-V Client Utility: This component provides a client application that can be used for assembling and disassembling SPIR-V modules as a standalone tool. The client application assembles and disassembles SPIR-V code from text to binary format and vice-versa. Furthermore, this component acts as an end-to-end application that demonstrates how the generated SPIR-V library can be used with complete examples.

As follows, we describe, in more detail, each of these components.

3.2 SPIR-V Template Library Generator

The objective of the SPIR-V template library generator is to automatically generate Java libraries that can be utilized by other software components (e.g., optimizing compilers, runtime systems, etc.) as a means to develop SPIR-V code. The generated API is **functional** (as in functional programming in which each new operation returns a new object of




Figure 2. Overview of the main components of the Beehive SPIR-V Toolkit. It provides three main components: 1) the SPIR-V Library Generator; 2) the SPIR-V Library, and 3) a SPIR-V Client. The purple sub-components are provided by the Beehive SPIR-V Toolkit. The SPIR-V Grammar file is provided by the Khronos Group in GitHub [16]. The green components represents two data structures that store all SPIR-V instructions and operands that will be used to generate the new Java types.

the requested type, such as an integer addition, and there is no mutation state of internal properties of the Java objects), and **composable** (new instructions are composed of other instructions, that are, in turn, built by creating and composing new objects of the requested types). We demonstrate the details of the composition of calls in Section 3.3.

Furthermore, the library generator is capable of creating an assembler and a disassembler that can adhere to various versions of the SPIR-V standards. Thus, any future extensions in the SPIR-V standard can be easily adopted in existing systems' software by updating to newly generated SPIR-V formats in a transparent and automatic manner. To enable this functionality, the SPIR-V library generator employs two open-source software components: (1) the Jackson Annotation Library [11], to parse the JSON grammar files and create a list of instructions and operands to be generated; and (2) the Apache Freemarker [2], to generate all Java types for all instructions and operands that were previously parsed.

3.2.1 Parsing the JSON files from the SPIR-V Standard Specification. As part of the tools for the SPIR-V standard, the Khronos Group provides, in their repository [31], the JSON files that specify the grammar for the SPIR-V Intermediate Language, and it defines all SPIR-V instructions, SPIR-V operands, and SPIR-V types.

Figure 3 presents the main structure of serialized JSON objects that reside in the JSON file. The blue blocks represent JSON objects that can be expanded to other SPIR-V category types, while the light-grey blocks represent the final objects that correspond to SPIR-V final types. As illustrated in Figure 3, the SPIR-V grammar is composed of a *Magic Number* along with a set of *Instructions* and a list of *Operand Kinds*.

To parse the SPIR-V grammar JSON file and generate the corresponding Java classes that will compose the SPIR-V library, we use the Java Jackson annotation framework, which is a popular and widely used Java framework to map Java objects into/from files. Thus, the mapped Java objects from the JSON file are structured following the object definition in the standardized SPIR-V grammar.

Furthermore, our library supports the OpenCL Extended Math Instructions set [3]. Note that, although the Beehive SPIR-V Toolkit implementation covers the whole SPIR-V grammar as defined in the SPIR-V core specification, it does not cover the complete list of SPIR-V extensions, such as GLSL, and AMD extensions. The reason is that at the current status of the Beehive SPIR-V Toolkit, our focus has been on providing support for parallel compute kernels (to be integrated with OpenCL and Intel Level Zero [27] runtime systems) rather than graphics processing. However, the proposed framework can be extended in the same way to generate all Java classes for graphics processing as well.

The result of the JSON parser process for the generation of the Beehive SPIR-V Toolkit is a list of two arrays to store instructions and operands with their corresponding fields. These two arrays are used in the next step (in the TSE) to generate all the Java types to represent SPIR-V instructions and operands.

3.2.2 Template System Engine (TSE). The core functionality of the SPIR-V Library Generator is the automatic generation of libraries through the TSE, which generates a Java library that uses function composition to build, at runtime, SPIR-V binary modules. The TSE component generates a set of Java classes, mapper classes and the utilities needed for the SPIR-V assembler and dissembler. As such, these Java templates follow three different categories within the TSE component: a) instructions templates; b) operand templates; and c) mappers templates.

The way the TSE component works is as follows: the engine takes, as inputs, the Java classes with the parsed JSON objects, and a set of Java template classes. The former classes map JSON objects to Java objects by the Jackson Annotation Framework (Section 3.2.1), while the latter Java classes are used for the template processing that will be used for generating all instructions, operands, and mappers. Note that this is an automatic process, and it builds all classes the first time that the library is built while accepting a grammar JSON file as input.




Figure 3. Block categories of the JSON file format that specifies the grammar of the SPIR-V Intermediate Language.

Listing 1. Pseudocode for the Instruction Generator.

```

1 Template t = openTemplate("instruction.template");
2   foreach (i: grammar.getInstructions())
3     do
4       templateFile = createFile(i)
5       t.eprocess(templateFile, i);
6       t.writeFile();
7   done

```

Templates for SPIR-V Instructions/Operands. Each SPIR-V instruction that was parsed from the JSON files corresponds to a new Java class that is generated based on a template within the TSE component. Listing 1 shows a pseudocode of the TSE component that generates all Java classes for all the SPIR-V instructions. Line 1 opens a file that represents the template in which the instructions will be written. Then, lines 2-7 traverse the array that contains all parsed instructions from the input JSON file to generate a Java class per instruction. Each file contains a set of special characters represented as strings that the TSE takes to substitute for specific values and new characters. The end result is a set of valid Java classes that compose the whole SPIR-V Java library.

As a result of the processing of all the SPIR-V instructions, the TSE component dynamically generates *366 new Java classes* (all instructions available in the SPIR-V 1.2), and *667 new Java classes* for the 1.6 version of the SPIR-V standard, while all instruction classes inherit from the same `SPIRVInstruction` base class that we provide.

Similarly to the templates for the instructions, the TSE component generates a set of Java types for all SPIR-V operands and SPIR-V kinds that will be then copied to a Java sub-package of the final library. In total, the TSE software component generates *34 Java classes* for the SPIR-V 1.2 and *44 new Java classes* for the SPIR-V 1.6 standard, that are mapped to different types of SPIR-V operands, such as `SPIRVBuiltin`, and `LiteralExtInsInteger`. Note that this process is fully automatic, and it is triggered every time the Beehive SPIR-V Toolkit is built from the source code.

Templates for Proxy Classes and Mappers. With the generated Java classes that compose the Java Library for SPIR-V instructions and operands, it is possible to dynamically build SPIR-V modules just by using the generated composable and functional API. However, the SPIR-V library can be also used to disassemble SPIR-V binary modules as well as to assemble modules expressed in a text format. To achieve that, the TSE software component also generates proxy classes that map both the description of the input SPIR-V module in a text format to binary (assembling); and the reverse action from binary modules to text (disassembling).

The TSE component generates three new classes per Java Proxy type: one module for the assembler, one module for the disassembler, and the Java classes that correspond to the mappers for instructions, operands, and extended OpenCL instructions (e.g., OpenCL math operations). We define these proxy classes as Java helpers for writing SPIR-V and disassembling SPIR-V binary modules. As a result, once the SPIR-V Library is generated, it is ready to be consumed by client applications.

Porting Different Standard Versions. First, we ported the SPIR-V 1.2 standard, because all major driver vendors (e.g., Intel) use this version. However, we also ported to the latest standard available, 1.6 (also called *unified*). We experienced a smooth transition, and we only needed to accommodate one of the Java templates due to the addition of a number at the beginning of a SPIR-V instruction. Since we map every SPIR-V instruction name to a Java class, we cannot assign a class name that starts with a number. Thus, in this case, we start with the symbol underscore for the generated Java type⁴.

3.3 SPIR-V Library

Through the generated Java classes and utility classes, the generated API has two main Java types: i) the instructions and operands (as we discussed in Section 3.2.2); and ii) the scopes. A scope data type represents a block of instructions that defines different visibility regions within the SPIR-V

⁴<https://github.com/beehive-lab/spirv-beehive-toolkit/commit/217208a2ef9b7a7ec4eb4d59e32c45a2344d8582>

Listing 2. Example of how to build a header for a SPIR-V module.

```

1  SPIRVModule module = new SPIRVModule(
2      new SPIRVHeader(
3          1, // SPIRV_MAJOR_VERSION
4          2, // SPIRV_MINOR_VERSION
5          32, // SPIRV_GENERATOR_ID
6          0, // SPIRV_INITIAL_BOUND
7          0)); // SCHEMA_BOUND

```

module. In turn, there are three main types of scope data types in SPIR-V: a) modules; b) functions; and c) blocks (basic block of instructions).

Module Scope. The outermost scope for a SPIR-V is a **module**, and it holds the global properties of the kernel (e.g., header, variables, addressing modes, capabilities, etc.) that affect the whole binary. The instructions allowed in this scope are the header declaration, the capabilities for the module (e.g., a compute **kernel**, a shader, enabling fp64, etc), decorators (e.g., to specify memory alignment of variables), types and composite declaration and function declarations. The module scope also handles the declaration of the device's local memory (e.g., shared memory on a GPU device).

Listing 2 shows an example of how to build the header for a new SPIR-V module using the Beehive SPIR-V Toolkit API. A SPIR-V module receives as a parameter, a SPIR-V header object, which is also a Java type provided by the SPIR-V Library. The SPIR-V header arguments specified in the constructor of the object correspond to the parameters specified in the SPIR-V standard, exactly in the same order. Thus, this makes it easier to follow the SPEC along with the API definition of the Beehive SPIR-V Toolkit.

Function Scope. Regarding the function scope, it handles the declaration of the parameters passed and declared at the function level and the function body. The instructions allowed in this scope are declaration SPIRVOpFunction and its parameters through the instruction SPIRVOpFunctionParameter.

Block Scope. Finally, the block scope handles the rest of the instructions of a basic block. In this paper, a basic block corresponds to a sequence of instructions that are executed one after the other and it does not contain control flow divergence. As soon as there is a new basic block, a new SPIR-V label must be instantiated, and therefore, the new instruction will be enclosed in a new block scope.

3.3.1 Enabling Function Composition. To enable the function composition of SPIR-V instructions, we designed a Java common interface (called SPIRVInstScope) for all types of SPIR-V scopes within the library. This interface contains abstract methods for adding new instructions to a SPIR-V module and registering new identifiers (IDs) in the SPIR-V module.

Listing 3. Java Instruction Scope Interface to allow instruction composition.

```

1  public interface SPIRVInstScope {
2      // Add a new instruction
3      SPIRVInstScope add(SPIRVInstruction var1);
4      // Generate a new ID for the next instruction
5      SPIRVId getNextId();
6  }

```

Listing 4. Example of obtaining new a new ID for the next instruction within a SPIR-V module.

```

1  SPIRVId idLabel = module.getNextId();
2  SPIRVOpLabel newLabel = new SPIRVOpLabel(idLabel);
3  functionScope.add(newLabel);

```

Listing 3 shows a simplified view of the SPIRVInstScope Java interface that the Beehive SPIR-V Toolkit library provides. There are two main methods for this interface: one for adding instructions (add), and another one for registering new IDs within the SPIR-V module that is being built (using the getNextId method).

Handling SSA Variables. In the Beehive SPIR-V Toolkit API, each instruction and kind have a unique SPIR-V identifier (ID). The reason is that SPIR-V binaries follow the SSA (Static Single Assignment) representation [29], in which each SPIR-V construct is assigned once.

Thus, to facilitate unique ID management, all new IDs are requested at the module level, which has the whole view of all instructions, operands, and kinds being used. Every time a new SPIR-V instruction is instantiated (a new Java object), the SPIR-V ID returned by the module is assigned to it (the new instruction). Additionally, since instructions are added within one of the three aforementioned scopes (module, function or block scopes), the moment that an instruction is added to the list of instructions that belong to the current scope, it also copies the ID into the module scope. In this way, the Beehive SPIR-V Toolkit library can perform checks on whether an ID is valid to be in the block that is being added.

Listing 4 shows an example of how to obtain a new ID for the next instruction within a function scope level. In this case, we want to add a new SPIR-V label instruction. Thus, we request, at the module level, a unique ID. Recall that new IDs are requested from the module level, which has the global view of all IDs being used. This ID is then used to instantiate a new label instruction.

The add, and getNextId methods in combination with all the instructions and operands generated by the TSE component, enable developers to compose SPIR-V modules and store them in binary format from Java.

3.3.2 Dis/Assembler. As we introduced in Section 3.2.2, the TSE component also generates two Java packages with all the logic regarding the assembler and disassembler. These two packages mainly contain the Proxy (or mapper) Java classes to transform from a) a text file that describes a SPIR-V module to a SPIR-V binary (assembler); and b) from a SPIR-V binary to a text file. Note that, if developers want to generate new SPIR-V modules in an instruction-by-instruction manner, the use of these mapper classes is not needed. However, these mapper Java classes are required in order to assemble or disassemble a SPIR-V binary code (from SPIR-V text to binary and vice-versa).

In order to generate the correct SPIR-V binary file from a text file that describes a SPIR-V module, the mapper Java classes contain helper methods. These methods are used to transform each token (i.e., instruction from the input file) to a SPIR-V instruction class by composing instructions of modules, functions and blocks (as explained in the previous section). Additionally, the Java mapper classes for the assembler keep a mapping between all unique IDs from the input text file in order to be used by the instructions that require these IDs as operands. For example, when a new variable is declared, the mapper Java class creates a mapping between the assigned string name and a new `SPIRVId` object that is created, and it is stored in an internal hash table. Then, when the requested ID is passed as an argument to any SPIR-V operation, the mapper recovers the `SPIRVId` by performing a lookup in the hash table.

Note that the input text file must represent a valid SPIR-V code. This is similar to the `spirv-as` LLVM utility command from the Khronos SPIR-V utilities [16], in which the SPIR-V text file can be manipulated and tested before integrating optimizations in the compiler pipelines of the optimizing compilers, thereby facilitating fast debugging. Since SPIR-V is an IR in binary format, it is not easy to try out new optimizations and reordering operations without changing and adapting the compiler infrastructure. However, the proposed approach via mappers allows developers of JIT compilers to test new optimizations, measure performance and then integrate the changes in the compilation pipeline while using the Java software ecosystem.

The disassembler mapper classes work in a similar way to the assembler mapper classes but in the reverse order. The mapping classes contain the logic to transform SPIR-V opcodes (integer values) into text format.

Listing 5 shows an example of how to disassemble a SPIR-V binary file into text. Line 1 creates a Java object for selecting the options for the disassembler. The constructor accepts some utilities that can be configured, such as syntax highlighting, indentation, turning off the header, etc. Lines 8-11 create the SPIR-V disassembler object and line 12 invokes the runner. When lines 8-11 are executed, it prints the SPIR-V text file that corresponds to the input binary in the selected output (e.g., standard output).

Listing 5. Code snippet that shows how to invoke the disassembler.

```

1  SPIRVDisassemblerOptions opt =
2      new SPIRVDisassemblerOptions(
3          true, // Syntax Highlight
4          true, // Show Inline Names
5          false, // Turn Off Indentation
6          true, // Should Group
7          false); // No Header
8  SPIRVTool spirvTool = new Disassembler(
9      reader, // input binary file
10     System.out, // output
11     opt); // options
12  spirvTool.run();

```

3.3.3 Validation Rules. The Beehive SPIR-V Toolkit also provides some validation methods for the generated SPIR-V binary modules. The validation rules provided are as follows: the SPIR-V module must have at least one function declared, at least one capability declared, the memory model must be defined and it should contain at least one entry point. Furthermore, the generated SPIR-V binary modules can be validated using the Khronos SPIR-V Tools, such as `spirv-val`. The validation of the generated modules is important to ensure the fidelity of the code generator. Our validator is complementary to the SPIR-V Khronos validator. For instance, our validator also checks for instruction capabilities and dependencies before storing the final SPIR-V binary. The dependencies are retrieved from the JSON file that describes the SPIR-V grammar. Through this validation, we were able to detect circular dependencies in the latest SPIR-V standard (1.6).

3.3.4 SPIR-V Client Utility. The Beehive SPIR-V Toolkit also contains a standalone module that includes a client application for assembling and disassembling SPIR-V code. The client application acts as a command line utility that invokes the assembler and disassembler components from the library.

4 Use case: Integration into TornadoVM

Once we developed the library generator, we integrated it into the TornadoVM JIT compiler. TornadoVM [7, 13] is a Java parallel programming framework that transforms, at runtime, Java bytecode into OpenCL and PTX code. To do so, TornadoVM accelerates a subset of Java by offloading code to be executed on CPUs, GPUs and FPGAs.

We extended the TornadoVM project with a new backend for dispatching SPIR-V code through the Intel Level Zero [27] API (a new low-level API developed by Intel to manage heterogeneous devices). Although the Level Zero API is not attached to any specific hardware accelerator, it is currently only available for Intel-integrated GPUs. Thus, we prototype

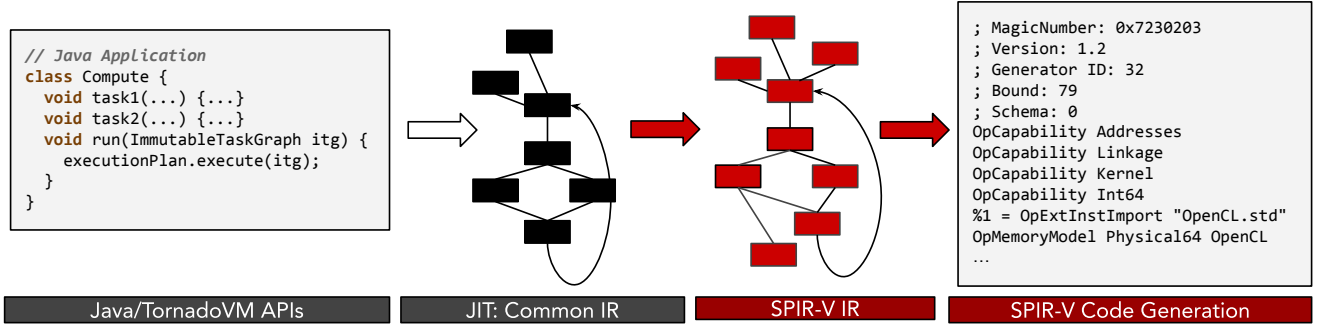


Figure 4. SPIR-V JIT Compilation Process in TornadoVM.

the SPIR-V backend for integrated GPUs with support for Level Zero and SPIR-V. In the future, we plan to also integrate the SPIR-V backend to be also dispatched through the OpenCL runtime.

SPIR-V Compilation Process for TornadoVM. Since the TornadoVM JIT compiler extends the Graal JIT compiler [10], and it is fully implemented in Java, we imported our library as a new dependency for the new SPIR-V backend in TornadoVM JIT compiler. We implemented the new backend following the templates for the existing OpenCL and CUDA PTX backends.

The compilation workflow for SPIR-V is shown in Figure 4. The Figure shows four blocks: the first block on the left shows a Java code snippet that represents a program using the TornadoVM parallel APIs for GPU and FPGA programming. When the code is first executed, the TornadoVM runtime system invokes the JIT compiler for each Java method to be compiled to the target backend (e.g., SPIR-V).

TornadoVM lowers the code from the Java bytecode to the target backend in different compilation phases or compilation tiers. The first tier is called the *sketcher*, and it contains a common representation for all backends (e.g., OpenCL and SPIR-V). During the *sketcher*, the TornadoVM JIT compiler applies common high-level optimizations, such as constant folding, evaluation of expressions, etc. We did not extend this phase to adopt the new SPIR-V Backend.

From the *sketcher*, the TornadoVM JIT compiler starts specializing the IR per backend. In our case, we extended with a new set of compiler optimizations and lowering phases to transform the high-tier compiler IR of the program into the SPIR-V code. To do so, we followed the compilation pipeline of TornadoVM with three new compilation tiers, named high-tier, mid-tier and low-tier, and added, in total, 58 new compilation phases that are specialized for the SPIR-V backend. These compilation phases include optimizations for performing fast math operations, vector operations, and inlining, among many others.

Once the TornadoVM JIT compiler optimizes the IR, we generate the corresponding SPIR-V code using our Beehive SPIR-V Toolkit library. To generate SPIR-V, we traverse the

final IR and build a new list of *lowerable* Java objects (Java types provided by GraalVM to generate code) with the specific instructions to generate SPIR-V code. The final process is to traverse the final list and invoke the generate method for each object in the list. The result of this process is a SPIR-V module that can be dispatched through the Level Zero API included in TornadoVM.

Advantages of the SPIR-V Backend in TornadoVM.

From our experience porting the SPIR-V backend in TornadoVM, we see that the structure is much simpler than the OpenCL C backend. This is because the SPIR-V backend generates code at the binary level, while the OpenCL C code reconstructs source code from the low-level IR of the TornadoVM JIT compiler.

For instance, the OpenCL C backend contains a lot of control to cover many corner cases when generating structured control flow from the unstructured control flow of the TornadoVM IR, and therefore, the Graal IR [24]. The SPIR-V backend of TornadoVM dramatically simplifies this process by allowing conditional and unconditional jumps to the specific compiler basic blocks.

5 Evaluation

We evaluated the library against the existing OpenCL Backend in TornadoVM using same device Intel-integrated GPU contained in an Intel i9-10885H CPU. The Intel driver used was 21.38.21026, which is the same for running OpenCL and SPIR-V applications. The OS used was Fedora Linux 34 with the Linux Kernel 5.16.18-100. TornadoVM contains a script for benchmarking, and it includes applications from many different domains such as machine learning, Fintech, linear algebra and physics [13].

Performance of the JIT Compiler and Code Generation.

Figure 5 shows the execution time for code generation (the total time that takes to build either a SPIR-V module or an OpenCL C kernel from the last compiler phase), and the driver compilation, which represents the total time that the GPU driver takes to compile to the final GPU binary. The SPIR-V backend employs the Beehive SPIR-V Toolkit library,




Figure 5. Performance evaluation between the code generator of TornadoVM for OpenCL and the SPIR-V. The SPIR-V backend uses the Beehive SPIR-V Toolkit library. The stack plot displays two types of metrics: a) the driver compilation, which involves the JIT compilation once the code is generated; and b) the code generation, which contains the total time to generate the code by using the proposed library. The total time is shown in nanoseconds. The lower, the better.

while the OpenCL C backend of TornadoVM uses an ad-hoc library of the project (with no decoupling).

In general, the SPIR-V backend performs slower than the OpenCL when generating the code (black section of the stacks in Figure 5). This is expected since the Beehive SPIR-V Toolkit API creates a significant number of Java objects in order to build SPIR-V modules. This is because each identified and each instruction is a new Java object. However, once the code has been generated, the driver compilation time decreases by up to 3.9x compared to OpenCL.

Although the function composition and modularity of the SPIR-V library come at a performance cost, the most efficient driver implementation of SPIR-V results in an end-to-end compilation performance increase of 2.72x (on average) compared to OpenCL. End-to-end performance speedups, compared to OpenCL, range from 2.2x% (*convolveImageArray* - the third group of bars), to 3x (*hilbert* - the fifth group of bars starting from the right-hand side).

Performance of the SPIR-V Backend. We also evaluated the total time that takes each benchmark to run for each backend. Figure 6 shows the speedup of each application using our new SPIR-V backend over the OpenCL C backend that was already implemented in the TornadoVM project. Thus, the higher, the better. We executed all benchmarks on the same Intel HD graphics for both backends.

To measure the performance, we executed the benchmark script that the TornadoVM provides, which performs a set of warm-up iterations and then provides the median time for all executions. Since we provide the median, the JIT compilation time, which happens in the first iteration, is excluded from these measurements.

As a disclaimer, note that the SPIR-V backend is dispatched through the Intel Level Zero API, while the OpenCL C backend is dispatched through the Intel OpenCL driver for the same GPU (Intel integrated GPU). This means that the block of threads to be deployed on the GPU might differ between the OpenCL and the Level Zero thread dispatchers.




Figure 6. Speedup of the our implementation of the SPIR-V backend for TornadoVM using the Beehive SPIR-V Toolkit library API over the existing OpenCL C backend on the integrated Intel GPU.

We see that, in general, the SPIR-V backend varies from 2% slowdown compared to the OpenCL C backend execution (for the *saxpy*, *addImage* and *dft* benchmarks) to 15% speedup (for the *juliaSet* benchmark).

In particular, the SPIR-V performs better than OpenCL for the *convolveArray* and *blurFilter* benchmarks, 1.47x and 1.52x respectively. In contrast, the *blackscholes* benchmark performs 50% slowdown compared to the OpenCL C.

Even though we are executing the benchmarks on the same device, we see different performances. One of the reasons is that the thread block we select for the SPIR-V backend is different compared to the one that TornadoVM selects for the OpenCL C backend. This clearly can influence performance [34].

6 Related Work

This section presents the most relevant work to the Beehive SPIR-V Toolkit concerning how it is implemented through its TSE and used as a library for SPIR-V code generation. To the best of our knowledge, this is the first Java library designed for SPIR-V code generation.

Template System Engine (TSE). SPIR-V Tools [16] uses a similar approach for mapping the code between SPIR-V text and SPIR-V binary. However, our approach also provides

a functional programming interface to dynamically build SPIR-V modules embedded in the TSE. We believe this API offers a cleaner and more concise way to compose SPIR-V applications.

LLVM *TableGen* [25] also uses a similar approach to our TSE. In LLVM, *TableGen* is a tool provided by the LLVM compiler infrastructure that can be used to generate code using the `llvm-tblgen` utility command. LLVM *TableGen* is a more powerful tool than our TSE because it can be used not only to generate operands and instructions (as the Beehive SPIR-V Toolkit does), but it can also generate instruction scheduling information and specific rules for each instruction. However, our TSE can be extended to accommodate this kind of functionality by providing the corresponding templates in Java.

J. Haavisto [17] presented an approach to generate SPIR-V binary modules using programs written in APL [19] (a non-imperative and array centric programming language) as a modelling language. There are two differentiation points with our work: i) our work exposes a generic and composable API, rather than the templates themselves generate the whole SPIR-V code, and ii) the TSE is generic to support other code instruction sets and IRs such as OpenCL or CUDA PTX.

SPIR-V ASM/DASM libraries. The Multi-Level IR Compiler Framework (MLIR) [23] contains a custom implementation of a SPIR-V library to generate SPIR-V binary modules from the MLIR-specific IR. The custom SPIR-V library generates a dialect of SPIR-V (LLVM IR with SPIR-V intrinsics)⁵. MLIR provides utilities to lower the SPIR-V dialect to standard SPIR-V modules. The SPIR-V dialect generated by the MLIR is designed to perform specific optimizations [35]. In contrast, the prime focus of the proposed TSE is to generate standard SPIR-V. However, the set of templates in Beehive SPIR-V Toolkit can be extended to include the emission of other dialects that can also interact with the MLIR GPU dialects for SPIR-V. To do so, we would need to extend the write methods of each SPIR-V instruction to emit the equivalent SPIR-V dialect instructions.

ViennaCL++ [6] can generate SPIR-V binary modules by invoking the LLVM SPIR-V compiler for input OpenCL kernels. The Beehive SPIR-V Toolkit library can generate SPIR-V binary modules via a composable and functional Java API. DCompute [36] is a SPIR-V code generator library that reuses the LLVM-based D Compiler (LDC) [25] to statically compile programs written in the D language to SPIR-V. The Beehive SPIR-V Toolkit, although it provides a Java API, it is language agnostic, making it suitable for reusing with other JVM languages such as Scala, R [32], JavaScript and NodeJS [37], Python or Ruby [26]. Furthermore, DCompute can compile high-level language constructs, such as lambda expressions. Unlikely, the proposed library is designed to operate at a lower-level, and it is intended for compiler engineers who

need to generate and debug SPIR-V code from managed run-time programming languages.

The vast majority of the SPIR-V compilers use LLVM tools [16] to generate, analyze and validate the generated code. Several compilers that belong to this category are the Intel SYCL compiler [5], HIPCL [4], ComputeCpp [8, 28], HipSYCL [1], triSYCL [20] and Intel oneAPI [9]. Our approach leverages the SPIR-V code generation as a high-level library for Java and JVM-based programming languages.

7 Conclusions

SPIR-V is an IR in a binary format to express parallel computations and graphics for execution on heterogeneous hardware. SPIR-V tools that assemble and disassemble from/to parallel programming and graphics models such as OpenCL and Vulkan already exist. However, they are primarily designed for C/C++ and use the LLVM compiler infrastructure and software ecosystem.

This paper presents the Beehive SPIR-V Toolkit, a Java programming framework that enables the generation of SPIR-V kernels from programming languages that can run on top of the JVM (e.g., Java, Truffle Python, Truffle Ruby or JavaScript). While the technique to automatically generate programs from JSON files is not new, it has been exploited to generate a complete API from a specification, allowing the adoption of new releases of the standards with minimal effort quickly. The Beehive SPIR-V Toolkit is an open-source project that generates a functional and composable API for building SPIR-V modules. This paper also described the overall architecture of the proposed library and presented a template system engine that can automatically generate a programming library from a well-specified grammar written in JSON. Furthermore, the Beehive SPIR-V Toolkit contains a client utility that can be used as a standalone tool for assembling, disassembling and validating a SPIR-V binary code.

We showcase the integration of our library into the TornadoVM project, by providing a new backend to compile Java to SPIR-V. Our experiments show that our library performs up to 3x compared to the existing OpenCL C code generator in TornadoVM, and the generated SPIR-V code performs up to 1.52x faster than the OpenCL C backend when running on an Intel integrated GPU. In future work, we plan to extend validation rules. Additionally, we want to expose an API for checking and validating different sections of the generated SPIR-V binary modules.

Acknowledgments

This work is partially funded by grants from Intel Corporation and the European Union's Horizon 2020 programme under grant agreement No 957286 (ELEGANT). Additionally, it is funded by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee for grant numbers 10048318 (AERO), 10048316 (INCORE), 10039809 (ENCRYPT) and 10039107 (TANGO)

⁵<https://groups.google.com/g/llvm-dev/c/n0vU7iHNis>

References

- [1] Aksel Alpay and Vincent Heuveline. 2020. SYCL beyond OpenCL: The Architecture, Current State and Future Direction of HipSYCL. In *Proceedings of the International Workshop on OpenCL* (Munich, Germany) (IWOCCL '20). Association for Computing Machinery, New York, NY, USA, Article 8, 1 pages. <https://doi.org/10.1145/3388333.3388658>
- [2] Apache. Last Access: May 2023. Apache FreeMarker Template Engine. <https://freemarker.apache.org/>
- [3] Ben Ashbaugh. Last Access: May 2022. OpenCL Extended Instruction Set Specification. https://registry.khronos.org/SPIR-V/specs/unified1/OpenCL.ExtendedInstructionSet.100.html#_a_id_math_a_math_extended_instructions
- [4] Michal Babej and Pekka Jääskeläinen. 2020. HIPCL: Tool for Porting CUDA Applications to Advanced OpenCL Platforms Through HIP. In *Proceedings of the International Workshop on OpenCL* (Munich, Germany) (IWOCCL '20). Association for Computing Machinery, New York, NY, USA, Article 18, 3 pages. <https://doi.org/10.1145/3388333.3388641>
- [5] Alexey Bader, James Brodman, and Michael Kinsner. 2019. A SYCL Compiler and Runtime Architecture. In *Proceedings of the International Workshop on OpenCL* (Boston, MA, USA) (IWOCCL '19). Association for Computing Machinery, New York, NY, USA, Article 16, 1 pages. <https://doi.org/10.1145/3318170.3318194>
- [6] Tai-Liang Chen, Shih-Huan Chien, and Jenq-Kuen Lee. 2018. ViennaCL++: Enable TensorFlow/Eigen via ViennaCL with OpenCL C++ Flow. In *Proceedings of the International Workshop on OpenCL* (Oxford, United Kingdom) (IWOCCL '18). Association for Computing Machinery, New York, NY, USA, Article 28, 2 pages. <https://doi.org/10.1145/3204919.3207894>
- [7] James Clarkson, Juan Fumero, Michail Papadimitriou, Foivos S. Zakkak, Maria Xekalaki, Christos Kotselidis, and Mikel Luján. 2018. Exploiting High-Performance Heterogeneous Hardware for Java Programs Using Graal. In *Proceedings of the 15th International Conference on Managed Languages & Runtimes* (Linz, Austria) (ManLang '18). Association for Computing Machinery, New York, NY, USA, Article 4, 13 pages. <https://doi.org/10.1145/3237009.3237016>
- [8] Codeplay. Last Access: May 2023. ComputeCpp. <https://developer.codeplay.com/products/computecpp/ce/home>
- [9] Intel Corporation. Last Access: May 2023. oneAPI. <https://www.intel.com/content/www/us/en/developer/tools/oneapi/overview.html>
- [10] Gilles Duboscq, Lukas Stadler, Thomas Würthinger, Doug Simon, Christian Wimmer, and Hanspeter Mössenböck. 2013. Graal IR: An extensible declarative intermediate representation. In *Proceedings of the Asia-Pacific Programming Languages and Compilers Workshop*.
- [11] FasterXML. Last Access: May 2023. Jackson Annotation Framework. <https://github.com/FasterXML/jackson>
- [12] Jean-Marie Favre. 2004. Towards a basic theory to model model driven engineering. In *3rd workshop in software model engineering, wisme*. Citeseer, 262–271.
- [13] Fumero, Juan and Papadimitriou, Michail and Zakkak, Foivos S. and Xekalaki, Maria and Clarkson, James and Kotselidis, Christos. 2019. Dynamic Application Reconfiguration on Heterogeneous Hardware. In *Proceedings of the 15th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments* (Providence, RI, USA) (VEE 2019). Association for Computing Machinery, New York, NY, USA, 165–178. <https://doi.org/10.1145/3313808.3313819>
- [14] Khronos Group. Last Access: May 2023. Khronos Group. <https://www.khronos.org/>
- [15] Khronos Group. Last Access: May 2023. SPIR-V Specifications. <https://www.khronos.org/spir/>
- [16] Khronos Group. Last Access: May 2023. SPIRV Tools. <https://github.com/KhronosGroup/SPIRV-Tools>
- [17] Juuso Haavisto. 2020. *Leveraging APL and SPIR-V languages to write network functions to be deployed on Vulkan compatible GPUs*. Master's thesis. Université de Lorraine. <https://hal.inria.fr/hal-03155647>
- [18] The Khronos Group Inc. Last access: May 2023. Vulkan 1.3.250 - A Specification. <https://registry.khronos.org/vulkan/specs/1.3/html/vkspec.html>
- [19] Kenneth E. Iverson. 1962. *A Programming Language*. John Wiley & Sons, Inc.
- [20] Ronan Keryell and Lin-Ya Yu. 2018. Early Experiments Using SYCL Single-Source Modern C++ on Xilinx FPGA: Extended Abstract of Technical Presentation. In *Proceedings of the International Workshop on OpenCL* (Oxford, United Kingdom) (IWOCCL '18). Association for Computing Machinery, New York, NY, USA, Article 18, 8 pages. <https://doi.org/10.1145/3204919.3204937>
- [21] Chris Lattner. 2008. LLVM and Clang: Next generation compiler technology. In *The BSD conference*, Vol. 5.
- [22] Chris Lattner and Vikram Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-Directed and Runtime Optimization* (Palo Alto, California) (CGO'04). IEEE Computer Society, USA, 75.
- [23] Chris Lattner, Jacques A. Pienaar, Mehdi Amini, Uday Bondhugula, River Riddle, Albert Cohen, Tatiana Shpeisman, Andy Davis, Nicolas Vasilache, and Oleksandr Zinenko. 2020. MLIR: A Compiler Infrastructure for the End of Moore's Law. *CoRR* abs/2002.11054 (2020). [arXiv:2002.11054](https://arxiv.org/abs/2002.11054) <https://arxiv.org/abs/2002.11054>
- [24] David Leopoldseder, Lukas Stadler, Christian Wimmer, and Hanspeter Mössenböck. 2015. Java-to-JavaScript Translation via Structured Control Flow Reconstruction of Compiler IR. *SIGPLAN Not.* 51, 2 (oct 2015), 91–103. <https://doi.org/10.1145/2936313.2816715>
- [25] LLVM.org. Last Access: May 2023. The LLVM Compiler Infrastructure. <https://llvm.org/>
- [26] Kevin Menard, Chris Seaton, and Benoit Daloze. 2018. Specializing Ropes for Ruby. In *Proceedings of the 15th International Conference on Managed Languages and Runtimes* (Linz, Austria) (ManLang '18). Association for Computing Machinery, New York, NY, USA, Article 10, 7 pages. <https://doi.org/10.1145/3237009.3237026>
- [27] Intel oneAPI. Last Access: May 2023. Level Zero SPEC. <https://spec.oneapi.io/level-zero/latest/index.html>
- [28] Ruymán Reyes and Victor Lomüller. 2014. SYCL: Single-source C++ accelerator programming. In *Advances in Parallel Computing. Volume 27: Parallel Computing: On the Road to Exascale*. 673–682. <https://doi.org/10.3233/978-1-61499-621-7-673>
- [29] B. K. Rosen, M. N. Wegman, and F. K. Zadeck. 1988. Global Value Numbers and Redundant Computations. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (POPL '88). Association for Computing Machinery, New York, NY, USA, 12–27. <https://doi.org/10.1145/73560.73562>
- [30] Mark Segal and Kurt Akeley. May 2022. The OpenGL Graphics System: A Specification (Version 4.6). <https://registry.khronos.org/OpenGL/specs/gl/glspec46.core.pdf>
- [31] SPIR-V-Khronos. Last Access: August 2023. SPIR-V Headers. <https://github.com/KhronosGroup/SPIRV-Headers>
- [32] Lukas Stadler, Adam Welc, Christian Humer, and Mick Jordan. 2016. Optimizing R Language Execution via Aggressive Speculation. In *Proceedings of the 12th Symposium on Dynamic Languages* (Amsterdam, Netherlands) (DLS 2016). Association for Computing Machinery, New York, NY, USA, 84–95. <https://doi.org/10.1145/2989225.2989236>
- [33] John E Stone, David Gohara, and Guochun Shi. 2010. OpenCL: A parallel programming standard for heterogeneous computing systems. *Computing in science & engineering* 12, 3 (2010), 66.
- [34] Yuri Torres, Arturo Gonzalez-Escribano, and Diego R. Llanos. 2013. UBench: Exposing the Impact of CUDA Block Geometry in Terms of Performance. *J. Supercomput.* 65, 3 (sep 2013), 1150–1163. <https://doi.org/10.1007/s11227-013-0921-z>

- [35] Nicolas Vasilache, Oleksandr Zinenko, Aart J. C. Bik, Mahesh Ravishankar, Thomas Raoux, Alexander Belyaev, Matthias Springer, Tobias Gysi, Diego Caballero, Stephan Herhut, Stella Laurenzo, and Albert Cohen. 2022. Composable and Modular Code Generation in MLIR: A Structured and Retargetable Approach to Tensor Compiler Construction. *ArXiv abs/2202.03293* (2022).
- [36] Nicholas Wilson. 2018. DCompute: Compiling D to SPIR-V for Seamless Integration with OpenCL. In *Proceedings of the International Workshop on OpenCL* (Oxford, United Kingdom) (*IWOCL '18*). Association for Computing Machinery, New York, NY, USA, Article 3, 3 pages. <https://doi.org/10.1145/3204919.3204922>
- [37] Thomas Würthinger, Christian Wimmer, Christian Humer, Andreas Wöß, Lukas Stadler, Chris Seaton, Gilles Duboscq, Doug Simon, and Matthias Grimmer. 2017. Practical Partial Evaluation for High-Performance Dynamic Language Runtimes. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Barcelona, Spain) (*PLDI 2017*). Association for Computing Machinery, New York, NY, USA, 662–676. <https://doi.org/10.1145/3062341.3062381>

Received 2023-07-23; accepted 2023-08-28